NOTES ON SEVERAL COMPLEX VARIABLES

J. L. Taylor

Department of Mathematics
University of Utah
July 27, 1994
Revised June 9, 1997

Notes from a 1993–94 graduate course
Revised for a 1996-97 graduate course

TABLE OF CONTENTS

1.	Holomorphic Functions	p. 3
2.	Local Theory	p. 10
3.	A Little Homological Algebra	p. 15
4.	Local Theory of Varieties	p. 24
5.	The Nullstellensatz	p. 30
6.	Dimension	p. 46
7.	Completion of Local Rings	p. 54
8.	Sheaves	p. 63
9.	Sheaf Cohomology	p. 71
10.	Coherent Algebraic Sheaves	p. 87
11.	Dolbeault Cohomology	p. 99
12.	Coherent Analytic Sheaves	p. 106
13.	Projective Varieties	p. 116
14.	Algebraic vs. Analytic Sheaves – Serre's Theorems	p. 124
15.	Stein Spaces	p. 136
16.	Fréchet Sheaves – Cartan's Theorems	p. 149
17.	The Borel-Weil-Bott Theorem	p. 166

1. Holomorphic Functions

There are a number of possible ways to define what it means for a function defined on a domain in $\mathbb{C}^n$ to be holomorphic. One could simply insist that the function be holomorphic in each variable separately. Or one could insist the function be continuous (as a function of several variables) in addition to being holomorphic in each variable separately. The a priori strongest condition would be to insist that a holomorphic function have a convergent expansion as a multi-variable power series in a neighborhood of each point of the domain. The main object of this chapter is to show that these possible definitions are all equivalent.

In what follows, $\Delta(a,r)$ will denote the polydisc of radius $r = (r_1, r_2, ..., r_n)$ about $a = (a_1, a_2, ..., a_n)$:

$$\Delta(a,r) = \{z = (z_1, z_2, ..., z_n) : |z_j - a_j| < r_j \quad j = 1, 2, ..., n\}$$

and $\bar{\Delta}(a,r)$ will denote the corresponding closed polydisc. Note that $\Delta(a,r)$ is just the Cartesian product of the open discs $\Delta(a_i, r_i) \subset \mathbb{C}$ and $\bar{\Delta}(a,r)$ is the Cartesian product of the closed discs $\bar{\Delta}(a_i, r_i)$.

1.1 Proposition (Cauchy Integral Formula). *If f is a function which is continuous in a neighborhood U of the closed polydisc $\bar{\Delta}(a,r)$ and holomorphic in each variable z_i at each point of U, then*

$$f(z) = \left(\frac{1}{2\pi i}\right)^n \int_{|\zeta_n - a_n| = r_n} \cdots \int_{|\zeta_1 - a_1| = r_1} \frac{f(\zeta_1, ..., \zeta_n) d\zeta_1 ... d\zeta_n}{(\zeta_1 - z_1)...(\zeta_n - z_n)}$$

for any $z \in \Delta(a,r)$.

This follows immediately from repeated application of the one variable Cauchy Theorem.

1.2 Proposition (Osgood's Lemma). *If a function is continuous in an open set $U \subset \mathbb{C}^n$ and is holomorphic in each variable at each point of U then at each point $a \in U$ there is a power series of the form*

$$\sum_{i_1, ..., i_n = 0}^{\infty} c_{i_1 ... i_n} (z_1 - a_1)^{i_1} ... (z_n - a_n)^{i_n}$$

which converges uniformly to f on every compact polydisc centered at a and contained in U.

Proof. Let $\bar{\Delta}(a,s)$ be a compact polydisc centered at a and contained in U. Since $\bar{\Delta}(a,s)$ is compact in U we may choose numbers $r_i > s_i$ such that the compact polydisc $\bar{\Delta}(a,r)$ with polyradius $r = (r_1, \ldots, r_n)$ is also contained in U. We substitute into the integrand of 1.1 the series expansion:

$$\frac{1}{(\zeta_1 - z_1)...(\zeta_n - z_n)} = \sum_{i_1, ..., i_n = 0}^{\infty} \frac{(z_1 - a_1)^{i_1} ... (z_n - a_n)^{i_n}}{(\zeta_1 - a_1)^{i_1 + 1} ... (\zeta_n - a_n)^{i_n + 1}}$$

This series converges uniformly in ζ and z for $|z_i - a_i| < s_i < r_i = |\zeta_i - a_i|$. Thus, the series in the integrand of 1.1 can be integrated termwise and the resulting series in z is uniformly convergent for $z \in \bar{\Delta}(a, s)$. The result is a series expansion of the form

$$f(z) = \sum_{i_1,...,i_n=0}^{\infty} c_{i_1...i_n}(z_1 - a_1)^{i_1}...(z_n - a_n)^{i_n}$$

uniformly convergent on $\bar{\Delta}(a, s)$, with

$$c_{i_1...i_n} = \left(\frac{1}{2\pi i}\right)^n \int_{|\zeta_n - z_n| = r_n} \cdots \int_{|\zeta_1 - z_1| = r_1} \frac{f(\zeta_1, ..., \zeta_n) d\zeta_1 ... d\zeta_n}{(\zeta_1 - a_1)^{i_1+1}...(\zeta_n - a_n)^{i_n+1}}$$

This completes the proof.

A function f defined on an open set $U \subset \mathbb{C}^n$ is called holomorphic at a if it has a power series expansion as in 1.2, convergent to f in some open polydisc centered at a. If f is holomorphic at each point of U then it is called holomorophic on U. Thus, Osgood's Lemma says that a continuous function on U which is holomorphic in each variable separately at each point of U is holomorphic on U. We will later prove that the continuity hypothesis is redundant.

1.3 Proposition (Cauchy's inequalities). *If f is a holomorphic function in a neighborhood of the closed polydisc $\bar{\Delta}(a, r)$ and $|f|$ is bounded by M in this polydisc, then for each multi-index $(i_1, ..., i_n)$*

$$\left| \frac{\partial^{i_1 + ... + i_n}}{\partial {z_1}^{i_1} ... \partial {z_n}^{i_n}} f(a) \right| \leq M(i_1!)...(i_n!){r_1}^{-i_1}...{r_n}^{-i_n}$$

Proof. If f is expressed as a power series convergent on $\bar{\Delta}(a, r)$ as in 1.2, then repeated differentiation yields

$$\frac{\partial^{i_1 + ... + i_n}}{\partial {z_1}^{i_1} ... \partial {z_n}^{i_n}} f(a) = (i_1!)...(i_n!)c_{i_1...i_n}$$

where, according to the proof of 1.2,

$$c_{i_1...i_n} = \left(\frac{1}{2\pi i}\right)^n \int_{|\zeta_n - z_n| = r_n} \cdots \int_{|\zeta_1 - z_1| = r_1} \frac{f(\zeta_1, ..., \zeta_n) d\zeta_1 ... d\zeta_n}{(\zeta_1 - a_1)^{i_1+1}...(\zeta_n - a_n)^{i_n+1}}$$

From the obvious bounds on the integrand of this integral it follows that

$$|c_{i_1...i_n}| \leq M{r_1}^{-i_1}...{r_n}^{-i_n}$$

and the Cauchy inequalities follow from this.

In what follows, the notation $|U|$ will be used to denote the volume of a set $U \subset \mathbb{C}^n$.

1.4 Theorem (Jensen's inequality). *If f is holomorphic in a neighborhood of $\bar\Delta = \bar\Delta(a,r)$ then*

$$\log|f(a)| \le \frac{1}{|\Delta|}\int_\Delta \log|f(z)|dV(z)$$

Proof. We recall the single variable version of Jensen's inequality:

$$\log|g(a)| \le \frac{1}{2\pi}\int_0^{2\pi} \log|g(a+\rho e^{i\theta})|d\theta$$

and apply it to f considered as a function of only z_1 with the other variables fixed at $a_2, ..., a_n$. This yields:

$$\log|f(a)| \le \frac{1}{2\pi}\int_0^{2\pi} \log|f(a_1+\rho_1 e^{i\theta_1}, a_2, ..., a_n)|d\theta_1$$

We now apply Jensen's inequality to the integrand of this expression, where f is considered as a function of z_2 with z_1 fixed at $a_a+\rho_1 e^{i\theta_1}$ and the remaining variables fixed at $a_3, ..., a_n$ to obtain:

$$\log|f(a)| \le \left(\frac{1}{2\pi}\right)^2 \int_0^{2\pi}\int_0^{2\pi} \log|f(a_1+\rho_1 e^{i\theta_1}, a_2+\rho_2 e^{i\theta_2}, a_3, ..., a_n)|d\theta_1 d\theta_2$$

Continuing in this way, we obtain:

$$\log|f(a)| \le \left(\frac{1}{2\pi}\right)^n \int_0^{2\pi} ... \int_0^{2\pi} \log|f(a_1+\rho_1 e^{i\theta_1}, a_2+\rho_2 e^{i\theta_2}, ..., a_n+\rho_n e^{i\theta_n})|d\theta_1 ... d\theta_n$$

Finally, we multiply both sides of this by $\rho_1 ... \rho_n$ and integrate with respect to $d\rho_1 ... d\rho_n$ over the set $\{\rho_i \le r_i; \quad i = 1, ..., n\}$ to obtain the inequality of the theorem.

In the next lemma we will write a point $z \in \mathbb{C}^n$ as $z = (z', z_n)$ with $z' \in \mathbb{C}^{n-1}$. Similarly we will write a polyradius r as $r = (r', r_n)$ so that a polydisc $\Delta(a,r)$ can be written as $\Delta(a', r') \times \Delta(a_n, r_n)$.

1.5 Lemma (Hartogs' lemma). *Let f be holomorphic in $\Delta(0,r)$ and let*

$$f(z) = \sum_k f_k(z') z_n^k$$

be the power series expansion of f in the variable z_n, where the f_k are holomorphic in $\Delta(0,r')$. If there is a number $c > r_n$ such that this series converges in $\bar\Delta(0,c)$ for each $z' \in \Delta(0,r')$ then it converges uniformly on any compact subset of $\Delta(0,r') \times \Delta(0,c)$. Thus, f extends to be holomorphic in this larger polydisc.

Proof.. Choose any point $a' \in \Delta(0,r')$, choose a closed polydisc $\bar\Delta(a', s') \subset \Delta(0,r')$ and choose some positive $b < r_n$. Then we may choose an upper bound $M > 1$ for f on the polydisc $\bar{\Delta}(a', s') \times \bar{\Delta}(0, b)$. It follows from Cauchy's inequalities that $|f_k(z')| \leq Mb^{-k}$ for all $z' \in \bar{\Delta}(a', s')$. Hence,

$$k^{-1} \log |f_k(z')| \leq k^{-1} \log M - \log b \leq \log M - \log b = M_0$$

for all $z' \in \bar{\Delta}(a', s')$ and all k. By the convergence of the above series at (z', c) we also have that $|f_k(z')|c^k$ converges to zero for each $z' \in \bar{\Delta}(a', s')$ and hence:

$$\limsup_k k^{-1} \log |f_k(z')| \leq -\log c$$

for all $z' \in \bar{\Delta}(a', s')$.

We have that the functions $k^{-1} \log |f_k(z')|$ are uniformly bounded and measurable in $\bar{\Delta}(a', s')$ and by Fatou's lemma

$$\begin{aligned} \limsup_k \int_{\Delta(a', s')} k^{-1} \log |f_k(z')| dV(z') &\leq \int_{\Delta(a', s')} \limsup_k k^{-1} \log |f_k(z')| dV(z') \\ &\leq -|\Delta(a', s')| \log c \end{aligned}$$

By replacing c by a slightly smaller number if neccessary we may conclude that there is a k_0 such that

$$\int_{\Delta(a', s')} k^{-1} \log |f_k(z')| dV(z') \leq -|\Delta(a', s')| \log c$$

for all $k \geq k_0$. Now by shrinking s' to a slightly smaller multiradius t' and choosing ϵ small enough we can arrange that

$$\Delta(a', t') \subset \Delta(w', t' + \epsilon) \subset \Delta(a', s')$$

for all $w' \in \Delta(a', \epsilon)$, from which we conclude that

$$\int_{\Delta(w', t'+\epsilon)} k^{-1} \log |f_k(z')| dV(z') \leq -|\Delta(a', t')| \log c + M_0 |\Delta(w', t' + \epsilon) - \Delta(w', t')|$$

Now if we choose a c_0 slightly smaller than c we may choose ϵ small enough that the right hand side of this inequality is less than $-|\Delta(w', t' + \epsilon)| \log c_0$. This yields

$$\int_{\Delta(w', t'+\epsilon)} k^{-1} \log |f_k(z')| dV(z') \leq -|\Delta(w', t' + \epsilon)| \log c_0$$

and this, in turn, through Jensen's inequality, yields

$$k^{-1} \log |f_k(w')| \leq -\log c_0$$

or

$$|f_k(w')| c_0^k \leq 1$$

for all $k \geq k_0$ and $w' \in \Delta(a', \epsilon)$. This implies that our series is uniformly convergent in $\Delta(a', \epsilon) \times \Delta(0, c_0)$ and since a' is arbitrary in $\Delta(0, r')$ and c_0 is an arbitrary positive number less than c, we have that our power series serves to extend f to a function which is continuous on $\Delta(0, r') \times \Delta(0, c)$ and certainly holomorphic in each variable. It follows from Osgood's lemma that this extension of f is, in fact, holomorphic.

1.6 Theorem (Hartog's theorem). *If a complex-valued function is holomorphic in each variable separately in a domain $U \subset \mathbb{C}^n$ then it is holomorphic in U.*

Proof. We prove this by induction on the dimension n. The theorem is trivial in one dimension. Thus, we suppose that $n > 1$ and that the theorem is true for dimension $n-1$.

Let a be any point in U and let $\bar{\Delta}(a, r)$ be a closed polydisc contained in U. As in the previous lemma, we write $z = (z', z_n)$ for points of $\mathbb{C}^n$ and write $\bar{\Delta}(a, r) = \bar{\Delta}(a', r') \times \bar{\Delta}(a_n, r_n)$.

Consider the sets

$$X_k = \{z_n \in \bar{\Delta}(a_n, r_n/2) : \quad |f(z', z_n)| \leq k \quad \forall z' \in \bar{\Delta}(a', r')\}$$

For each k this set is closed since $f(z', z_n)$ is continuous in z_n for each fixed z'. On the other hand, since $f(z', z_n)$ is also continuous in z' and, hence bounded on $\bar{\Delta}(a', r')$ for each z_n, we have that $\bar{\Delta}(a_n, r_n/2) \subset \bigcup_k X_k$. It follows from the Baire category theorem that for some k the set X_k contains a neighborhood $\Delta(b_n, \delta)$ of some point $b_n \in \Delta(a_n, r_n/2)$.

We now know that f is separately holomorphic and uniformly bounded in the polydisc $\Delta(a', r') \times \bar{\Delta}(b_n, \delta)$. It follows from Cauchy's inequalities in each variable separately that the first order complex partial derivatives are also uniformly bounded on compact subsets of this polydisc. This, together with the Cauchy-Riemann equations implies that all first order partial derivatives are bounded on compacta in this polydisc and this implies uniform continuity on compacta by the mean value theorem. We conclude from Osgood's theorem that f is holomorphic on $\Delta(a', r') \times \Delta(b_n, \delta)$ and, in fact, has a power series expansion about (a', b_n) which converges uniformly on compact subsets of this polydisc.

Now choose $s_n > r_n/2$ so that $\Delta(b_n, s_n) \subset \Delta(a_n, r_n)$. Then $f(z', z_n)$ is holomorphic in z_n on $\Delta(b_n, s_n)$ for each $z' \in \Delta(a', r')$ and, hence, its power series expansion about (a', b_n) converges as a power series in z_n on $\Delta(b_n, s_n)$ for each fixed point $z' \in \Delta(a', r')$. It follows from Hartog's lemma that f is actually holomorphic on all of $\Delta(a', r') \times \Delta(b_n, s_n)$. Since a is in this set and a was an arbitrary element of U, the proof is complete.

We introduce the first order partial differential operators

$$\frac{\partial}{\partial z_j} = \frac{1}{2}\left(\frac{\partial}{\partial x_j} - i\frac{\partial}{\partial y_j}\right) \quad \frac{\partial}{\partial \bar{z}_j} = \frac{1}{2}\left(\frac{\partial}{\partial x_j} + i\frac{\partial}{\partial y_j}\right)$$

Clearly, the Cauchy-Riemann equations mean that a function is holomorphic in each variable separately in a domain U if and only if its first order partial derivatives exist in U and $\frac{\partial}{\partial \bar{z}_j} f = 0$ in U for each j. We may write these equations in a more succinct form by using differential forms. We choose as a basis for the complex one forms in $\mathbb{C}^n$ the forms

$$dz_j = dx_j + idy_j, \quad d\bar{z}_j = dx_j - idy_j, \quad j = 1, ..., n$$

then the differential df of a function f decomposes as $\partial f + \bar{\partial} f$, where

$$\partial f = \sum_j \frac{\partial f}{\partial z_j} dz_j, \quad \bar{\partial} f = \sum_j \frac{\partial f}{\partial \bar{z}_j} d\bar{z}_j$$

and the Cauchy-Riemann equations become simply $\bar{\partial} f = 0$. Thus, a function defined in a domain U is holomorphic in U provided its first order partial derivatives exist and $\bar{\partial} f = 0$ in U. A much stronger resultis true: A distribution ψ defined in a domain in $\mathbb{C}^n$ which satisfies the Cauchy-Riemann equations $\bar{\partial}\psi = 0$ in the distribution sense is actually a holomorphic function. We will not prove this here, but it follows from appropriate regularity theorems for elliptic PDE's.

If U is an open subset of $\mathbb{C}^n$ then the space of all holomorphic functions on U will be denoted by $\mathcal{H}(U)$. This is obviously a complex algebra under the operations of pointwise addition and multiplication of functions. It also has a natural topology - the topology of uniform convergence on compact subsets of U. We end this section by proving a couple of important theorems about this topology.

The topology of uniform convergence on compacta is defined as follows. For each compact set $K \subset U$ we define a seminorm $|| \cdot ||_K$ on $\mathcal{H}(U)$ by

$$||f||_K = \sup_{z \in K} |f(z)|$$

Then a neighborhood basis at $g \in \mathcal{H}(U)$ consists of all sets of the form

$$\{f \in \mathcal{H}(U) : ||f - g||_K < \epsilon\}$$

for K a compact subset of U and $\epsilon > 0$. If $\{K_n\}$ is an increasing sequence of compact subsets of U with the property that each compact subset of U is contained in some K_n, then a basis for this topology is also obtained using just the sets of the above form with K one of the $K_n's$ and ϵ one of the numbers m^{-1} for m a positive interger. Thus, there is a countable basis for the topology at each point and, in fact, $\mathcal{H}(U)$ is a metric space in this topology. Clearly, a sequence converges in this topology if and only if it converges uniformly on each compact subset of U. A nice application of Osgoods's Lemma is the proof that $H(U)$ is complete in this topology:

1.7 Theorem. *If $\{f_n\}$ is a sequence of holomorphic functions on U which is uniformly convergent on each compact subset of U, then the limit is also holomorphic on U.*

Proof. For holomorphic functions of one variable this is a standard result. Its proof is a simple application of Morera's Theorem. In the several variable case we simply apply this result in each variable separately (with the other variables fixed) to conclude that the limit of a sequence of holomorphic functions, converging uniformly on compacta, is holomorphic in each variable separately. Such a limit is also obviously continuous and so is holomorphic by Osgood's Lemma (we could appeal to Hartog's Theorem but we really only need the much weaker Osgoods Lemma).

A topological vector space with a topology defined by a sequence of seminorms, as above, and which is complete in this topology is called a *Frechet space*. Thus, $\mathcal{H}(U)$ is a Frechet space. It is, in fact, a Montel space. The content of this statement is that it is a Frechet space with the following additional property:

1.8 Theorem. *Every closed bounded subset of $\mathcal{H}(U)$ is compact, where $A \subset \mathcal{H}(U)$ is bounded if $\{||f||_K; f \in A\}$ is bounded for each compact subset K of U.*

Proof. Since $H(U)$ is a metric space we need only prove that every bounded sequence in $\mathcal{H}(U)$ has a convergent subsequence. Thus, let $\{f_n\}$ be a bounded sequence in $\mathcal{H}(U)$. By the previous theorem, we need only show that it has a subsequence that converges uniformly on compacta to a continuous function - the limit will then automatically be holomorphic as well. By the Ascoli-Arzela Theorem, a bounded sequence of continuous functions on a compact set has a uniformly convergent subsequence if it is equicontinuous. It follows from the Cauchy estimates that $\{f_n\}$ has uniformly bounded partial derivatives on each compact set. This, together with the mean value theorem implies that $\{f_n\}$ is equicontinuous on each compact set and, hence, has a uniformly convergent subsequence on each compact set. We choose a sequence $\{K_m\}$ of compact subsets of U with the property that each compact subset of U is contained in some K_m. We then choose inductively a sequence of subsequences $\{f_{n_{m,i}}\}_i$ of $\{f_n\}$ with the property that $\{f_{n_{m+1,i}}\}_i$ is a subsequence of $\{f_{n_{m,i}}\}_i$ and $\{f_{n_{m,i}}\}_i$ is uniformly convergent on K_m for each m. The diagonal of the resulting array of functions converges uniformly on each K_m and, hence, on each compact subset of U. This completes the proof.

If U is a domain in C^n and $F : U \to C^m$ is a map, then F is called holomorphic if each of its coordinate functions is holomorphic.

1. Problems

1. Prove that the composition of two holomorphic mappings is holomorphic.
2. Formulate and prove the identity theorem for holomorphic functions of several variables.
3. Formulate and prove the maximum modulus theorem for holomorphic functions of several variables.
4. Formulate and prove Schwarz's lemma for holomorphic functions of several variables.
5. Prove that a function which is holomorphic in a connected neighborhood of the boundary of a polydisc in $\mathbb{C}^n$, $\quad n > 1$, has a unique holomorphic extension through the interior of the polydisc.
6. Prove that if a holomorphic function on a domain (a connected open set) is not identically zero, then the set where it vanishes has $2n$-dimensional Lebesgue measure zero (hint: use Jensen's inequality).

2. Local theory

Let X be a topological space and x a point of X. If f and g are functions defined in neighborgoods U and V of x and if $f(y) = g(y)$ for all y in some third neighborhood $W \subset U \cap V$ of x, then we say that f and g are *equivalent* at x. The equivalence class consisting of all functions equivalent to f at x is called the *germ* of f at x.

The set of germs of complex valued functions at x is clearly an algebra over the complex field with the algebra operations defined in the obvious way. In fact, this algebra can be described as the inductive limit $\varinjlim F(U)$ where $F(U)$ is the algebra of complex valued functions on U and the limit is taken over the directed set consisting of neighborhoods of x. The germs of continuous functions at x obviously form a subalgebra of the germs of all complex valued functions at x and, in the case where $X = \mathbb{C}^n$, the germs of holomorphic functions at x form a subalgebra of the germs of $\mathcal{C}^\infty$ functions at x which, in turn, form a subalgebra of the germs of continuous functions at X.

We shall denote the algebra of holomorphic functions in a neighborhood $U \subset \mathbb{C}^n$ by $\mathcal{H}(U)$ and the algebra of germs of holomorphic functions at $z \in \mathbb{C}^n$ by $\mathcal{H}_z$ or by ${}_n\mathcal{H}_z$ in case it is important to stress the dimension n. We have that $\mathcal{H}_z = \varinjlim \mathcal{H}(U)$ where the limit is taken over a system of neighborhoods of z. We also have as an immediate consequence of the definition of holomorphic function that:

2.1 Proposition. *The algebra ${}_n\mathcal{H}_0$ may be described as the algebra $\mathbb{C}\{z_1, ..., z_n\}$ of convergent power series in n variables.*

Another important algebra of germs of functions is the algebra ${}_n\mathcal{O}_z$ of germs of regular functions at $z \in \mathbb{C}^n$. Here we give $\mathbb{C}^n$ the *Zariski topology.* This is the topology in which the closed sets are exactly the algebraic subvarieties of $\mathbb{C}^n$. By an algebraic subvariety of $\mathbb{C}^n$ we mean a subset which is the set of common zeroes of some set of complex polynomials in $\mathbb{C}[z_1, ..., z_n]$. A regular function on a Zariski open set U is a rational function with a denominator which does not vanish on U. The algebra of regular functions on a Zariski open set U will be denoted $\mathcal{O}(U)$. The algebra of germs of regular functions at $z \in \mathbb{C}^n$ is then ${}_n\mathcal{O}_z = \varinjlim \mathcal{O}(U)$. One easily sees that:

2.2 Proposition. *The algebra ${}_n\mathcal{O}_z$ is just the ring of fractions of the algebra $\mathbb{C}[z_1, ..., z_n]$ with respect to the multiplicative set consisting of polynomials which do not vanish at z.*

By a local ring we will mean a ring with a unique maximal ideal. It is a trivial observation that:

2.3 Proposition. *The algebras ${}_n\mathcal{O}_z$ and ${}_n\mathcal{H}_z$ are local rings and in each case the maximal ideal consists of the elements which vanish at z.*

The algebras ${}_n\mathcal{O}_z$ and ${}_n\mathcal{H}_z$ are, in fact, Noetherian rings (every ideal is finitely generated). For ${}_n\mathcal{O}_z$ this is a well known and elementary fact from commutative algebra. We will give the proof here because the main ingredient (Hilbert's basis theorem) will also be needed in the proof that ${}_n\mathcal{H}_z$ is Noetherian.

We will use the elementary fact that if M is a finitely generated module over a Noetherian ring A then every submodule and every quotient module of M is also finitely generated.

2.4 Theorem (Hilbert basis theorem). *If A is a Noetherian ring then the polynomial ring $A[x]$ is also Noetherian.*

Proof. Let I be an ideal in $A[x]$ and let J be the ideal of A consisting of all leading coeficients of elements of I. Since A is Noetherian, J has a finite set of generators $\{a_1, ..., a_n\}$. For each i there is an $f_i \in I$ such that $f_i = a_i x^{r_i} + g_i$ where g_i has degree less than r_i.

Let $r = \max_i r_i$ and let $f = ax^m + g$ (deg $g < m$) be an element of I of degree $m \geq r$. We may choose $b_i, ..., b_n \in A$ such that $a = \sum_i b_i A_i$. then $f - \sum_i b_i f_i x^{m-r_i}$ belongs to I and has degree less than m. By iterating this process we conclude that every element of I may be written as the sum of an element of the ideal generated by $\{f_1, ..., f_n\}$ and a polynomial of degree less than r belonging to I. However, the polynomials of degree less than r form a finitely generated A module and, hence, the submodule consisting of its intersection with I is also finitely generated. A generating set for this A module, together with $\{f_1, ..., f_n\}$ provides a set of generators for I as an $A[x]$-module. This completes the proof.

The above result and induction show that $\mathbb{C}[z_1, ..., z_n]$ is Noetherian. This implies that $\mathcal{O}_z$ is Noetherian as follows: If I is an ideal of $\mathcal{O}_z$ and $J = I \cap \mathbb{C}[z_1, ..., z_n]$ then J generates I as an $\mathcal{O}_z$-module but J is finitely generated as a $\mathbb{C}[z_1, ..., z_n]$-module. It follows that I is finitely generated as an $\mathcal{O}_z$-module. In summary:

2.5 Theorem. *The polynomial algebra $\mathbb{C}[z_1, ..., z_n]$ and the local algebra ${}_n\mathcal{O}_z$ are Noetherian rings.*

We now proceed to develop the tools needed to prove that ${}_n\mathcal{H}_z$ is Noetherian.

A holomorhic function defined in a neighborhood of 0 is said to be *regular* of order k in z_n at 0 provided $f(0, ..., 0, z_n)$ has a zero of order k at 0.

2.6 Theorem. *If f is holomorphic in a neighborhood U of 0 in $\mathbb{C}^n$ and regular of order k in z_n at 0, then there is a polydisc $\Delta(0, r') \times \Delta(0, r_n)$ such that for each $z' \in \Delta(0, r')$, as a function of z_n, $f(z', z_n)$ has exactly k zeroes in $\Delta(0, r_n)$, counting multipicity, and no zeroes on the boundary of $\Delta(0, r_n)$.*

Proof. Choose r_n small enough that the only zeroes of $f(0, z_n)$ on $\bar{\Delta}(0, r_n)$ occur at $z_n = 0$. Set

$$\delta = \inf\{|f(0, z_n)| : \quad |z_n| = r_n\}$$

and choose r' small enough that

$$|f(z', z_n) - f(0, z_n)| < \delta \text{ whenever } z' \in \Delta(0, r'), \quad |z_n| = r_n$$

Then Rouche's theorem implies that for each $z' \in \Delta(0, r')$ the functions $f(0, z_n)$ and $f(z', z_n)$ of z_n have the same number of zeroes in the disc $\Delta(0, r_n)$. This completes the proof.

A thin subset of $\mathbb{C}^n$ is a set which locally at each point of $\mathbb{C}^n$ is contained in the zero set of a holomorphic function.

2.7 Theorem (Removable singularity theorem). *If f is bounded and holomorphic in an open set of the form $U - T$ where U is open and T is thin, then f has a unique holomorphic extension to all of U.*

Proof. This is a local result and needs only to be proved for a neighborhood of each point $a \in U$. We may assume that T intersected with some neighborhood of a is contained in the zero set of a holomorphic function g which we may assume is regular in z_n of some order k (otherwise we simply perform a coordinate change). By the proceeding result we may choose a neighborhood of the form $\Delta(a,r) = \Delta(a',r') \times \Delta(a_n, r_n)$ and assume that for each $z' \in \Delta(a',r')$ the set T meets $\{z'\} \times \bar{\Delta}(a_n, r_n)$ in at most k points all of which lie in $\{z'\} \times \Delta(a_n, r_n)$. Then the function

$$h(z', z_n) = \frac{1}{2\pi i} \int_{|\zeta_n - a_n| = r_n} \frac{f(z', \zeta_n)}{\zeta_n - z_n} d\zeta_n$$

is holomorphic in all of $\Delta(a,r)$ and agrees with f off T by the one variable removable singularity theorem and the Cauchy integral theorem.

A *Weierstrass polynomial* of degree k in z_n is a polynomial $h \in {}_{n-1}\mathcal{H}_0[z_n]$ of the form

$$h(z) = z_n^k + a_1(z')z_n^{k-1} + ... + a_{k-1}(z')z_n + a_k(z')$$

where $z = (z', z_n)$ and each a_i is a non-unit in ${}_{n-1}\mathcal{H}_0$.

2.8 Theorem (Weierstrass preparation theorem). *If $f \in {}_n\mathcal{H}_0$ is regular of order k in z_n, then f has a unique factorization as $f = uh$ where h is a Weierstrass polynomial of degree k in z_n and u is a unit in ${}_n\mathcal{H}_0$.*

Proof. Choose some representative of f and a polydisc $\Delta(0,r)$ in which $f(z', z_n)$ has exactly k zeroes (none on the boundary) as a function of z_n for each $z' \in \Delta(0, r')$ as in 2.6. Label the zeroes $b_1(z'), ..., b_k(z')$. The polynomial we are seeking is

$$h(z) = \prod_{j=1}^{k} (z_n - b_j(z')) = z_n^k + a_1(z')z_n^{k-1} + ... + a_{k-1}(z')z_n + a_k(z')$$

Now the functions $b_j(z')$ need not even be continuous because of the arbitrary choices made in labeling the zeroes of f. However, the functions $a_j(z')$ are, in fact, holomorphic. To see this, note that these functions are the elementary symmetric functions of the b_j's and these, in turn, may be written as polynomials in the power sums s_m where

$$s_m(z') = \sum_{j=1}^{k} b_j(z')^m = \frac{1}{2\pi i} \int_{|\zeta| = r_n} \zeta^m \frac{\partial f(z', \zeta)}{\partial \zeta} \frac{d\zeta}{f(z', \zeta)}$$

These functions and, consequently, the $a_j(z')$ are holomorphic in $\Delta(0, r')$. Note that the b_j's all vanish at $z' = 0$ and, thus, so do the a_j's. We now have that h is a Weierstrass polynomial.

The proof will be complete if we can show that $u = f/h$ is holomorphic and non-vanishing in $\Delta(0,r)$. For each fixed $z' \in \Delta(0,r')$ the function $f(z',z_n)/h(z',z_n)$ is holomorphic in z_n in $\Delta(0,r_n)$ since numerator and denominator have exactly the same zeroes in this polydisc. Furthermore, h is bounded away from zero on $\Delta(0,r') \times \partial\Delta(0,r_n)$. This and the maximum modulus principal imply that f/h is bounded on $\Delta(0,r)$. Since it is holomorphic in this set except where h vanishs, the removable singularities theorem (2.7) implies that it extends to be holomorphic and non-vanishing in the entire polydisc. The uniqueness is clear from the construction. This completes the proof.

2.9 Theorem (Weierstrass division theorem). *If $h \in {}_{n-1}\mathcal{H}[z_n]$ is a Weierstrass polynomial of degree k and $f \in {}_n\mathcal{H}$, then f can be written uniquely in the form $f = gh + q$ where $g \in {}_n\mathcal{H}$ and $q \in {}_{n-1}\mathcal{H}[z_n]$ is a polynomial in z_n of degree less than k. Furthermore, if f is a polynomial in z_n then so is g.*

Proof. Choose representatives of f and h (still call them f and h) which are defined in a neighborhood of a polydisc $\bar{\Delta}(0,r)$ which is chosen small enough that $h(z',z_n)$ has exactly k zeroes in $\Delta(0,r_n)$ as a function of z_n for each $z' \in \bar{\Delta}(0,r')$ with none occuring on $|z_n| = r_n$. Then the function

$$g(z) = \frac{1}{2\pi i}\int_{|\zeta|=r_n} \frac{f(z',\zeta)}{h(z',\zeta)}\frac{d\zeta}{\zeta - z_n}$$

is holomorphic in $\Delta(0,r)$ as is the function $q = f - gh$. The function q may be written as

$$q(z) = \frac{1}{2\pi i}\int_{|\zeta|=r_n} \frac{f(z',\zeta)}{h(z',\zeta)}\frac{h(z',\zeta) - h(z',z_n)}{\zeta - z_n}d\zeta$$

But the function

$$\frac{h(z',\zeta) - h(z',z_n)}{\zeta - z_n}$$

is a polynomial in z_n of degree less than k (with coeficients which are functions of ζ) which shows that q is a polynomial in z_n of degree less than k.

To show that this representation is unique, suppose we have two representations

$$f = gh + q = g_1h + q_1$$

with q and q_1 both polynomials of degree less than k in z_n. Then $q - q_1 = h(g_1 - g)$ is a polynomial of degree less than k in z_n with at least k zeroes for each fixed value of $z' \in \Delta(0,r')$. This is only possible if it is identically zero.

Now if f itself is a polynomial in z_n then the usual division algorithm for polynomials over a commutative ring gives a representation of f as above with g a polynomial in z_n. The uniqueness says that this must coincide with the representation given above. This completes the proof.

2.10 Theorem. *The ring ${}_n\mathcal{H}$ is a Noetherian ring.*

Proof. We proceed by induction on the dimension n. For $n = 0$ we are talking about a field, $\mathbb{C}$, which has no non-trivial ideals and is trivially Noetherian. Suppose that ${}_{n-1}\mathcal{H}$ is Noetherian. Then we know from Hilbert's basis theorem that ${}_{n-1}\mathcal{H}[z_n]$ is Noetherian. Suppose that I is an ideal in ${}_n\mathcal{H}$. Let h be any non-zero element of I. By performing a coordinate change if necessary, we may assume that h is regular in z_n and from the Weierstrass preparation theorem we may (after multiplying by a unit if neccessary) assume that h is a Weierstrass polynomial in z_n. In other words, h belongs to $I \cap_{n-1} \mathcal{H}[z_n]$ which is an ideal in the Noetherian ring ${}_{n-1}\mathcal{H}[z_n]$ and is, thus, finitely generated by say $g_1, ..., g_m$. Now by the Weierstrass division theorem any element $f \in I$ may be written as $f = gh + q$ with $g \in {}_n\mathcal{H}$ and $q \in I \cap {}_{n-1}\mathcal{H}[z_n]$. But this means that h and q both belong to $I \cap_{n-1} \mathcal{H}[z_n]$ and, hence, to the ideal generated by $g_1, ..., g_m$. Therefore f belongs to the ideal generated by $g_1, ..., g_m$ and we conclude that this set of elements generates I. this completes the proof.

2. Problems

1. A unique factorization domain is an integral domain in which each element has a unique (up to units) factorization as a finite product of irreducible factors. Prove that if A is a unique factorization domain then so is $A[x]$.
2. Prove that $\mathbb{C}[z_1, ..., z_n]$ and ${}_n\mathcal{O}_0$ are unique factorization domains.
3. Prove that ${}_n\mathcal{H}_0$ is a unique factorization domain.
4. Prove Nakayama's lemma: If M is a finitely generated module over a local ring A with maximal ideal $\mathfrak{m}$ and if $\mathfrak{m}M = M$ then $M = 0$. Hint: Prove that if M has a generating set with k elements with $k > 0$ then it also a generating set with $k - 1$ elements.
5. Prove the implicit function theorem: If f is holomorphic in a neighborhood of $a = (a', a_n)$, $f(a) = 0$, and $\dfrac{\partial f}{\partial z_n}(a) \neq 0$ then there is a polydisc $\Delta(a, r) = \Delta(a', r') \times \Delta(a_n, r_n)$ and a holomorphic map $g : \Delta(a', r') \to \Delta(a, r)$ such that $g(a') = a$ and, for each $z \in \Delta(a, r)$, $f(z) = 0$ if and only if $z = g(z')$ for some $z' \in \Delta(a', r')$. Hint: Use the Weierstrass preparation theorem with $k = 1$.

3. A Little Homological Algebra

Let A be a commutative algebra over a field and let M and N be A-modules. We will denote the tensor product of M and N as vector spaces over the field simply by $M \otimes N$. The tensor product of M and N as A modules is denoted $M \otimes_A N$ and is the cokernel of the map $M \otimes A \otimes N \to M \otimes N$ defined by $m \otimes a \otimes n \to am \otimes n - m \otimes an$. Clearly, $M \otimes_A N$ is a covariant functor of each of its arguments for fixed values of the other argument.

The functor $M \otimes (\cdot)$ from the category of vector spaces to itself is exact. That is, if the sequence

$$0 \to N_1 \to N_2 \to N_3 \to 0$$

is exact (the kernel of each map is equal to the image of the preceding map) then so is the sequence

$$0 \to M \otimes N_1 \to M \otimes N_2 \to M \otimes N_3 \to 0$$

On the other hand, the functor $M \otimes_A (\cdot)$ from the category of A-modules to itself fails to be exact. It is right exact (kernel equals image at the right hand and middle stages) as is easily seen by chasing the diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & M \otimes A \otimes N_1 & \longrightarrow & M \otimes A \otimes N_2 & \longrightarrow & M \otimes A \otimes N_3 & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & M \otimes N_1 & \longrightarrow & M \otimes N_2 & \longrightarrow & M \otimes N_3 & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & M \otimes_A N_1 & \longrightarrow & M \otimes_A N_2 & \longrightarrow & M \otimes_A N_3 & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

However, it is not generally true that this functor preserves exactness at the left stage.

The tensor product relative to A is an important functor and the fact that it is right exact but not exact cannot be ignored. This circumstance requires careful analysis and the development of tools that allow one to deal effectively with the problems it poses.

A similar problem arises with the functor $\hom_A$. Again, if M and N are modules over A then the space of linear maps from M to N is denoted $\hom(M, N)$ while the A module of A-module homomorphisms from M to N is denoted $\hom_A(M, N)$. It can be described as the kernel of the map $\hom(M, N) \to \hom(A \otimes M, N)$ defined by $\phi \to \{a \otimes m \to a\phi(m) - \phi(am)\}$. The functor $\hom(M, \cdot)$ is covariant and exact while the functor $\hom(\cdot, N)$ is contravariant and exact. On the other hand, the functors $\hom_A(M, \cdot)$ and $\hom_A(\cdot, N)$ are not generally exact. It is true that they are both left exact, as is easily seen from the definition and a diagram chase like the one above. Here again, the fact that $\hom_A(\cdot, N)$ is not exact leads to the need to develop tools to deal with the problems that this poses.

The first step in this program is to understand exactly when the two functors in question are exact.

3.1 Definition. *An A-module M for which $M \otimes_A (\cdot)$ is an exact functor from the category of A-modules to itself is called a flat A-module.*

An A-module M for which $\hom_A(M, \cdot)$ is an exact functor is called a projective A-module.

An A-module for which $\hom_A(\cdot, M)$ is an exact functor is called an injective A-module.

Of course, each of these functors already satisfies two of the three conditions for exactness. Thus, M is flat if and only if for every injection $i : N_1 \to N_2$ the induced morphism $M \otimes_A N_1 \to M \otimes_A N_2$ is an injection . Similarly, M is projective if and only if for every surjection $N_2 \to N_3$ the induced morphism $\hom_A(M, N_2) \to \hom_A(M, N_3)$ is a surjection (ie. that every morphism from M to N_3 lifts to N_2). Finally, an A-module M is injective if and only if for every injection $N_1 \to N_2$ the induced morphism $\hom_A(N_2, M) \to \hom_A(N_1, M)$ is surjective (ie. every morphism from N_1 to M extends to N_2).

3.2 Proposition. *Let M be an A-module. Then there are natural isomorphisms*

$$A \otimes_A M \to M \quad \textit{and} \quad M \to \hom_A(A, M)$$

Proof. The map $a \otimes m \to am : A \otimes M \to M$ has kernel which contains the image of the map

$$a \otimes b \otimes m \to ab \otimes m - a \otimes bm : A \otimes A \otimes M \to A \otimes M$$

In fact the two are equal since, if $\sum a_i m_i = 0$, then $\sum a_i \otimes m_i$ is the image under the latter map of $\sum 1 \otimes a_i \otimes m_i$. It follows from the definition of $\otimes_A$ that $a \otimes m \to am : A \otimes M \to M$ induces an isomorphism from $A \otimes_A M$ to M.

The case of $\hom(A, M)$ is even easier. Each element $m \in M$ determines a homomorphism $a \to am : A \to M$ and every homomorphism from A to M arises in this way from a unique element m. In fact, m is just the image of 1 under the homomorphism. This completes the proof.

A trivial consequence of this proposition is that the functors $M \otimes_A (\cdot)$ and $\hom_A(M, \cdot)$ are exact in the case where $M = A$. This is also clearly true if M is a finite direct sum of copies of A. It is only slighltly less trivial that this continues to hold if M is an arbitrary direct sum of copies of A. A module M which is a direct sum of an (arbitrary) set of copies of A is called a free A-module . Thus, the functors $M \otimes_A (\cdot)$ and $\hom_A(M, \cdot)$ are exact when M is any free A- module. Finally, it is an easy consequence of the definition of direct summand that if M is a direct summand of a free A-module then $M \otimes_A (\cdot)$ is an exact functor. The same arguments applied to the functor $\hom_A(M, \cdot)$ show that it is also exact when M is a direct summand of a free A-module. Thus, direct summands of free modules are both flat and projective.

Since A is an algebra over a field, an A-module M is free if and only if it has the form $M = A \otimes X$ where X is a vector space over the same field. Here, the module action is given by the action of A on the left factor in the tensor product; that is, $a(b \otimes x) = ab \otimes x$. Note, every A-module M is a quotient of a projective, in fact a free, A-module. In fact the morphism

$$\epsilon : A \otimes M \to M, \quad \epsilon(a \otimes m) = am$$

expresses M as a quotient of the free A-module $A \otimes M$. If M happens to be projective, the identity morphism from M to M lifts to $A \otimes M$ and embedds M as a direct summand of the free module $A \otimes M$. In summary, we have proved:

3.3 Proposition.

(a) *An A-module is projective if and only if it is a direct summand of a free A-module;*
(b) *every projective module is flat;*
(c) *every A module is a quotient of a projective, in fact a free, A-module.*

If $P_0 \to M$ is a morphism which expresses M as a quotient of a projective module P_0 and if K is the kernel of this morphism then we may express K as a quotient of a projective module P_1. This yields an exact sequence

$$P_1 \to P_0 \to M \to 0$$

by continuing in this way we construct a *projective resolution of M*, that is, an exact sequence of the form:

$$\cdots \to P_n \to P_{n-1} \to \cdots \to P_1 \to P_0 \to M \to 0$$

in which each P_i is projective. Thus, we have proved:

3.4 Proposition. *Each A module has a projective resolution.*

For technical reasons it is useful to have a construction of projective resolutions which is functorial, that is, one in which the complex of projectives

$$\cdots \to P_n \to P_{n-1} \to \cdots \to P_1 \to P_0 \to 0$$

is a functor of the module M. There are many ways of doing this. We briefly describe one, the Hochschild resolution, which has some nice properties and is used extensively. This is the resolution

$$\dots \xrightarrow{\delta_{n+1}} F_n(M) \xrightarrow{\delta_n} \dots \xrightarrow{\delta_2} F_1(M) \xrightarrow{\delta_1} F_0(M) \xrightarrow{\epsilon} M \longrightarrow 0$$

where $F_n(M)$ is the free A-module $(\otimes^{n+1}A) \otimes M$, $\epsilon(a \otimes m) = am$, and

$$\delta_n(a_0 \otimes \cdots \otimes a_n \otimes m) = \sum_{i=0}^{n-1}(-1)^i a_0 \otimes \cdots \otimes a_i a_{i+1} \otimes \cdots \otimes a_n \otimes m + (-1)^n a_0 \otimes \cdots \otimes a_{n-1} \otimes a_n m$$

A simple calculation shows that $\delta_n \circ \delta_{n+1} = 0$ for $n \geq 1$ and $\epsilon \circ \delta_1 = 0$, so that this is a complex. To show that this complex is exact and, hence, provides a projective resolution of M we construct a contracting homotopy for it as a sequence of vector spaces by defining $s_n : F_n(M) \to F_{n+1}(M)$ by

$$s_n(a_0 \otimes a_1 \otimes \cdots \otimes a_n \otimes m) = 1 \otimes a_0 \otimes a_1 \otimes \cdots \otimes a_n \otimes m$$

and $s_{-1} : M \to F_0(M)$ by

$$s_{-1}(m) = 1 \otimes m$$

Note that though the maps δ_n are A-module homomorphisms the maps s_n are only linear maps and not A-module homomorphisms. A direct calculation shows that

$$s_{n-1} \circ \delta_n + \delta_{n+1} \circ s_n = 1, \quad n \geq 1$$

and

$$s_{-1} \circ \epsilon + \delta_1 \circ s_0 = 1$$

and this is exactly what is meant by the statement that $\{s_n\}$ is a contracting homotopy for the above complex. It is immediate that a complex with a contracting homotopy is exact. We shall write $F(M)$ for the complex

$$\dots \xrightarrow{\delta_{n+1}} F_n(M) \xrightarrow{\delta_n} \dots \xrightarrow{\delta_2} F_1(M) \xrightarrow{\delta_1} F_0(M) \longrightarrow 0$$

and

$$F(M) \xrightarrow{\epsilon} M \longrightarrow 0$$

for the corresponding resolution of M. To make sense of this notation, just think of ϵ as being a morphism between two complexes, where M is thought of as the complex whose only non-zero term is the degree zero term which is the module M. Finally, note that the each of the functors F_n is, by construction, an exact functor from A-modules to free A-modules and, hence, F is an exact functor from A-modules to complexes of free A-modules.

To summarize the above discussion, we have

3.5 Proposition. *The Hochschild functor functor $M \to F(M)$ is an exact functor from A-modules to complexes of free A-modules and for each M*

$$F(M) \xrightarrow{\epsilon} M \longrightarrow 0$$

is a free resolution of M.

We have used the term *complex* several times in this discussion. Actually two types of complexes occur and it is time to be more precise. A chain complex C of A-modules is a sequence of modules and morphisms of the form

$$\dots \xrightarrow{\delta_{n+2}} C_{n+1} \xrightarrow{\delta_{n+1}} C_n \xrightarrow{\delta_n} C_{n-1} \xrightarrow{\delta_{n-1}} \dots$$

If C is a chain complex then its n^{th} homology is

$$H_n(C) = \ker \delta_n / \mathrm{im}\ \delta_{n+1}$$

A cochain complex is just a chain complex with the modules indexed in increasing order instead of decreasing order. Also, it is traditional to index cochain complexes with superscripts. Thus, a cochain complex C is a sequence of modules and morphisms of the form

$$\dots \xrightarrow{\delta^{n-2}} C^{n-1} \xrightarrow{\delta^{n-1}} C^n \xrightarrow{\delta^n} C^{n+1} \xrightarrow{\delta^{n+1}} \dots$$

If C is a cochain complex then its n^{th} cohomology is

$$H^n(C) = \ker \delta_n / \mathrm{im}\ \delta_{n-1}$$

We can now define the functors tor_n^A and ext_A^n.

3.6 Definition. *Let M and N be A-modules. Then we define*

$$\mathit{tor}_n^A(M,N) = H_n(M \otimes_A F(N)) \text{ and } \mathit{ext}_A^n(M,N) = H^n(\hom_A(F(M),N))$$

A quick look at the first two terms of the complexes $M \otimes_A F(N)$ and $\hom_A(F(M),N)$ shows that:

3.7 Proposition. $\mathit{tor}_0^A(M,N) = M \otimes_A N$ *and* $\mathit{ext}_A^0(M,N) = \hom_A(M,N)$.

Another simple fact is that the isomorphism $n \otimes m \to m \otimes n : M \otimes N \to N \otimes M$ extends to an isomorphism between the complexes $M \otimes_A F(N)$ and $N \otimes_A F(M)$ and, hence, determines an isomorphism $\mathrm{tor}_n^A(M,N) \to \mathrm{tor}_n^A(N,M)$.

3.8 Theorem. *Let $0 \to M_1 \to M_2 \to M_3 \to 0$ be a short exact sequence of A-modules and let N be any A-module. Then there is a natural long exact sequence*

$$\cdots \to \mathit{tor}_{n+1}^A(M_3,N) \to \mathit{tor}_n^A(M_1,N) \to \mathit{tor}_n^A(M_2,N) \to \mathit{tor}_n^A(M_3,N) \to \cdots$$

$$\cdots \to \mathit{tor}_1^A(M_3,N) \to M_1 \otimes_A N \to M_2 \otimes_A N \to M_3 \otimes_A N \to 0$$

Again, the same result holds with the roles of M and N reversed.

Proof. If we take the tensor product of the short exact sequence $0 \to M_1 \to M_2 \to M_3 \to 0$ with the Hochschild complex $F(N)$ for N we obtain a short exact sequence of complexes

$$0 \to M_1 \otimes_A F(N) \to M_2 \otimes_A F(N) \to M_3 \otimes_A F(N) \to 0$$

That this is exact follows from the fact that $F(N)$ is a complex of free A-modules. A simple diagram chase proves the standard result that every short exact sequence of complexes induces a long exact sequence of the corresponding homology. This completes the proof.

The same sort of arguments yield analogous results for ext:

3.9 Theorem. *Let $0 \to M_1 \to M_2 \to M_3 \to 0$ be a short exact sequence of A-modules and let N be any A-module. Then there are natural long exact sequences*

$$0 \to \mathit{hom}_A(M_3,N) \to \mathit{hom}_A(M_2,N) \to \mathit{hom}_A(M_1,N) \to \mathit{ext}_A^1(M_3,N) \to \cdots$$

$$\to \mathit{ext}_A^n(M_3,N) \to \mathit{ext}_A^n(M_2,N) \to \mathit{ext}_A^n(M_1,N) \to \mathit{ext}_A^{n+1}(M_3,N) \to \cdots$$

and

$$0 \to \mathit{hom}_A(N,M_1) \to \mathit{hom}_A(N,M_2) \to \mathit{hom}_A(N,M_3) \to \mathit{ext}_A^1(N,M_1) \to \dots$$

$$\to \mathit{ext}_A^n(N,M_1) \to \mathit{ext}_A^n(N,M_2) \to \mathit{ext}_A^n(N,M_3) \to \mathit{ext}_A^{n+1}(N,M_1) \to \cdots$$

3.10 Theorem. *Let M be a module. Then M is flat if and only if $tor_n^A(M,N) = 0$ for all $n > 0$ and all modules N. Of course, the same statement holds with M and N reversed.*

The module M is projective if and only if $ext_A^n(M,N) = 0$ for all $n > 0$ and all modules N.

The module M is injective if and only if $ext_A^n(N,M) = 0$ for all $n > 0$ and all modules N.

Proof. That $tor_n^A(M,N) = 0$ for all N and all $n > 0$ if M is flat follows from the fact that $M \otimes_A (\cdot)$ preserves exactness and, in particular, preserves the exactness of the Hochschild resolution. The reverse implication follows from the long exact sequence for tor. The proofs of the results for ext are completely analogous.

The following results are trivial consequences of Theorem 3.10 and the existence of the long exact sequences:

3.11 Theorem. *If $0 \to K \to P \to M \to 0$ is a short exact sequence of modules with P projective and N is any module, then there are natural isomorphisms*

$$tor_{n+1}^A(M,N) \simeq tor_n^A(K,N) \quad n > 0$$

$$tor_1^A(M,N) \simeq ker\{K \otimes_A N \to P \otimes_A N\}$$

$$ext_A^{n+1}(M,N) \simeq ext_A^n(K,N) \quad n > 0$$

$$ext_A^1(M,N) \simeq coker\{hom_A(P,N) \to hom_A(K,N)\}$$

This result is the basis for an induction argument that shows that tor and ext can be computed from any projective resolution: Let

$$\cdots \to P_n \to P_{n-1} \to \cdots \to P_1 \to P_0 \to M \to 0$$

be a projective resolution of M and let K_n be the kernel of $P_n \to P_{n-1}$ for $n > 0$, K_0 the kernel of $P_0 \to M$ and set $K_{-1} = M$. Then Theorem 3.11, applied to the short exact sequences

$$0 \to K_n \to P_n \to K_{n-1} \to 0, \quad n \ge 0$$

yields

$$\mathrm{tor}_1(K_{p-2}, N) \simeq \ker\{K_{p-1} \otimes N \to P_{n-1} \otimes N\} \simeq H_p(P \otimes N), \quad p \ge 1$$

and

$$\mathrm{tor}_{n+1}(K_q, N) \simeq \mathrm{tor}_n(K_{q+1}, N) \quad n \ge 0, q \ge 0$$

By induction, this yields

$$\mathrm{tor}_p(M,N) = \mathrm{tor}_p(K_{-1}, N) \simeq \mathrm{tor}_1(K_{p-2}, N) \simeq H_p(P \otimes N)$$

A similar argument works for ext. This leads to:

3.12 Theorem. *Given any projective resolution $P \to M$, as above, and any module N there are natural isomorphisms*

$$tor_n^A(M,N) \simeq H_n(P \otimes_A N) \text{ and } ext_A^n(M,N) \simeq H^n(hom_A(P,N))$$

The result for tor holds if P is just a flat resolution of M.

Finally, we return to the study of the algebras ${}_n\mathcal{O}$ and ${}_n\mathcal{H}$ with a result which shows that they have particularly simple free resolutions. First, note that if A is Noetherian then any finitely generated A-module has a resolution by free finitely generated modules, that is, a resolution of the form

$$\cdots \xrightarrow{\delta_{n+1}} A^{k_n} \xrightarrow{\delta_n} A^{k_{n-1}} \xrightarrow{\delta_{n-1}} \cdots \xrightarrow{\delta_1} A^{k_0} \longrightarrow M \longrightarrow 0$$

Such a resolution is called a *chain of syzygies for M*. A chain of syzygies, as above, is said to *terminate at the n^{th} stage* if the kernel of δ_{n-1} is free. In this case, the kernel of δ_{n-1} may be used to replace A_n resulting in a chain in which the terms beyond the n^{th} one are all zero.

3.13 Theorem (Hilbert's syzygy theorem). *If A is the algebra ${}_n\mathcal{H}$ or the algebra ${}_n\mathcal{O}$ and M is a finitely generated A-module then any chain of sysygies for M terminates at the n^{th} stage.*

Proof. The proofs are not different in the two cases so we shall just do the case of $A = {}_n\mathcal{H}$. Let $I_j \subset \mathcal{H} = {}_n\mathcal{H}$ be the ideal generated by the germs $z_1, \dots z_j$ of the first j coordinate functions. Thus, I_n is the maximal ideal of $\mathcal{H}$. Also, let $K_p \subset \mathcal{H}^{k_p}$ denote the kernel of the map δ_p of the syzygy.

Our objective is to prove that K_{n-1} is free. To this end, let $f_1, \dots, f_s$ be a minimal set of generators of K_{n-1}. Thus, no proper subset of this set generates K_{n-1}. We consider the map $\alpha : \mathcal{H}^s \to \mathcal{H}^{k_{n-1}}$ defined by $\alpha(g_1, \dots, g_s) = \sum g_i f_i$. The image of α is K_{n-1}. Thus, without loss of generality we may replace the original syzygy with one in which $\delta_n = \alpha$. We now show that the kernel K_n of this map is necessarily zero. Note that $K_n \subset I_n\mathcal{H}^s$ since, otherwise, there exist $g_1, \cdots g_s \in \mathcal{H}$ such that

$$\sum g_i f_i = 0$$

with some g_j a unit. Then the set of generators of K_{n-1} could be reduced by throwing out the corresponding f_j. We shall show that this implies that $K_n = I_n K_n$ which, by Nakayama's lemma, shows that $K_n = 0$, as desired.

The fact that $K_n \subset I_n\mathcal{H}^s$ implies $K_n = I_n K_n$ follows immediately from the case $k = j = n$ of the equality

$$K_p \cap I_j\mathcal{H}^{k_p} = I_j K_p \quad \text{for} \quad 1 \le j \le p$$

which we shall prove by induction on j. We need only prove that $K_p \cap I_j\mathcal{H}^{k_p} \subset I_j K_p$ since the reverse containment is clear.

Suppose that $j = 1$ and $f \in K_p \cap I_1\mathcal{H}^{k_p}$ for some $p \geq 1$. Then $\delta_p(f) = 0$ and $f = z_1 g$ for some $g \in \mathcal{H}^{k_p}$. Since $z_1\delta_p(g) = \delta_p(f) = 0$, and $\mathcal{H}$ is an integral domain, it follows that $\delta_p(g) = 0$ and that $g \in K_p$. Then $f = z_1 g \in I_1 K_p$ as desired. Now assume that the above equality holds for some j and all $p \geq j$. Suppose that $f \in K_p \cap I_{j+1}\mathcal{H}^{k_p}$ for some $p \geq j+1$. This means that $\delta_p(f) = 0$ and

$$f = z_1 g_1 + \cdots + z_{j+1} g_{j+1}$$

for some $g_1, \ldots, g_{j+1} \in \mathcal{H}^{k_p}$. From these two facts it follows that

$$z_{j+1}\delta_p(g_{j+1}) = -z_1\delta_p(g_1) - \cdots - z_j\delta_p(g_j) \in \mathcal{H}^{k_{p-1}}$$

This implies that each monomial in the power series expansion of each component of $z_{j+1}\delta_p(g_{j+1})$ is divisible by one of the germs $z_1, \ldots, z_j$ and it follows that the same thing is true of $\delta_p(g_{j+1})$ itself. Hence,

$$\delta_p(g_{j+1}) = z_1 g_1' + \cdots + z_j g_j'$$

for elements $g_1', \ldots, g_j' \in \mathcal{H}^{k_{p-1}}$. Since $\delta_{p-1}\delta_p = 0$ this implies that

$$\delta_p(g_{j+1}) \in K_{p-1} \cap I_j\mathcal{H}^{k_{p-1}}$$

Since $p - 1 \geq j$ it follows from the induction hypothesis that $\delta_p(g_{j+1}) \in I_j K_{p-1} = I_j\delta_p(\mathcal{H}^{k_p})$. In other words,

$$\delta_p(g_{j+1}) = z_1\delta_p(h_1) + \cdots + z_j\delta_p(h_j)$$

for elements $h_1, \ldots, h_j \in \mathcal{H}^{k_p}$. If we set

$$h_{j+1} = g_{j+1} - z_1 h_1 - \cdots - z_j h_j \in \mathcal{H}^{k_p}$$

then $\delta_p(h_{j+1}) = 0$ so that $h_{j+1} \in K_p$. Also,

$$f - z_{j+1}h_{j+1} = z_1(g_1 + z_{j+1}h_1) + \cdots + z_j(g_j + z_{j+1}h_j)$$

and, hence, $f - z_{j+1}h_{j+1} \in K_p \cap I_j\mathcal{H}^{k_p}$ which is $I_j K_p$ by the induction hypothesis. Hence, $f \in I_{j+1}K_p$ as required. this completes the induction and the proof.

3.14 Corollary. *If A is ${}_n\mathcal{H}$ or ${}_n\mathcal{O}$ then every finitely generated A module has a free resolution of the form:*

$$0 \to A^{k_n} \to A^{k_{n-1}} \to \cdots \to A^{k_0} \to M \to 0$$

3.15 Corollary. *If A is ${}_n\mathcal{H}$ or ${}_n\mathcal{O}$ then for every pair of A-modules M and N we have $tor_p^A(M,N) = 0 = ext_A^p(M,N)$ for $p > n$.*

3. Problems

1. Prove that if X is a vector space and $\hom(A,X)$ is given the obvious A-module structure, then $\hom(A,X)$ is an injective A-module. Show that every A module is a submodule of a module of this form. Then prove that a module is injective if and only if it is a direct summand of a module of the form $\hom(A,X)$.
2. Fix $\lambda \in \mathbb{C}$. Find a resolution of the form given in corollary 3.14 for the one dimensional $\mathbb{C}[z]$-module, $\mathbb{C}_\lambda$, on which each $p \in \mathbb{C}[z]$ acts as multiplication by the scalar $p(\lambda)$.
3. If V is any vector space and $L \in \mathrm{end}(V)$ is any linear transformation, then we can make V into a $\mathbb{C}[z]$-module V_L by letting $p \in \mathbb{C}[z]$ act on V as the linear transformation $p(L)$. Show that $\mathrm{tor}_1^A(\mathbb{C}_\lambda, V_L) = \ker\,(\lambda - L)$ and $\mathrm{tor}_0^A(\mathbb{C}_\lambda, V_L) = \mathrm{coker}\,(\lambda - L)$. Thus, $\lambda - L$ is invertible if and only if both of these tor groups vanish.
4. Prove a result analogous to the result of problem 3 but with tor replaced by ext.
5. Verify (if you have never done so before) that a short exact sequence of complexes yields a long exact sequence of homology.
6. Show that each non-zero element of $\mathrm{ext}_A^1(M,N)$ corresponds to a non-trivial extension of M by N, that is, to a short exact sequence

$$0 \to N \to Q \to M \to 0$$

which does not split.

4. Local Theory of Varieties

A subset V of $\mathbb{C}^n$ is a *subvariety* of $\mathbb{C}^n$ if for each point $\lambda \in V$ there is a neighborhood U_λ of λ and functions $f_1, \dots, f_k$ holomorphic in U_λ such that $V \cap U_\lambda = \{z \in U_\lambda : f_1(z) = \cdots = f_k(z) = 0\}$. Algebraic subvarieties are defined in the same way except that Zariski open sets are used and regular functions replace holomorphic functions. However, since we can always multiply through by the common denominator of a finite set of regular functions, we can always choose the functions $f_1, \dots, f_k$ that define the variety on U_λ to be polynomials in the case of an algebraic subvariety. If V and W are subvarieties of $\mathbb{C}^n$ and $V \subset W$ then we will also say that V is a subvariety of W. If V is also a closed subset of W then we will call it a closed subvariety of W.

Warning! In much of the several complex variables literature, subvarieties are required to be closed. Our terminology is consistent with the algebraic geometry literature.

Note that V is a *closed* subvariety of an open set $U \subset \mathbb{C}^n$ if and only if for each point $\lambda \in U$ there is a neighborhood U_λ of λ and functions $f_1, \dots, f_k$ holomorphic in U_λ such that $V \cap U_\lambda = \{z \in U_\lambda : f_1(z) = \cdots = f_k(z) = 0\}$.

It is obvious that finite unions and intersections of closed subvarieties of an open set U are also subvarieties of U. We may also define the *germ* of a closed subvariety of U at a point $\lambda \in U$. That is, we define two closed subvarieties V and W of U to be *equivalent at* λ if there is a neighborhood U_λ of λ such that $V \cap U_\lambda = W \cap U_\lambda$. The germ of a closed subvariety V at λ is then the equivalence class containing V. We will say V is the *germ of a variety* at $\lambda \in \mathbb{C}^n$ if V is the germ of some closed subvariety of some neighborhood of λ.

Given finitely many germs of varieties, $V_1, \cdots, V_k$ at λ, we may choose (by taking intersections, if necessary) a common neighborhood U of λ in which these germs have representatives. The germ of the intersection of these varieties is then well defined independent of the choice of U and the representatives of the V_i. We will call this the intersection, $V_1 \cap \cdots \cap V_k$, of the germs $V_1, \cdots, V_k$. The union of finitely many germs of varieties is defined in like manner, as is the relation " $\subset$ ".

4.1 Definition. *If V is the germ of a holomorphic (algebraic) variety at 0 then $\operatorname{id} V$ is defined to be the ideal of ${}_n\mathcal{H}$ consisting of all germs which vanish on V. On the other hand, if $\mathcal{I}$ is an ideal of ${}_n\mathcal{H}$ (${}_n\mathcal{O}$), then $\operatorname{loc} \mathcal{I}$ is defined to be the germ of the subvariety of a neighborhood of 0 defined by the vanishing of a finite set of generators for $\mathcal{I}$.*

One needs to check that these definitions make sense; that is, are $\operatorname{id} V$ and $\operatorname{loc} \mathcal{I}$ well defined? However, this is easy to do. The next theorem lists elementary properties of these ideas that follow directly from the definitions:

4.2 Theorem. *The following relationships hold between ideals of ${}_n\mathcal{H}$ (${}_n\mathcal{O}$) and germs of holomorphic (algebraic) subvarieties at zero.*

(a) $V_1 \subset V_2 \implies \operatorname{id} V_1 \supset \operatorname{id} V_2$.
(b) $\mathcal{I}_1 \subset \mathcal{I}_2 \implies \operatorname{loc} \mathcal{I}_1 \supset \operatorname{loc} \mathcal{I}_2$.
(c) $V = \operatorname{loc} \operatorname{id} V$.
(d) $\mathcal{I} \subset \operatorname{id} \operatorname{loc} \mathcal{I}$ *but they are not generally equal.*
(e) $\operatorname{id}(V_1 \cup V_2) = (\operatorname{id} V_1) \cap (\operatorname{id} V_2) \supset (\operatorname{id} V_1) \cdot (\operatorname{id} V_2)$.
(f) $\operatorname{id}(V_1 \cap V_2) \supset (\operatorname{id} V_1) + (\operatorname{id} V_2)$.

(g) $\operatorname{loc}(\mathcal{I}_1 \cdot \mathcal{I}_2) = \operatorname{loc}(\mathcal{I}_1 \cap \mathcal{I}_2) = \operatorname{loc}(\mathcal{I}_1) \cup \operatorname{loc}(\mathcal{I}_2)$.
(h) $\operatorname{loc}(\mathcal{I}_1 + \mathcal{I}_2) = \operatorname{loc}(\mathcal{I}_1) \cap \operatorname{loc}(\mathcal{I}_2)$.

A germ V of a variety is called *reducible* if $V = V_1 \cup V_2$ where V_1 and V_2 are both subvarieties contained properly in V. If V is not reducible then it is called irreducible.

Another elementary result which follows easily from the above is:

4.3 Theorem. *A germ of a subvariety V is irreducible if and only if* $\operatorname{id} V$ *is a prime ideal.*

4.4 Theorem. *A germ of a subvariety can be uniquely (up to order) written as $V_1 \cup \cdots \cup V_k$ where each V_i is an irreducible germ of a subvariety, each V_i is a proper subgerm of V and no V_i is contained in a distinct V_j for $i \neq j$.*

Proof. Suppose V is a germ of a subvariety and it cannot be written as a finite union of irreducible germs of subvarieties. Then V itself is not irreducible and so it can be written as $V_1 \cup V_2$ for two germs of subvarieties which are properly contained in V. Then at least one of V_1 and V_2 also fails to be a finite union of irreducibles. We assume without loss of generality that this is V_1 and then write V_1 as the union of two proper subvarieties. Continuing in this way will produce an infinite decreasing sequence of subvarieties each properly contained in the next. This means that the corresponding sequence of ideals of these subvarieties will form an infinite ascending chain of ideals of our Noetherian local ring (${}_n\mathcal{H}$ or ${}_n\mathcal{O}$), which is impossible. Thus we have proved by contradiction that every germ of a subvariety is a finite union of irreducibles. By deleting redundant subvarieties we may get the other conditions satisfied.

If
$$V = V_1 \cup \cdots \cup V_k = V_1' \cup \cdots \cup V_m'$$
are two nonredundant ways of writing V as a union of irreducibles then for each i, $V_i = (V_i \cap V_1') \cup \cdots \cup (V_i \cap V_m')$ which implies that V_i is contained in one of the V_j'. Likewise, each V_j' is contained in one of the V_i's. Since the decompositions are non-redundant we conclude that each V_i is equal to some V_j and vice-versa. In other words, the decomposition is unique up to order. This completes the proof.

Of particular interest are non-singular subvarieties. We shall describe these in the holomorphic case first where the structure is particularly simple.

If U and U' are open subsets of $\mathbb{C}^n$ then a *biholomorphic mapping* from U to U' is a holomorphic map with a holomorphic inverse. A *holomorphic submanifold* of an open set U is a relatively closed subset V such that for each point $\lambda \in V$ there is a neighborhood U_λ and a biholomorphic map $F : U_\lambda \to \Delta(0; r)$ onto some open polydisc in $\mathbb{C}^n$ such that $F(\lambda) = 0$ and
$$F(U_\lambda \cap V) = \{a \in \Delta(0; r) : z_{k+1}, \ldots, z_n = 0\}$$
for some integer k. The integer k is called the dimension of the submanifold at λ.

Thus, a submanifold is a subvariety that locally, up to biholomorphic equivalence, looks like a complex linear subspace. A germ of a holomorphic subvariety is called non-singular if it is the germ of a holomorphic submanifold. In order to characterize holomorphic submanifolds, we need the complex version of the familiar implicit mapping theorem from advanced calculus. If U is a domain in $\mathbb{C}^n$ and $F : U \to \mathbb{C}^m$ is a holomorphic map with coordinate functions $f_1, \dots, f_m$ then the *Jacobian* of F is the $m \times n$ matrix $J_F(z) = \left(\dfrac{\partial f_i}{\partial z_j}(z)\right)$.

4.5 Theorem (Implicit mapping theorem). *Let F be a holomorphic mapping as above and suppose $\lambda \in U$ and $F(\lambda) = 0$. Suppose also that the last m columns of $J_F(\lambda)$ form a non-singular $m \times m$ matrix. Then there is a polydisc*

$$\Delta(\lambda; r) = \Delta(\lambda'; r') \times \Delta(\lambda''; r'') \subset \mathbb{C}^{n-m} \times \mathbb{C}^m$$

and a holomorphic map $G : \Delta(\lambda'; r') \to \Delta(\lambda''; r'')$ such that $G(\lambda') = \lambda''$ and $F(z) = 0$ for $z = (z', z'') \in \Delta(\lambda; r)$ if and only if $G(z') = z''$.

Proof. When $m = 1$ this is the implicit function theorem which is a simple corollary of the Weierstrass preparation theorem in the case where the function is regular of degree one in its last variable. We prove the general case by induction on m. Thus, we assume that the result is true for $m - 1$ and proceed to prove it for m.

Let $J_F(\lambda) = (J'_F(\lambda), J''_F(\lambda))$ be the separation of $J_F(\lambda)$ into its first $n-m$ columns and its last m columns. The hypothesis is that $J''_F(\lambda)$ is non-singular. By a linear change of variables in the range space $\mathbb{C}^m$ we may assume that $J''_F(\lambda)$ is the $m \times m$ identity matrix. Then, since $\partial f_m / \partial z_n = 1$ at λ, it follows from the implicit function theorem that there is a polydisc $\Delta(\lambda; r)$ and a holomorphic mapping

$$h : \Delta(\lambda_1, \dots, \lambda_{n-1}; r_1, \dots, r_{n-1}) \to \Delta(\lambda_n, r_n)$$

such that $h(\lambda_1, \dots, \lambda_{n-1}) = \lambda_n$ and $f_m(z) = 0$ for $z \in \Delta(\lambda; r)$ exactly when $z_n = h(z_1, \dots, z_{n-1})$. Then we may define a holomorphic mapping

$$F' : \Delta(\lambda_1, \dots, \lambda_{n-1}; r_1, \dots, r_{n-1}) \to \mathbb{C}^{m-1}$$

by defining its coordinate functions $f'_1, \dots, f'_{m-1}$ to be

$$f'_i(z_1, \dots, z_{n-1}) = f_i(z_1, \dots, z_{n-1}, h(z_1, \dots, z_{n-1}))$$

Then $F'(\lambda_1, \dots, \lambda_{n-1}) = 0$ and the Jacobian $J_{F'}(\lambda)$ has the property that its last $m - 1$ columns form an $(m-1) \times (m-1)$ identity matrix. It follows from the induction hypothesis that, after possibly shrinking the polydisc, there is a holomorphic mapping

$$G' : \Delta(\lambda'; r') \to \Delta(\lambda_{n-m+1}, \dots, \lambda_{n-1}; r_{n-m+1}, \dots, r_{n-1})$$

such that $G'(\lambda') = (\lambda_{n-m+1}, \dots, \lambda_{n-1})$ and such that $F'(z_1, \dots, z_{n-1}) = 0$ for a point $(z_1, \dots, z_{n-1}) \in \Delta(\lambda_1, \dots, \lambda_{n-1}; r_1, \dots, r_{n-1})$ precisely when $(z_{n-m+1}, \dots, z_{n-1}) = G'(z')$. Since $F(z) = 0$ for $z \in \Delta(\lambda; r)$ precisely when $z_n = h(z_1, \dots, z_{n-1})$ and $F'(z_1, \dots, z_{n-1}) = 0$, the mapping

$$G(z') = (G'(z'), h(z', G'(z')))$$

has the required properties. This completes the proof.

4.6 Inverse mapping theorem. *If F is a holomorphic mapping from a neighborhood U of $\lambda \in \mathbb{C}^n$ into $\mathbb{C}^n$ and if $J_F(\lambda)$ is non-singular, then, on some possibly smaller neighborhood U' of λ, F is a biholomorphic mapping to some neighborhood of $F(\lambda)$.*

Proof. This follows immediately from the implicit mapping theorem applied to the mapping $H : \mathbb{C}^n \times U \to \mathbb{C}^n$ defined by $H(z', z'') = F(z'') - z'$.

4.7 Theorem. *If F is a holomorphic mapping from a domain U in $\mathbb{C}^n$ into $\mathbb{C}^m$ and if J_F has constant rank k in U, then for each point $\lambda \in U$ there is a neighborhood U_λ of λ in which F is biholomorphically equivalent to the mapping $(z_1, \ldots, z_n) \to (z_1, \ldots, z_k, 0, \ldots, 0)$ from a neighborhood of zero in $\mathbb{C}^n$ to a neighborhood of zero in $\mathbb{C}^m$. Thus, $\{z \in U : F(z) = 0\}$ is a submanifold of dimension $n - k$ in U_λ and, for each $\lambda \in U$, $F(U_\lambda)$ is a submanifold of dimension k in a neighborhood of $F(\lambda)$.*

Proof. We may assume that λ and $F(\lambda)$ are both the origin. After a linear change of coordinates, we may assume that the upper left hand $k \times k$ submatrix of $J_F(z)$ is nonsingular in a neighborhood U' of 0. We then define a new mapping G from U' into $\mathbb{C}^n$ by

$$G(z_1, \ldots, z_n) = (f_1(z_1, \ldots, z_n), \ldots, f_k(z_1, \ldots, z_n), z_{k+1}, \ldots, z_n)$$

Then J_G is non-singular in U' and thus G is a biholomorphic mapping of one neighborhood U'' of $0 \in \mathbb{C}^n$ to another. Then $F \circ G^{-1}$ has the form $(z_1, \ldots, z_k, f'_{k+1}, \ldots, f'_m)$. Since the Jacobian of $F \circ G^{-1}$ also has rank k throughout U'', it follows that the functions $\partial f'_j / \partial z_i$ vanish identically for $i > k$ and, thus, that the functions f'_j are functions of $z_1, \ldots, z_k$ alone. If we set

$$H(z_1, \ldots, z_m) = (z_1, \ldots, z_k, z_{k+1} - f'_{k+1}(z_1, \ldots, z_k), \ldots, z_m - f'_m(z_1, \ldots, z_k))$$

then H is a biholomorphic map on some neighborhood of $0 \in \mathbb{C}^m$ and

$$H \circ F \circ G^{-1}(z_1, \ldots, z_n) = (z_1, \ldots, z_k, 0, \ldots, 0)$$

on a neighborhood U of $0 \in \mathbb{C}^n$, as required. This completes the proof.

Note that for any holomorphic mapping F on a domain U, the set on which J_F has rank less than or equal to k is a subvariety of U since it is defined by the condition that the determinants of certain submatrices of J_F vanish. Let $V = \{z \in U : F(z) = 0\}$ and for each k let $V_k = \{z : \text{rank} J_F(z) \le k\}$ and let j be the largest integer k for which V_k is a proper subvariety of U. Then $U - V_j$ is an open dense set on which rank J_F is equal to $j + 1$. The set $V_0 = V \cap (U - V_j)$ consists of regular points of V and is a submanifold of dimension $n - j$ of $U - V_j$. It would be nice to know that V_0 is an open dense subset of V, since this would show that most points of a variety are regular. However, is quite possible for V and V_j to coincide, in which case, our discussion so far tells us nothing about the regular points of V. In this situation we have made a bad choice of a mapping F to define our subvariety. We need to be able to describe a variety as the set of common zeroes of a set of functions which are chosen in a way that gives us much more detailed information about the local structure of our subvariety at a point. We also need such a description for the proof of Hilbert's nullstellensatz. We obtain such an *optimal choice of functions defining a subvariety* in the next section and use it to prove the nullstellensatz and several other important facts concerning varieties.

We end this section with a brief discussion of the local ring of the germ of a variety.

4.8 Definition. *If V is a subvariety of $\mathbb{C}^n$ then a holomorphic (regular) function f on V is a complex valued function on V with the property that for each point $\lambda \in V$ there is a neighborhood U_λ of λ such that f extends to be holomorphic (regular) in U_λ. The algebra of functions holomorphic (regular) on V will be denoted ${}_V\mathcal{H}$ (${}_V\mathcal{O}$) while the algebra of germs at $\lambda \in V$ of functions holomorphic (regular) on neigborhoods in V of λ will be denoted ${}_V\mathcal{H}_\lambda$ (${}_V\mathcal{O}_\lambda$). The later is called the local ring of V at λ.*

The following is immediate from the definition:

4.9 Theorem. *If V is a holomorphic (algebraic) subvariety of $\mathbb{C}^n$ and V_λ is its germ at λ then ${}_V\mathcal{H}_\lambda = {}_n\mathcal{H}_\lambda / \operatorname{id} V_\lambda$ and this ring is a noetherian local ring. The analogous statement holds for ${}_V\mathcal{O}_\lambda$.*

4.10 Definition. *If V and W are holomorphic (algebraic) subvarieties of $\mathbb{C}^n$ and $\mathbb{C}^m$ and $F : V \to W$ is a mapping then F is called holomorphic (regular) if each of its coordinate functions is a holomorphic (regular) complex valued function on V. A holomorphic (regular) function with a holomorphic inverse is called biholomorphic (biregular).*

Again it is immediate from the definition that:

4.11 Theorem. *If V and W are holomorphic (algebraic) subvarieties of $\mathbb{C}^n$ and $\mathbb{C}^m$ and $F : V \to W$ is a mapping and $\lambda \in V$ then F is a holomorphic (regular) mapping on some neighborhood of λ if and only if $F^*(g) = g \circ F$ defines an algebra homomorphism F^* from ${}_W\mathcal{H}_{F(\lambda)}$ to ${}_V\mathcal{H}_\lambda$ (${}_W\mathcal{O}_{F(\lambda)}$ to ${}_V\mathcal{O}_\lambda$).*

A somewhat deeper result is the following:

4.12 Theorem. *If V and W are holomorphic subvarieties of $\mathbb{C}^n$ and $\mathbb{C}^m$, $\lambda \in V$, $\mu \in W$ and $\phi : {}_W\mathcal{H}_\mu \to {}_V\mathcal{H}_\lambda$ is any algebra homomorphism, then ϕ is induced as in Theorem 4.11 by a holomorphic mapping from a neighborhood of λ in V to a neighborhood of μ in W. The analogous result also holds for algebraic subvarieties.*

Proof. The algebra homomorphism ϕ maps units to units. However, every element f in either local ring has the property that there is a unique complex number c such that $f - c$ is a non-unit. It follows that ϕ also maps non-units to non-units, ie. f is in the maximal ideal of ${}_V\mathcal{H}_\lambda$ if and only if $\phi(f)$ is in the maximal ideal of ${}_W\mathcal{H}_\mu$.

If $w_1, \dots, w_m$ are the germs in ${}_W\mathcal{H}_\mu$ of the restrictions of the coordinate functions in $\mathbb{C}^m$ to W then $\phi(w_1), \dots, \phi(w_m)$ are germs in ${}_V\mathcal{H}_\lambda$ and thus are represented by functions which extend to be holomorphic in some neighborhood of λ in $\mathbb{C}^n$. Let $f_1, \dots, f_m$ be such holomorphic extensions. Then they are the coordinate functions of a holomorphic map F from a neighborhood of λ in $\mathbb{C}^n$ into $\mathbb{C}^m$. Since, for each j, $\phi(w_j - w_j(\mu)) = (f_j)|_V - w_j(\mu)$ which belongs to the maximal ideal of ${}_W\mathcal{H}_\mu$, it follows that $f_j(\lambda) = w_j(\mu)$ and, hence, that $F(\lambda) = \mu$. Thus, the mapping F induces an algebra homomorphism $F^* : {}_m\mathcal{H}_\mu \to {}_n\mathcal{H}_\lambda$. If we follow this by restriction to V then we have a homomorphism $\tilde{F}^* : {}_m\mathcal{H}_\mu \to {}_V\mathcal{H}_\lambda$. On the other hand, if we precede ϕ by restriction from $\mathbb{C}^m$ to W, then it also determines a homomorphism $\tilde{\phi} : {}_m\mathcal{H}_\mu \to {}_V\mathcal{H}_\lambda$. Clearly the proof will be complete if we can show that $\tilde{\phi} = \tilde{F}^*$ and that F maps some neighborhood of λ in V into W.

Note that, by the construction of F, $\tilde{F}^*(w_j) = \tilde{\phi}(w_j)$ for each j. This implies that the two homomorphisms agree on polynomials. Since, for each k every element of ${}_m\mathcal{H}_\mu$ is a polynomial of degree k plus an element of the k^{th} power of the maximal ideal of ${}_m\mathcal{H}_\mu$, and since each of $\tilde{F}^*$ and $\tilde{\phi}$ maps the k^{th} power of the maximal ideal of ${}_m\mathcal{H}_\mu$ to the k^{th} power of the maximal ideal of ${}_V\mathcal{H}_\lambda$, we conclude that $\tilde{F}^*(f) - \tilde{\phi}(f)$ belongs to the k^{th} power of the maximal ideal of ${}_V\mathcal{H}_\lambda$ for every $f \in {}_m\mathcal{H}_\mu$ and every positive integer k. However, by Nakayama's lemma the intersection of all powers of the maximal ideal in a Noetherian local ring is zero. Therefore, $\tilde{F}^* = \tilde{\phi}$.

It remains to prove that F maps a neighborhood of λ in V into W. However, $\operatorname{id} W$ is in the kernel of $\tilde{\phi} = \tilde{F}^*$ by the definition of $\tilde{\phi}$. Hence, for every $f \in \operatorname{id} W$ we have $(f \circ F)|_V = \tilde{F}^*(f) = 0$, ie. f vanishes on $F(V)$. If we apply this fact to a finite set of generators of $\operatorname{id} W$ we conclude that a suitably small neighborhood of λ in V is mapped by F into W.

Two germs of varieties are said to be biholomorphically equivalent if there is a biholomorphic map between suitable representative neighborhoods. In view of the preceding theorem we have:

4.13 Corollary. *Two germs of varieties are biholomorphically equivalent if and only if their local rings are isomorphic as algebras.*

4. Problems

1. Give an example which shows that the implicit function theorem, the inverse mapping theorem and the Weierstrass preparation theorem fail in the algebraic case.
2. Consider the polynomial on $\mathbb{C}^2$ defined by $p(z, w) = z^2 - w^3$. Prove that p is irreducible in both ${}_2\mathcal{H}$ and ${}_2\mathcal{O}$ and, hence, generates a prime ideal in each algebra.
3. Show that in either ${}_2\mathcal{H}$ or ${}_2\mathcal{O}$ the ideal generated by the polynomial p in problem 2 is $\operatorname{id} V$ where $V = \{(z, w) \in \mathbb{C}^2 : p(z, w) = 0$.
4. With V as in problem 2, show that there are irreducible elements $f, g \in {}_V\mathcal{H}$ (${}_V\mathcal{O}$) such that $f^2 = g^3$. Conclude that these algebras are not unique factorization domains.
5. With V as in problem 2, let M be the maximal ideal of ${}_V\mathcal{H}$. Show that M is generated by two of its elements in such a way the the resulting morphism ${}_V\mathcal{H}^2 \to M$ has kernel isomorphic to M as a ${}_V\mathcal{H}$ -module. Conclude that the Hilbert syzygy theorem fails to hold for ${}_V\mathcal{H}$.

5. The Nullstellensatz

In rings A for which it is true, the Nullstellensatz says that

$$\operatorname{id}\operatorname{loc} I = \sqrt{I} = \{x \in A : x^n \in I \text{ for some n }\}$$

for each ideal I. This is true for the rings ${}_n\mathcal{H}$ and ${}_n\mathcal{O}$ as well as for the ring of polynomials $\mathbb{C}[z_1, \dots, z_n]$. In each case, it is trivial from the definition that $\sqrt{I} \subset \operatorname{id}\operatorname{loc} I$. Also in each case, the theorem can easily be reduced to the case of prime ideals. In fact, a primary ideal is an ideal I whose radical $\sqrt{I}$ is prime and, in a Noetherian ring, each ideal I has a primary decomposition $I = \cap_{j=1}^m I_j$. Thus, if we assume the Nullstellensatz for prime ideals, then we have

$$\operatorname{id}\operatorname{loc} I = \operatorname{id}(\bigcup_{j=1}^m \operatorname{loc} I_j) = \cap_{j=1}^m \operatorname{id}\operatorname{loc} I_j \subset \cap_{j=1}^m \operatorname{id}\operatorname{loc} \sqrt{I_j} = \cap_{j=1}^m \sqrt{I_j} = \sqrt{I}$$

which is the Nullstellensatz for general ideals since we already have the reverse containment.

Our goal in this section is to prove the Nullstellensatz for the ring ${}_n\mathcal{H}$. However, we first prove the much easier result that the Nullstellensatz holds for ${}_n\mathcal{O}$. It is easy to see that the Nullstellensatz for $\mathbb{C}[z_1, \dots, z_n]$ implies the Nullstellensatz for ${}_n\mathcal{O}$ (this is left as an exercise) so we shall prove the Nullstellensatz for $\mathbb{C}[z_1, \dots, z_n]$. We first need a couple of lemmas from commutative algebra.

5.1 Lemma. *Let $A \subset B \subset C$ be rings with A Noetherian. Suppose that C is a finitely generated A-algebra and C is integral over B. Then B is also a finitely generated A algebra.*

Proof. We first show that C is actually a finitely generated B-module. The fact that C is finitely generated as an A-algebra means that it is also finitely generated as a B-algebra and, hence, that every element of C is a polynomial in a finite set of generators $c_1, \dots, c_k$ with coeficients in B. However, the fact that C is integral over B means that for each $c \in C$ there is an integer n_c such that every polynomial in c is equal to one of degree less than or equal to n_c. Thus, the algebra generated over B by c_1 is a finitely generated B-module. An induction argument on the number of generators now shows that C is a finitely generated B-module as claimed.

Now let $x_1, \dots, x_l$ generate C as an A-algebra and $y_1, \dots, y_m$ generate C as a B-module. Then there exist elements b_{ij} and b_{ijk} in B such that

$$x_i = \sum_j b_{ij} y_j, \quad y_i y_j = \sum_k b_{ijk} y_k$$

Let B_0 be the algebra generated over A by the b_{ij} and the b_{ijk}. We have $A \subset B_0 \subset B$ and B_0 is Noetherian since A is Noetherian and B_0 is a quotient of a polynomial ring over A. The above equations show that each element of C is a linear combination of the y_i with coeficients from B_0 so that C is a finitely generated B_0-module. Since B_0 is Noetherian and B is a submodule of C it follows that B is also a finitely generated B_0-module. Since B_0 is finitely generated as an A-algebra it follows that B is finitely generated as an A-algebra. This completes the proof.

5.2 Lemma (weak Nullstellensatz). *Let k be an algebraically closed field and A a finitely generated k-algebra. Then for each maximal ideal M of A, the inclusion of k into A/M is an isomorphism.*

Proof. Let $x_1, \dots, x_n$ be generators of A/M over k and assume that they are chosen such that $x_1, \dots, x_j$ are algebraically independent over k and each of the others is algebraic over the field $F = k(x_1, \dots, x_j)$. Thus, A/M is a finite algebraic field extension of F and, hence, a finitely generated F-module. It follows from the previous lemma applied to $k \subset F \subset A/M$ that F is a finitely generated k-algebra - that is, there is a finite set of fractions $y_i = f_i/g_i$, with $f_i, g_i \in k[x_1, \dots, x_j]$ such that every element of F is a polynomial in the y_i. Now if the set $\{x_1, \dots, x_j\}$ is not the empty set then there is an irreducible polynomial $h \in k[x_1, \dots, x_j]$ prime to all the g_i's since the g_i's have no common factors with $1 + \prod_i g_i$. However, $h^{-1} \in F$ and so it must be a polynomial in the y_i's. This implies that $h^{-1}k \in k[x_1, \dots, x_j]$ where k is some product of the $g_i's$, which is impossible if h is relatively prime to all the $g_i's$ in $k[x_1, \dots, x_j]$. The resulting contradiction shows that $\{x_1, \dots, x_j\}$ must be empty. But in this case A/M is algebraic over k and, hence, equal to k since k is algebraically closed.

5.3 Theorem (Nullstellensatz for polynomial algebras). *If $A = \mathbb{C}[z_1, \dots, z_n]$ then for each ideal $I \subset A$ we have*

$$\sqrt{I} = \operatorname{id} \operatorname{loc} I$$

Proof. We may assume I is prime. We know that $I \subset \operatorname{id} \operatorname{loc} I$ and so we need only prove the reverse containment. Let $f \in A$ be any element not in I and let $B = A/I$ and $C = B_f$, the algebra of fractions over B with denominators that are powers of the image of f in B. Now let M be a maximal ideal of C. Since C is finitely generated over $\mathbb{C}$ it follows from the previous lemma that $C/M = \mathbb{C}$. Then the images of $z_1, \dots, z_n$ in C/M are the coordinates of a point $\lambda \in \mathbb{C}^n$. It is clearly a point in $\operatorname{loc} I$ (since the maximal ideal it determines contains I by construction) and a point at which $f(\lambda) \neq 0$ (since f is invertible in C). Thus, f is not in $\operatorname{id} \operatorname{loc} I$ and the proof is complete.

The above proof depended heavily on the fact that quotients of $\mathbb{C}[z_1, \dots, z_n]$ are finitely generated algebras over the ground field. We have no such finite generation conditions in the holomorphic case and must use entirely different methods. The proof of the Nullstellensatz for ${}_n\mathcal{H}$ depends on a fairly precise description of the locus of a prime ideal as the germ of a certain kind of finite branched cover of a neighborhood in $\mathbb{C}^m$. We now proceed to develop this description.

In what follows we will be making fairly heavy use of field theory. Since this discussion will not make much sense to someone who doesn't know a certain amount of field theory, we present below a list of facts from this subject that we will use implicitly or explicitly. This is presented as a study guide for those who need to brush up on the subject. Proofs can be found in any algebra book with a good treatment of field theory (eg: Hungerford).

If $K \subset F$ are fields and $x \in F$ then x is called *algebraic* over K if it is the root of a polynomial with coeficients in K. The field F is called an *algebraic field extension* of K if each of its elements is algebraic over K.

F1 Theorem. *The following are equivalent for fields $K \subset F$:*

(a) *F is an algebraic field extension of K and is finitely generated over K;*
(b) *F is generated over K by finitely many elements which are algebraic over K;*
(c) *F is a finite dimensional vector space over K.*

F2 Theorem of the primitive element. *If K is an infinite field and F is a finitely generated algebraic extension of K then F is generated by a single element which may be chosen to be a linear combination of any given set of generators.*

The subfield of F generated by $x_1, \dots, x_n$ is denoted $K(x_1, \dots, x_n)$.

F3 Theorem. *If $x \in F$ is algebraic over K then x is a root of a unique irreducible monic polynomial with coefficients in K. If the degree of that polynomial is n, then $\{x^{n-1}, \dots, x, 1\}$ is a basis for $K(x)$ as a vector space over K.*

The unique irreducible monic polynomial having x as a root is called the *minimal polynomial* of the element x.

If p is a polynomial with coeficients in K then a *splitting field* F for p is an extension field of K which is generated over K by the roots of p and in which p factors as a product of linear factors (so that all possible roots of p are included in F).

F4 Theorem. *If p is a polynomial with coeficients in K then there is a splitting field for p which is unique up to isomorphism. If p is the minimal polynomial of an algebraic element x over K, then a splitting field for p may be chosen which is a field extension of $K(x)$.*

The Galois group of a field extension $K \subset F$ is the group of automorphisms of F which leave all elements of K fixed. The extension is called a Galois extension if K is exactly the set of elements fixed by the Galois group.

F5 Theorem. *If F is a splitting field for some polynomial p with coeficients in K and if K has characteristic zero, then F is a Galois extension of K. Every element of the Galois group of such an extension is uniquely determined by a permutation of the roots of p. If p is irreducible then the Galois group acts transitively on the roots.*

If p is a polynomial with coeficients in K then the *discriminant* of p is the product

$$d = \prod_{i \neq j} (x_i - x_j)$$

where $x_1, \dots, x_n$ are the roots of p in some splitting field.

F6 Theorem. *If p is a polynomial with coeficients in K then its discriminant is a well defined element of K and it is non-zero if and only if p has no multiple roots.*

F7 Theorem. *If p is an irreducible polynomial with coeficients in K and if K has characteristic zero then p has no multiple roots and, thus, has non-vanishing discriminant.*

F8 Theorem. *If p is a polynomial with coeficients in K and if its roots in some splitting field are $x_1, \dots, x_n$ then the discriminant of p is the square of the Vandermonde determinant*

$$\begin{vmatrix} x_1^{n-1} & x_1^{n-2} & \dots & x_1 & 1 \\ x_2^{n-1} & x_2^{n-2} & \dots & x_2 & 1 \\ \cdot & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \dots & \cdot & \cdot \\ x_n^{n-1} & x_n^{n-2} & \dots & x_n & 1 \end{vmatrix}$$

A ring is said to be *integrally closed* if it is integrally closed in its quotient field. That is, if whenever a monic polynomial with coeficients in the ring has a root in the quotient field of the ring then that root actually lies in the ring.

F9 Theorem. *A unique factorization domain is integrally closed.*

The next result may be less well known than its predecessors and so we shall actually prove it.

F10 Theorem. *If A is a unique factorization domain with quotient field K of characteristic zero and if $F = K(x)$ is the field extension generated by an element x, integral over A, with minimal polynomial p of degree n, then every element of F which is integral over A belongs to the A submodule of F generated by the elements $\frac{x^{n-1}}{d}, \dots, \frac{x}{d}, \frac{1}{d}$, where d is the discriminant of p.*

Proof. For an element $f(x) \in K(x)$, integral over A, we wish to find coeficients $a_0, \dots, a_{n-1}$ such that

$$a_{n-1}x^{n-1} + a_{n-2}x^{n-2} + \dots + a_1 x + a_0 = d \cdot f(x)$$

Let $x_1 = x$ and let $x_2, \dots, x_n$ be the other roots of p in a splitting field for p. We may then write down a system of n equations, each of which is a copy of the one above but with x replaced by x_j in the j^{th} equation. If we consider this as a system of equations in which the unknowns are the elements $a_1, \dots, a_n$, Kramer's rule give as solution

$$a_j = d \begin{vmatrix} x_1^{n-1} & x_1^{n-2} & \dots & x_1 & 1 \\ x_2^{n-1} & x_2^{n-2} & \dots & x_2 & 1 \\ \cdot & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \dots & \cdot & \cdot \\ x_n^{n-1} & x_n^{n-2} & \dots & x_n & 1 \end{vmatrix}^{-1} \begin{vmatrix} x_1^{n-1} & x_1^{n-2} & \dots & f(x_1) & \dots & x_1 & 1 \\ x_2^{n-1} & x_2^{n-2} & \dots & f(x_2) & \dots & x_2 & 1 \\ \cdot & \cdot & \dots & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \dots & \cdot & \dots & \cdot & \cdot \\ x_n^{n-1} & x_n^{n-2} & \dots & f(x_n) & \dots & x_n & 1 \end{vmatrix}$$

where in the second determinant the $f(x_i)$ replace the j^{th} column of the first determinant. Now, of course, the first determinant is the Vandermonde determinant which has square equal to d by F8. Thus, a_j is the product of the Vandermonde and the determinant obtained from the Vandermonde by replacing its j^{th} column with the column formed by the $f(x_i)$. Clearly this product is left fixed by any permutation of the roots $x_1, \dots, x_n$ since this just amounts to applying the same permutation to the rows in both matrices. Thus, the elements a_j so determined are fixed by the Galois group of the splitting field of p and belong to K by F5. However, since x and $f(x)$ are integral over A so are all the x_i

and $f(x_i)$, again by F5. It follows that the a_j are also integral over A since they lie in the ring generated by the x_i and $f(x_i)$. However, A is integrally closed in its quotient field K by F9 and, hence, $a_j \in A$ for $j = 1, \ldots, n$. This completes the proof.

We now return to our study of a prime ideal $P \subset {}_n\mathcal{H}$.

Recall that a holomorphic function f in a neighborhood of 0 is called regular in z_n if $f(0, \ldots, 0, z_n)$ is not identically zero. In what follows, we consider ${}_j\mathcal{H}$ for $j \le n$ to be the subring of ${}_n\mathcal{H}$ consisting of functions that depend only on the first j variables.

5.4 Definition. *An ideal $I \subset {}_n\mathcal{H}$ is called regular in the variables $z_{m+1}, \ldots, z_n$ if ${}_m\mathcal{H} \cap I = 0$ and for each $j \in \{m+1, \ldots, n\}$ there is an element $f_j \in {}_j\mathcal{H} \cap I$ which is regular in z_j.*

5.5 Lemma. *For each non-zero ideal $I \subset {}_n\mathcal{H}$ there is a choice of a complex linear coordinate system for $\mathbb{C}^n$ and an $m < n$ such that I is regular in the variables $z_{m+1}, \ldots, z_n$.*

Proof. We can choose a non-zero $f_n \in I$ and then by a suitable linear change of coordinates arrange that f_n is regular in z_n. Suppose we have chosen $f_{j+1}, \ldots, f_n$ satisfying the conditions of Definition 5.4. Then either ${}_j\mathcal{H} \cap I = 0$, in which case we are through, or there is a nonzero $f_j \in {}_j\mathcal{H} \cap I$. The function f_j can be made regular in z_j by a linear change of coordinates that involves only the first j coordinates and, hence, does not effect the regularity of the functions chosen previously. The Lemma follows by induction.

The notion of an ideal being regular in the variables $z_{m+1}, \ldots, z_n$ seems to depend on the ordering of these variables. However, the next lemma shows that it depends only on the decomposition $\mathbb{C}^n = \mathbb{C}^m \times \mathbb{C}^{n-m}$ and not on the choice of coordinate systems within the two factors.

5.6 Lemma. *An ideal $I \subset {}_n\mathcal{H}$ is regular in the variables $z_{m+1}, \ldots, z_n$ if and only if ${}_m\mathcal{H}$ is isomorphic to its image ${}_m\tilde{\mathcal{H}}$ in ${}_n\tilde{\mathcal{H}} = {}_n\mathcal{H}/I$ and ${}_n\tilde{\mathcal{H}}$ is an integral algebraic extension of ${}_m\tilde{\mathcal{H}}$ generated by the images of $z_{m+1}, \ldots, z_n$ in ${}_n\mathcal{H}/I$.*

Proof. Obviously we have that ${}_m\mathcal{H} \cap I = 0$ if and only if the natural map of ${}_m\mathcal{H}$ to ${}_m\tilde{\mathcal{H}}$ is an isomorphism. Thus, we need to prove that the existence of the functions f_j in Definition 5.4 is equivalent to the fact that the images of $z_{m+1}, \ldots, z_n$ in ${}_n\tilde{\mathcal{H}}$ are all integral over ${}_m\tilde{\mathcal{H}}$. However, if the image $\tilde{z}_j$ of z_j in ${}_n\tilde{\mathcal{H}}$ is integral over ${}_m\tilde{\mathcal{H}}$ then it is a root of a monic polynomial with coeficients in ${}_m\tilde{\mathcal{H}}$. Lifting these coeficients to ${}_m\mathcal{H}$ results in a Weierstrass polynomial in ${}_j\mathcal{H} \cap I$ which will serve as the element f_j. Now on the other hand, if there are $f_j \in {}_j\mathcal{H} \cap I$, regular in z_j, then we may assume that they are Weierstrass polynomials by the Weierstrass preparation theorem. Since these polynomials belong to I this implies that the image $\tilde{z}_j$ of each z_j in ${}_n\tilde{\mathcal{H}}$ for $j > m$ is integral over ${}_{j-1}\tilde{\mathcal{H}}$. Furthermore, by the Weierstrass division theorem, for any $f \in {}_j\mathcal{H}$ there are elements $g \in {}_j\mathcal{H}$ and $r \in {}_{j-1}\mathcal{H}[z_j]$ such that $f = f_j g + r$. On passing to residue classes mod I and remembering that $f_j \in I$, we conclude that every element of ${}_j\tilde{\mathcal{H}}$ is a polynomial in $\tilde{z}_j$ over the ring ${}_{j-1}\tilde{\mathcal{H}}$. Thus, ${}_j\tilde{\mathcal{H}}$ is the integral algebraic extension of the ring ${}_{j-1}\tilde{\mathcal{H}}$ by the element $\tilde{z}_j$. We now have that ${}_n\tilde{\mathcal{H}}$ is obtained from ${}_m\tilde{\mathcal{H}}$ by successive integral algebraic extensions by the elements $\tilde{z}_j$ for $j = m+1, \ldots, n$. The theorem of transitivity of integral dependence now implies that ${}_n\tilde{\mathcal{H}} = {}_m\tilde{\mathcal{H}}[\tilde{z}_{m+1}, \ldots, \tilde{z}_n]$ where each of the elements $\tilde{z}_j$ is integral over ${}_m\tilde{\mathcal{H}}$. This completes the proof.

We can do even better in the case that our ideal, regular in $z_{m+1},\dots,z_n$, is a prime ideal P. In this case, ${}_n\tilde{\mathcal{H}} = {}_n\mathcal{H}/P$ is an integral domain and has a field of quotients ${}_n\tilde{\mathcal{M}}$. Also ${}_m\tilde{\mathcal{H}} \simeq {}_m\mathcal{H}$ have fields of quotients ${}_m\tilde{\mathcal{M}} \simeq {}_m\mathcal{M}$ and ${}_n\tilde{\mathcal{M}}$ is an algebraic field extension of ${}_m\tilde{\mathcal{M}}$ by the elements $\tilde{z}_{m+1},\dots,\tilde{z}_n$.

5.7 Definition. *A prime ideal is called strictly regular in the variables $z_{m+1},\dots,z_n$ if it is regular in these variables and ${}_n\tilde{\mathcal{M}}$ is generated over ${}_m\tilde{\mathcal{M}}$ by the single element $\tilde{z}_{m+1}$.*

5.8 Lemma. *Any non-zero prime ideal P can be made strictly regular in some set of variables $z_{m+1},\dots,z_n$ by a linear change of coordinates.*

Proof. By a linear change of coordinates we can make P regular in some set of variables $z_{m+1},\dots,z_n$. Then, as noted above, ${}_n\tilde{\mathcal{M}}$ is an algebraic field extension of ${}_m\tilde{\mathcal{M}}$ by the elements $z_{m+1},\dots,z_n$. By the theorem of the primitive element, ${}_n\tilde{\mathcal{M}}$ is actually generated over ${}_m\tilde{\mathcal{M}}$ by a single element which may be chosen to be a linear combination of $z_{m+1},\dots,z_n$. Another linear change of variables effecting only these coordinates can be used to transform this element into z_{m+1}. Such a change of variables does not change the fact that P is regular in the variables $z_{m+1},\dots,z_n$ by Lemma 5.6. This completes the proof.

Let P be a prime ideal of ${}_n\mathcal{H}$. We will choose a particularly nice finite set of elements of P which determine P in a certain fashion although they are not quite a set of generators. We assume that coordinates have been chosen so that P is strictly regular in the variables $z_{m+1},\dots,z_n$. Then the $\tilde{z}_j$ are all integral over ${}_m\tilde{\mathcal{H}}$. For each j we let p_j be the minimal polynomial of $\tilde{z}_j$. A priori, p_j has coeficients in the quotient field ${}_m\tilde{\mathcal{M}}$ of ${}_m\tilde{\mathcal{H}}$. However, all roots in a splitting field of p_j are also integral over ${}_m\tilde{\mathcal{H}}$ by the transitivity of the Galois group and, hence, the coeficients of p_j, being elementary symmetric functions of the roots, are integral over ${}_m\tilde{\mathcal{H}}$. Since the latter algebra is a unique factorization domain and, hence, integrally closed, we conclude that the coeficients of each p_j are actually in ${}_m\tilde{\mathcal{H}}$ and , hence, may be considered elements of ${}_m\mathcal{H}$. Then each polynomial $p_j(z_j)$ is an element of ${}_m\mathcal{H}[z_j] \subset {}_n\mathcal{H}$ which belongs to P since it vanishes mod P. The $p_j(z_j)$ are some of the elements of P that we are seeking.

We use F10 to choose the remaining elements. For $j = m+2,\dots n$ the image $\tilde{z}_j$ of z_j in ${}_n\tilde{\mathcal{M}} = {}_m\tilde{\mathcal{M}}(\tilde{z}_{m+1})$ is integral over ${}_m\tilde{\mathcal{H}}$ and, thus by F10, $d \cdot \tilde{z}_j = s_j(\tilde{z}_{m+1})$ where d is the discriminant of p_{m+1} and s_j is a unique polynomial of degree less than the degree of p_{m+1} with coeficients in ${}_m\tilde{\mathcal{H}} \simeq {}_m\mathcal{H}$. Then $q_j = d \cdot z_j - s_j(z_{m+1})$ belongs to ${}_m\mathcal{H}[z_{m+1}, z_j] \cap P$. The q_j together with the p_j are the elements we need to adequately describe P and $\operatorname{loc} P$.

We introduce two sub ideals of P that will play a role in what follows. We let I denote the ideal generated by $p_{m+1},\dots,p_n, q_{m+2},\dots,q_n$ and I' denote the ideal generated by $p_{m+1}, q_{m+2},\dots,q_n$.

Lemma 5.9. *With d the discriminant of p_{m+1}, as above, there is an integer ν such that for any $f \in {}_n\mathcal{H}$ there is a polynomial $r \in {}_m\mathcal{H}[z_{m+1}]$ with degree less than* degree p_{m+1} *such that $d^\nu \cdot f - r \in I$. Furthermore,*

$$d^\nu \cdot P \subset I \subset P \quad \text{and} \quad d^\nu \cdot I \subset I' \subset I$$

Proof. The polynomial p_n is a Weierstrass polynomial and so, for any $f \in {}_n\mathcal{H}$, the Weierstrass division theorem allows us to write $f = p_n g_n + r_n$ for some $g_n \in {}_n\mathcal{H}$ and some $r_n \in {}_{n-1}\mathcal{H}[z_n]$ of degree less than the degree of p_n. We now apply the Weierstrass division theorem to each coeficient of r_n with the divisor being the Weierstrass polynomial $p_{n-1} \in {}_m\mathcal{H}[z_{n-1}] \subset {}_{n-2}\mathcal{H}[z_{n-1}]$. If we gather together the terms this yields

$$f = p_n g_n + p_{n-1} g_{n-1} + r_{n-1}$$

with $g_j \in {}_n\mathcal{H}$ and $r_{n-1} \in {}_{n-2}\mathcal{H}[z_{n-1}, z_n]$. By repeating this argument as long as we have p_j's we eventually get

$$f = p_n g_n + \cdots + p_{m+1} g_{m+1} + r_{m+1}$$

with $g_j \in {}_n\mathcal{H}$ and $r_{m+1} \in {}_m\mathcal{H}[z_{m+1}, \ldots, z_n]$. Also note that the degree of r_j is less than the degree of p_j for $j = m+1, \ldots, n$.

Now for each $j = m+1, \ldots, n$ we have $d \cdot z_j = q_j + s_j$ and so, by the binomial theorem, $d^k \cdot z_j^k = h_{jk} q_j + s_j^k$ for each integer k and some $h_{jk} \in {}_n\mathcal{H}$. If we apply this to each power of each z_j appearing in r_{m+1} and if we choose $\nu = \sum(\text{degree } p_j - 1)$ then we conclude that

$$d^\nu \cdot r_{m+1} = q_{m+2} h_{m+2} + \cdots + q_n h_n + r_0$$

for some $h_j \in {}_n\mathcal{H}$ and $r_0 \in {}_m\mathcal{H}[z_{m+1}]$. Another application of the Weierstrass division theorem gives us $r_0 = p_{m+1} h_{m+1} + r$ where $h_{m+1} \in {}_n\mathcal{H}$ and $r \in {}_m\mathcal{H}[z_{m+1}]$ with degree of r less than the degree of p_{m+1}. Finally, this gives us

$$d^\nu \cdot f = \sum_{j=m+1}^{n} d^\nu g_j p_j + \sum_{j=m+2}^{n} h_j q_j + h_{m+1} p_{m+1} + r$$

which implies $d^\nu \cdot f - r \in I$ as required.

Obviously $I \subset P$. Now suppose that $f \in P$ and apply the above conclusion to f. Then $r \in P$ and so $r = 0$ since it has degree less than that of p_{m+1} which is the polynomial of lowest degree in $P \cap {}_m\mathcal{H}[z_{m+1}]$. Thus, $d^\nu \cdot f \in I$ and we conclude that $d^\nu \cdot P \subset I$.

Since $I' \subset I$ is obvious, to finish the proof we need to prove that $d^\nu \cdot I \subset I'$ or, in other words, that $d^\nu \cdot p_j \in I'$ for $j = m+1, \ldots, n$. To show this, we again apply the binomial theorem to the expression $(d \cdot z_j)^k = (q_j + s_j)^k$ for each power of z_j appearing in the polynomial p_j to obtain $d^\nu \cdot p_j = h_j q_j + r'$ where $r' \in {}_m\mathcal{H}[z_{m+1}]$. Another application of the Weierstrass division theorem gives us $r' = h_{m+1} p_{m+1} + r$ where $r \in {}_m\mathcal{H}[z_{m+1}]$ and degree of r is less than the degree of p_{m+1}. As before it follows that $r \in P$ and , Hence, $r = 0$ since p_{m+1} has minimal degree for an element of $P \cap {}_m\mathcal{H}[z_{m+1}]$. It follows that $d^\nu \cdot p_j = h_j q_j + h_{m+1} p_{m+1} \in I'$ as required. This completes the proof.

This has the following as an immediate corollary:

5.10 Corollary. *If D is the locus of the ideal in ${}_n\mathcal{H}$ generated by d then*

(i) $\operatorname{loc} P \subset \operatorname{loc} I \subset \operatorname{loc} I'$;
(ii) $D \cup \operatorname{loc} P = D \cup \operatorname{loc} I = D \cup \operatorname{loc} I'$.

5.11 Definition. *Let V and W be closed holomorphic subvarieties of open subsets of $\mathbb{C}^n$ and $\mathbb{C}^m$, respectively and let $\pi : V \to W$ be a finite to one, proper holomorphic map. Then π is said to be a finite branched holomorphic cover if there are dense open subsets $W_0 \subset W$ and $V_0 = \pi^{-1}(W_0) \subset V$ such that $W - W_0$ is a subvariety of W and $\pi : V_0 \to W_0$ is a locally biholomorphic mapping.*

Let $\pi : V \to W$ is a finite branched holomorphic cover, as above, and let p be a point of W_0. Let $\pi^{-1}(p) = \{q_1, \cdots, q_k\}$ and suppose we have chosen for each i a neighborhood U_i of q_i in V_0 such that $U_i \cap U_j = \emptyset$ for $i \neq j$. If A is a neighborhood of p with compact closure in W_0, then $\pi^{-1}(\bar{A}) - \bigcup U_i$ is a compact subset of V_0. The collection of sets of this form is closed under finite intersection and so, if they are all non-empty, then there is a point q in their intersection. Then, necessarily, $\pi(q) = p$ and, hence, q be one of the q_i. This is not possible since q is in the complement of each U_i. It follows that for some choice of A the set $\pi^{-1}(A)$ is contained in the disjoint union of the open sets $\pi^{-1}(A) \cap U_i$. Now suppose that p is a point of W_0 and the U_i have been chosen so that $\pi|_{U_i}$ is a biholomorphic map onto a neighborhood of p for each i. If A is chosen to be a subset of $\bigcap \pi(U_i)$, then $\pi^{-1}(A)$ is the disjoint union of the sets $\pi^{-1}(A) \cap U_i$ and, for each i, π is a biholomorphic map of $\pi^{-1}(A) \cap U_i$ onto A. Thus, each point of W_0 has a neighborhood which, under π, is covered by a finite number of biholomorphic copies of itself. A map with this property is a finite holomorphic covering map. In particular, it is a finite covering map in the topological sense.

Note that the number of points in the inverse image of a point of W_0 is locally constant on W_0. If W_0 is connected then this number is a constant r and, in this case, we say that $\pi : V \to W$ is a finite branched holomorphic cover of pure order r and $\pi : V_0 \to W_0$ is a finite holomorphic covering map of pure order r.

Now suppose W_0 is locally connected in W (so that each point of W has a neighborhood which intersects W_0 in a connected set). Also suppose that $p \in W$ but p is no longer necessarily in W_0. Suppose $\{U_1, \cdots, U_k\}$ is a pairwise disjoint collection of open subsets of V with $q_i \in U_i$. As above, we may choose a neighborhood A of p so that $\pi^{-1}(A) \subset \bigcup U_i$. By replacing each U_i with $U_i \cap \pi^{-1}(A)$, we may assume that

$$\pi^{-1}(A) = \bigcup U_i$$

We may also choose A so that $A_0 = W_0 \cap A$ is connected. By the above paragraph, $\pi^{-1}(A_0)$ is a finite holomorphic cover of A_0 of pure order. If we set $U_i' = \pi^{-1}(A_0) \cap U_i = U_i \cap V_0$, then U_i' is dense in U_i. It follows that the restriction of π to U_i is a finite branched holomorphic cover of A. The next proposition summarizes the preceding discussion:

5.12 Proposition. *If $\pi : V \to W$ is a finite branched holomorphic cover with W_0 and $V_0 = \pi^{-1}(W_0)$ as in Definition 5.11, then*

(a) *$\pi : V_0 \to W_0$ is a finite holomorphic covering map; that is, each point $p \in W_0$ has a neighborhood A such that $\pi^{-1}(A)$ is a finite disjoint union of open sets on each of which π is a biholomorphic map onto A;*

(b) *if W_0 is locally connected in W, $p \in W$ and $q \in \pi^{-1}(p)$, then there are arbitrarily small neighborhoods U of q and $A = \pi(U)$ of p such that $\pi : U \to A$ is a finite branched holomorphic cover of pure order.*

With $\pi : V \to W$ as in part (b) of the above proposition and $q \in V$, we know that there are arbitrirly small neighborhoods of q on which π is a finite branched holomorphic cover of pure order. For small enough neighborhoods the order must stabilize at some positive integer $o_q(\pi)$. We call this integer the *branching order* of π at q.

5.13 Lemma. *Suppose $p_{m+1}, \dots, p_n$ are monic polynomials with coefficients in ${}_m\mathcal{H}_U$ for some open set $U \subset \mathbb{C}^m$ and with non-zero discriminants $d_{m+1}, \dots, d_n$ and let $V = \{z \in \mathbb{C}^n : p_j(z_j) = 0, \quad j = m+1, \dots n\}$. Then the projection $\mathbb{C}^n \to \mathbb{C}^m$ exhibits V as a finite branched holomorphic cover of U.*

Proof. We let D be the union of the zero sets of the discriminants d_j. Then $U_0 = U - D$ is an open dense subset of U and we set $V_0 = \pi^{-1}(U_0)$ where $\pi : V \to U$ is the restriction of the projection $\mathbb{C}^n \to \mathbb{C}^m$ to V. We need to show that V_0 is dense in V, that π is proper and finite to one on V and locally biholomorphic on V_0.

Now let $K \subset U$ be compact and set $L = \pi^{-1}(K)$. We claim that L is a bounded subset of $\mathbb{C}^n$. Clearly the j^{th} components of points of L are bounded if $j \le m$ since then they are j^{th} components of points of K. For $j > m$ the j^{th} components of points of L are bounded because they satify the monic polynomial equation $p_j(z_j) = 0$. If we divide this equation by $z_j^{n_j - 1}$, where n_j is the order of p_j, we may use the resulting equation to estimate $|z_j|$ in terms of the coefficients of p_j, on the set where $|z_j| \ge 1$. We conclude that $|z_j|$ is bounded by the maximum of 1 and the sum of the suprema of the absolute values of the coefficients of p_j on K. Thus, the set of j^{th} components of points of L is a bounded set for all j and, therefore, L is bounded. The set L is also closed in $\mathbb{C}^n$ since it is just the set of points in $\mathbb{C}^n$ which map to K under $\pi : \mathbb{C}^n \to \mathbb{C}^m$ and at which each p_j vanishes. Hence, L is compact. Thus, π is a proper holomorphic map. The fact that each p_j has a finite, non-empty set of zeroes for each fixed value of $(z_1, \cdots, z_m)$ shows that it is finite to one and surjective.

Let $a = (a', a'')$ be a point of $V \subset U \times \mathbb{C}^{n-m}$ with a' the corresponding point of U. Each polynomial p_j is monic and, hence, $p_j(z_j)$ is regular of some order greater than zero in the variable z_j at a. Given $\epsilon > 0$, Theorem 2.6 implies that we may choose a polydisc $\bar{\Delta}(a, r) = \bar{\Delta}(a', r') \times \bar{\Delta}(a'', r'') \subset U \times \mathbb{C}^{n-m}$, with $r''_j < \epsilon$, such that for each j and for $z' \in \bar{\Delta}(0, r')$ the roots of $p_j(z', z_j)$ all lie in the interior of $\bar{\Delta}(a'', r''_j)$ and the number of these roots is constant, counting multiplicity, as a function of $z' \in \bar{\Delta}(a', r')$. This is a kind of continuity of the roots result. It implies in particular that for each point of V there are arbitrarily nearby points that lie over points of the open dense set U_0. In other words, V_0 is dense in V.

Since U_0 is exactly the set of $z' \in U$ at which the roots of all the p_j are distinct, the inverse image V_0 of this set under π is the subset of $U \times \mathbb{C}^{n-m}$ on which each p_j vanishes but its derivative with respect to z_j does not vanish. Thus, the map $F : \mathbb{C}^n \to \mathbb{C}^{n-m}$ defined by

$$F(z_1, \cdots, z_n) = (p_{m+1}(z_1, \cdots, z_m, z_{m+1}), \cdots, p_n(z_1, \cdots, z_m, z_n))$$

has Jacobian J_F in which the last $n - m$ columns form a diagonal matrix with entries that do not vanish on V_0. Thus, J_F has rank $n - m$ in an open set containing V_0. It follows from the implicit mapping theorem that, for each $\lambda \in V_0$, there is a neigborhood A_λ of λ

in $\mathbb{C}^n$, a neighborhood B_λ of $\pi(\lambda)$ in $\mathbb{C}^m$ and a holomorphic map $G : B_\lambda \to A_\lambda$ such that the points of A_λ where F vanishes (i. e. the points of $V_0 \cap A_\lambda$) are exactly the points in the image of G. It follows that G is a holomorphic inverse for the restriction of π to $V_0 \cap A_\lambda$. Thus, π locally has a holomorphic inverse on V_0. In other words π is locally biholomorphic on V_0 and exhibits V as a branched holomorphic cover of U This completes the proof.

Our further study of finite branched holomorphic covers depends on the following technical lemma.

5.14 Lemma. *Given positive integers n and r, there exists a finite set $\{f_1, \cdots, f_q\}$ of linear functionals on $\mathbb{C}^n$ such that, for any set of r distinct points $\{z_1, \cdots, z_r\} \subset \mathbb{C}^n$ there is some i for which the numbers $f_i(z_1), \cdots, f_i(z_r)$ are distinct.*

Proof. We may assume $r \geq 2$. Choose an integer $q > \frac{1}{2}r(r-1)(n-1)$ and choose a set of linear functionals on $\mathbb{C}^n$, $\{f_1, \cdots, f_q\}$, such that every subset of this set with n elements is linearly independent. If the elements f_i are interpreted as the rows of a $q \times n$ matrix, then this is just the condition that each $n \times n$ submatrix has non vanishing determinant. Such a choice is clearly possible since the union of the zero sets in $\mathbb{C}^{nq}$ of these determinants is a subvariety of dimension less than nq.

Now, given distinct integers j and k between 1 and r, the set of linear functionals f on $\mathbb{C}^n$ for which $f(z_j - z_k) = 0$ is a linear subspace of dimension $n-1$ and, hence, it may contain at most $n-1$ of the functionals $f_1, \cdots, f_q$ since any set of n of these is linearly independent. There are $\frac{1}{2}r(r-1)$ unordered pairs of distinct points in the set $\{z_1, \cdots, z_r\}$ so there are at most $\frac{1}{2}r(r-1)(n-1)$ integers i for which an equation of the form

$$f_i(z_j - z_k) = 0, \quad \text{with} \quad j \neq k$$

can be satisfied. By the choice of q, there must be at least one index i between 1 and q such that no such equation holds. For this i, the functional f_i separates the points $z_1, \cdots, z_r$.

5.15 Lemma. *Let $\pi : V \to W$ be a finite branched holomorphic cover where W is a domain in $\mathbb{C}^m$ and let D be a proper subvariety of W, with the property that $V_0 = V - \pi^{-1}(D)$ is dense in V and π is locally biholomorphic on V_0. If V_1 is a connected component of V_0, then the closure $\bar{V}_1$ of V_1 in V is a holomorphic subvariety of V.*

Proof. The lemma is purely a local statement and so we may assume that W is a polydisc in $\mathbb{C}^m$ and that V is a closed subvariety of some polydisc in $\mathbb{C}^n$.

The set $W_0 = W - D$ is connected (by Problem 5.2) and V_0 has finitely many components, each of which is a finite unbranched holomorphic cover of W_0 and one of which is V_1. Let the points of V_1 over a point $w \in W_0$ be labeled $\lambda_1(w), \ldots, \lambda_k(w)$ and for a function f holomorphic on V let p be the polynomial in the indeterminant x defined by

$$p(w, x) = \prod_{j=1}^{k}(x - f(\lambda_j(w)))$$

For w in a neighborhood of each point of W_0 it is possible to choose the labeling of the $\lambda_j(w)$ in such a way that these functions are holomorphic, although it may not be possible to do this globally. However, since the coefficients of the polynomial p are independent of the labeling of the roots, they are well defined and holomorphic in all of W_0. In fact, if f is holomorphic in all of V, it is locally bounded there, which implies that the coefficients of p are as well. The generalized removable singularities theorem (Theorem 2.7) then implies that they extend to be holomorphic in all of W. Now we have that p is a polynomial with coefficients holomorphic in W and with the property that, whenever $w \in W_0$, the roots of $p(w, x)$ are exactly the values assumed by f on the set $\pi^{-1}(w) \cap V_1$. In particular, $p(\pi(z), z)$ vanishes on V_1. By continuity, this function also vanishes on $\bar{V}_1$.

Since W_0 is connected, π is a cover of pure order r for some r. We apply the previous lemma to obtain linear functionals $f_1, \cdots, f_q$ such that any set of r distinct points in $\mathbb{C}^n$ can be separated by some one of the functionals f_i. We then let $p_1, \cdots, p_q$ be the polynomials constructed, as above, for the functions $f_1, \cdots, f_q$. Each of the functions $p_j(\pi(z), f_j(z))$ vanishes on $\bar{V}_1$ and so we let V^* be the subvariety of V on which they all vanish. We endeavor to prove that $\bar{V}_1 = V^*$. To this end, let a_1 be a point of V^* and let $a_1, a_2, \cdots, a_k$ be the distinct points of $\pi^{-1}(\pi(a_1))$. Note that $k \leq r$ and so there is an i for which the numbers $f_i(a_1), \cdots, f_i(a_k)$ are all distinct. These are all roots of the polynomial $p_i(\pi(a_1), x)$. Then for any $w \in W_0$, sufficiently near $\pi(a_1)$, there must be a root of $p_i(w, x)$ near $f_i(a_1)$. But by the construction of p_i this root must be a value assumed by f_i on $\pi^{-1}(w) \cap V_1$. It follows that for any sequence $\{w_s\} \subset W_0$, converging to $\pi(a_1)$ there is a sequence $\{z_s\} \subset V_1$ with $\pi(z_s) = w_s$ and with $f_i(z_s)$ converging to $f_i(a_1)$. Since π is a proper map, the sequence $\{z_s\}$ must have a limit point $z \in \bar{V}_1$. Necessarily $f_i(z) = f_i(a_1)$ and, hence, $z = a_1$ since f_i separates the points $a_1, \cdots, a_k$. Thus, $a_1 \in \bar{V}_1$ and $\bar{V}_1 = V^*$. This completes the proof.

The branch locus of a finite branched holomorphic cover is the set on which the branching order is at least 2.

5.16 Proposition. *If $\pi : V \to W$ is a finite branched holomorphic cover over a domain W in $\mathbb{C}^m$ and k is any positive integer, then the subset of V on which π has branching order at least k is a closed subvariety of V. In particular, the branch locus of π is a closed subvariety of V. Furthermore, the image under π of the branch locus is a closed subvariety of W.*

Proof. With V_0 and W_0 as above and $V_1 = V_0$ we choose linear functionals $f_1, \cdots, f_q$ and polynomials $p_1, \cdots, p_q$ as in the proof of the previous lemma. Then for any $w \in W$ the roots of the polynomial p_j are exactly the values that f_j assumes on $\pi^{-1}(w)$. A root is a multiple root of p_j of multiplicity k if and only if the polymomials $p_j^{(s)}$ vanish there for $s \leq k$. It follows that $\lambda \in \pi^{-1}(w)$ is a point of branching order at least k if and only if the functions $p_j^{(s)}(\pi(z), f_j(z))$ vanish at λ for $s \leq k$ and for all j. This proves the first part of the Proposition. The second part follow from the fact that the set where the branching order is at least two is the set where all the p_j have multiple roots and this is the set where the discriminants of all the polynomials p_j vanish. The discriminants are functions of $w \in W$ only and the result follows.

5.17 Theorem. *Let $P \subset {}_n\mathcal{H}$ be a prime ideal which is strictly regular in the variables $z_{m+1}, \ldots, z_n$. Then the projection $\mathbb{C}^n = \mathbb{C}^m \times \mathbb{C}^{n-m} \to \mathbb{C}^m$ exhibits* $\operatorname{loc} P$ *as the germ of* *a finite branched holomorphic cover of pure order r over $\mathbb{C}^m$, where r is the order of the defining polynomial for z_{m+1}.*

Proof. We let the polynomials $p_{m+1}, \dots, p_n, q_{m+2}, \dots, q_n$ be as in Lemma 5.9. Choose a polydisc Δ centered at 0 on which each of these germs has a representative and on which each element of a generating set for P (containing the $p's$ and $q's$) also has a representative. We let $\Delta = \Delta' \times \Delta''$ be the usual decomposition induced by representing $\mathbb{C}^n$ as $\mathbb{C}^m \times \mathbb{C}^{n-m}$. We replace each p_j and each q_k by its representative on Δ. We let V be the locus of common zeroes of the set of representatives on Δ of our generating set for P. Then the germ of V at 0 is $\operatorname{loc} P$. We let W be the zero set in $\Delta \cap \mathbb{C}^{m+1}$ of $p_{m+1}(z_{m+1})$. We also let D be the zero set in Δ of the discriminant d of p_{m+1}. Since d depends only on the first m variables, we have $D = D' \times \Delta''$ where D' is the zero set in Δ' of d. By Lemma 5.13 the projection $\mathbb{C}^{m+1} = \mathbb{C}^m \times \mathbb{C} \to \mathbb{C}^m$ induces a holomorphic branched covering of pure order r, $\pi_2 : W \to \Delta'$. Also, if $W_0 = W - (D \cap W)$ then W_0 is the set on which this covering is regular (locally biholomorphic).

Another finite branched holomorphic covering is obtained by applying Lemma 5.13 to the set

$$V_1 = \{z \in \Delta : p_{m+1}(z_{m+1}) = 0, \dots, p_n(z_n) = 0\}$$

and the map $\pi_1 : V_1 \to \Delta'$ induced by the canonical projection $\mathbb{C}^n = \mathbb{C}^m \times \mathbb{C}^{n-m} \to \mathbb{C}^m$. This is regular on the set $V_1 - (E \cap V_1)$, where $E = E' \times \Delta''$ is the union of the zero sets of the discriminants of $p_{m+1}, \cdots, p_n$. We obtain the covering map π that we are looking for by restricting π_1 to the subvariety $V \subset V_1$. It remains to show that π is a finite branched holomorphic cover of pure order r.

It is clear that π is a finite to one, proper, holomorphic map. The density of the regular points and pure order r are not yet clear. We will exploit the fact that π factors as $\pi = \pi_2 \circ \pi_3$ where π_2 is the finite branched holomorphic cover introduced above and $\pi_3 : V \to W$ is induced by the canonical projection $\mathbb{C}^n \to \mathbb{C}^{m+1}$. Note that π_3 is also the germ of a proper finite to one holomorphic map. We let $V_0 = V - (E \cap V)$. We will show that V_0 is a dense open subset of V on which π is locally biholomorphic. First note that, by corollary 5.10, since $D \subset E$, the germ of V_0 at 0 is $\operatorname{loc} I' - (E \cap \operatorname{loc} I')$. Thus, we may choose Δ small enough that $(z', z_{m+1}, \dots, z_n) \in \Delta$ is in V_0 exactly when $z' \notin E'$ and

$$p_{m+1}(z', z_{m+1}) = 0 \quad \text{and}$$

$$q_j(z', z_{m+1}, z_j) = d(z')z_j - s_j(z', z_{m+1}) = 0 \quad \text{for} \quad j = m+2, \dots, n.$$

Since $d(z') \neq 0$ when $z' \notin E'$, the map $\pi_3 : V_0 \to W_0$ is a biholomorphic mapping with inverse given by the equations $z_j = d(z')^{-1} s_j(z', z_{m+1}),\ j = m+2, \dots, n$. Thus, $\pi : V_0 \to \Delta' - E'$ is a locally biholomorphic covering map of degree r and the proof will be complete if we can show that the closure $\bar{V}_0$ of V_0 in V_1 is V.

Now $\pi_1 : V_1 \to \Delta'$ is a finite branched holomorphic cover which is regular on the set $V_1 - (E \cap V_1)$. Thus, π_1 is a locally biholomorphic covering map from $V_1 - (E \cap V_1)$ to $\Delta' - E'$. Since it is also a locally biholomorphic covering map from V_0 to $\Delta' - E'$, it follows that V_0 must be a component of $V_1 - E \cap V_1$ Now by Lemma 5.15 the set $\bar{V}_0$ is a subvariety of V_1. Therefore, $\bar{V}_0$ is also a subvariety of V. Clearly $V = \bar{V}_0 \bigcup (V \cap D)$ and since V is irreducible and not contained in D we must have $\bar{V}_0 = V$. This completes the proof.

A thin subset W of a subvariety V is a set which is contained in a closed subvariety of V with dense complement.

5.18 Corollary. *For any subvariety V of a domain in $\mathbb{C}^n$, the set of singular points is a thin set.*

Proof. This is a local result. At any point p decompose the germ of the variety V into irreducible germs of subvarieties: $V = \bigcup V_i$. Then, in a sufficiently small neighborhood of p, the singular set of V will be the union of the intersection sets $V_i \cap V_j$ for $i \neq j$ and the singular sets of the $V_i's$ themselves. If the neighborhood of p is chosen small enough, then Theorem 5.17 implies that each V_i has a representative that is a finite branched holomorphic cover of a polydisc in $\mathbb{C}^m$ for some m. If V_{i0} is the regular part of V_i for the covering map, then V_{i0} is dense in V_i and the singular set of V_i is contained in the complement of V_{i0} in V_i which is a closed subvariety. Furthermore, Each set $V_i \cap V_j$ for $i \neq j$ is a proper closed subvariety of V_i which necessarily meets V_{i0} in a proper closed subvariety of V_{i0}. However, V_{i0} is locally biholomorphic to a polydisc in $\mathbb{C}^m$ and so the complement of a closed proper subvariety is dense. Since V_{i0} is a dense open subset of V_i, it follows that the complement of $V_i \cap V_j$ in V_i is also dense in V_i.

The above paragraph shows that the set of singular points of V which lie in V_i is a thin subset of V_i. It also follos from the above that a thin subset of V_i is also a thin subset of V. Since the union of finitely many thin sets is clearly thin, we conclude that the set of all singular points of V is a thin set.

5.19 Theorem (Nullstellensatz). *If I is an ideal of ${}_n\mathcal{H}$ then $\operatorname{id}\operatorname{loc} I = \sqrt{I}$.*

Proof. To complete the proof we need only prove that for a prime ideal P we have $\operatorname{id}\operatorname{loc} P \subset P$. Suppose $f \in \operatorname{id}\operatorname{loc} P$ and suppose coordinates have been chosen so that P is strictly regular in $z_{m+1}, \dots, z_n$. By Lemma 5.9 there is a polynomial r with coeficients in ${}_m\mathcal{H}$ of degree less than the degree of the minimal polynomial p_{m+1} of z_{m+1} such that $d^\nu f(z) - r(z', z_{m+1}) \in P$. But then r must vanish on $V = \operatorname{loc} P$ since f does. However, by Theorem 5.17, the projection on the first m coordinates exhibits V as the germ of a finite branched holomorphic cover of $\mathbb{C}^m$ of pure degree equal to the degree of p_{m+1}. If we choose a sufficiently small polydisc Δ, on which f, r, and p_{m+1} may all be replaced by representatives, then this implies that r has as many distinct roots as the degree of p_{m+1} on an open dense subset of Δ and, unless $r = 0$, this contradicts the fact that the degree of r is less than the degree of p_{m+1}. Hence, $d^\nu f \in P$, which implies $f \in P$ since P is prime and $d \notin P$. This completes the proof.

The following three theorems represent very useful refinements of the information in Theorem 5.17 and their proofs include some nice applications of the Nullstellensatz.

5.20 Theorem. *If P is a prime ideal that is regular in the variables $z_{m+1}, \dots, z_n$ and $\pi : \operatorname{loc} P \to \mathbb{C}^k$ is the germ of the holomorphic mapping induced by the projection $\mathbb{C}^n \to \mathbb{C}^k$ where $m \leq k \leq n$ then the image of π is a germ of a holomorphic subvariety of $\mathbb{C}^k$ and π is the germ of a finite branched holomorhic cover of its image.*

Proof. Let $V = \operatorname{loc} P$. The projections $\mathbb{C}^n \to \mathbb{C}^k$ and $\mathbb{C}^n \to \mathbb{C}^m$ induce germs $\pi : \operatorname{loc} P \to \mathbb{C}^k$ and $\pi' : \operatorname{loc} P \to \mathbb{C}^m$ of holomorphic mappings. By Problem 3, π' is the germ of a finite branched holomorphic cover of $\mathbb{C}^m$. With the polynomials p_j as in the proof of Theorem 5.17, consider the germ

$$W = \{z \in \mathbb{C}^k : p_{m+1}(z_{m+1}) = \cdots = p_k(z_k) = 0\}$$

By Theorem 5.13 the projection $\mathbb{C}^k \to \mathbb{C}^m$ induces the germ $\pi'' : W \to \mathbb{C}^m$ of a finite branched holomorphic cover. Since $p_j(z_j) \in P$ it is clear that $\pi(V) \subset W$ so that $\pi' = \pi'' \circ \pi$. Now choose representatives so that $\pi' : V \to U$ and $\pi'' : W \to U$ are finite branched holomorphic covers of a connected open set $U \subset \mathbb{C}^m$. For a suitable proper subvariety $D \subset U$ the sets $V_0 = \pi'^{-1}(U - D)$ and $W_0 = \pi''^{-1}(U - D)$ are sets on which π' and π'' are locally biholomorphic covering maps. Since $\pi' = \pi'' \circ \pi$, it is clear that π is a finite, proper holomorphic mapping and $\pi|_{V_0}$ is a locally biholomorphic covering map. The image $\pi(V_0)$ is the union of some connected components of W_0 and, hence, its closure in W is a subvariety by Lemma 5.15. Since V_0 is dense in V and π is a proper continuous mapping, it follows that the closure of $\pi(V_0)$ in W is $\pi(V)$. Thus, $\pi(V)$ is a subvariety and π is a finite branched holomorphic cover. This completes the proof.

5.21 Theorem. *Let I be an ideal of ${}_n\mathcal{H}$ and set $V = \operatorname{loc} I$. Then the following three conditions on I are equivalent:*

(i) *For each $j = m+1, \dots, n$ there is an element $f_j \in {}_j\mathcal{H} \cap I$ which is regular in z_j;*
(ii) *For each irreducible component V_i of V the projection $\mathbb{C}^n \to \mathbb{C}^m$ induces the germ $\pi_i : V_i \to \pi_i(V_i)$ of a finite branched holomorphic cover and $\pi_i(V_i)$ is the germ of a holomorphic subvariety of $\mathbb{C}^n$;*
(iii) *If L is the germ of $\{z \in \mathbb{C}^n : z_1 = \cdots = z_m = 0\}$, then $L \cap V = \{0\}$.*

Proof. We first show that (i) implies (ii). Suppose that I satisfies (i). Then we may enlarge the sequence $\{f_j\}$ to the extent possible by sequentially choosing non-zero functions $f_j \in {}_j\mathcal{H} \cap I$ for $j = m, m-1, \dots$ as long as ${}_j\mathcal{H} \cap I \neq 0$ and by changing variables in the first j coordinates at each stage, if necessary, to make f_j regular in z_j. The result will be a sequence $0 \neq f_j \in {}_j\mathcal{H} \cap I$ for $j = k+1, \dots, n$ and some $k \le m$ with ${}_k\mathcal{H} \cap I = 0$. This means that I is regular in the variables $z_{k+1}, \dots, z_n$. Now $V = \bigcup_{i=1}^r V_i$ where the V_i are the irreducible components of V. If $P_i = \operatorname{id} V_i$, then each P_i is a prime ideal containing f_j for $j = k+1, \dots, n$. As above, after a change of variables in the first k coordinates, we may assume that P_i is regular in some variables $z_{k_i+1}, \dots, z_n$ with $k_i \le k$. Then by the previous theorem the projection $\mathbb{C}^n \to \mathbb{C}^m$ induces the germ $\pi_i : V_i \to \pi_i(V_i)$ of a finite branched holomorphic cover of the germ of a subvariety $\pi_i(V_i)$. Thus, condition (ii) is satisfied and we have proved that (i) implies (ii).

Now suppose condition (ii) is satisfied. Let $V_1, \dots, V_r$ be the irreducible components of V, as before. The projection $\mathbb{C}^n \to \mathbb{C}^m$ induces the germ $\pi_i : V_i \to \pi_i(V_i)$ of a finite branched holomorphic cover of the germ of a subvariety $\pi_i(V_i)$ of a neighborhood of zero. If $L = \{z \in \mathbb{C}^n : z_1 = \cdots = z_m = 0\}$ then $L \cap V_i = \pi^{-1}(0)$ which must be the germ of a finite set containing 0 since π_i is finite to one. But the germ at 0 of a finite set containing 0 is 0. Thus, (ii) implies (iii).

Now suppose I satisfies condition (iii) for some $m < n$. This means that the ideal J generated by I and $z_1, \dots, z_m$ satisfies $\operatorname{loc} J = \{0\}$. Since z_n vanishes at 0 it belongs to $\operatorname{id} \operatorname{loc} J$ which is $\sqrt{J}$ by the Nullstellensatz. This implies that $z_n^\nu \in J$ for some ν. In other words, $z_n^\nu = f_n + z_1 g_1 + \cdots + z_m g_m$ for some $f_n \in I$ and $g_1, \ldots, g_m \in {}_n\mathcal{H}$. Thus we have produced an $f_n \in I$ that is regular in the variable z_n. The rest of the proof amounts to iterating this procedure as long as we can. In order to carry out the next step, we must prove that if $I \subset {}_n\mathcal{H}$ satisfies (iii) for some $m < n$ then the ideal $I \cap {}_{n-1}\mathcal{H}$ satisfies (iii) in $\mathbb{C}^{n-1}$ for the same m. Since we have just showed that I contains an element regular in z_n and we know that (i) implies (ii) we conclude that the projection $\mathbb{C}^n \to \mathbb{C}^{n-1}$ maps V to the germ V_{n-1} of a holomorphic subvariety of $\mathbb{C}^{n-1}$. Condition (iii) for I implies that $L \cap V_{n-1} = 0$. If we set $I_{n-1} = I \cap {}_{n-1}\mathcal{H}$ then clearly $I_{n-1} \subset \operatorname{id} V_{n-1}$. But if $f \in \operatorname{id} V_{n-1}$ then when f is considered as an element of ${}_n\mathcal{H}$, constant in z_n, it vanishes on V and, by the Nullstellensatz, some power of it belongs to I and, hence, to I_{n-1}. It follows that $I_{n-1} \subset \operatorname{id} V_{n-1} \subset \sqrt{I}_{n-1}$ and from this that $\operatorname{loc} I_{n-1} = V_{n-1}$. Thus, I_{n-1} satisfies condition (iii) as desired. Either $m = n-1$ or we may now conclude as above that there is $f_{m-1} \in I_{n-1} = {}_{n-1}\mathcal{H} \cap I$ which is regular in z_{n-1}. Clearly we can repeat this procedure $n - m$ times to achieve condition (i). This completes the proof.

We say that a germ $\pi : V \to W$ of a holomorphic map between two germs of varieties is *finite* if $\pi^{-1}(0) = (0)$.

5.22 Theorem. *A germ $\pi : V \to W$ of a holomorphic mapping between two germs of holomorphic varieties is finite if and only if for each irreducible component V_i of V, the image $\pi(V_i)$ is the germ of a holomorphic subvariety of W and $\pi : V_i \to \pi(V_i)$ is a finite branched holomorphic cover.*

Proof. Suppose π is finite. We may represent V and W by subvarieties of neighborhoods in $\mathbb{C}^n$ and $\mathbb{C}^m$ ($m \le n$) in such a way that the projection $\mathbb{C}^n \to \mathbb{C}^m$ induces π, by Problem 8. Then $\pi^{-1}(0) = L \cap V$ where L is the germ of $\{z \in \mathbb{C}^n : z_1 = \cdots = z_m = 0\}$. Since π is finite we have $L \cap V = 0$. It then follows from the previous theorem that for each irreducible component V_i of V, $\pi : V_i \to \pi(V_i)$ is the germ of a finite branched holomorphic cover of the germ $\pi(V_i)$ of a subvariety. The reverse implication is obvious and so the proof is complete.

5. Problems

1. Assuming the Nullstellensatz for $\mathbb{C}[z_1, \ldots, z_n]$, prove it for ${}_n\mathcal{O}$.
2. Prove that if $U \subset \mathbb{C}^n$ is a connected open set and $D \subset U$ is a closed subvariety, then $U - D$ is also connected.
3. Prove that if $P \subset {}_n\mathcal{H}$ is a prime ideal which is just regular in the variables $z_{m+1}, \ldots, z_n$, then the projection $\mathbb{C}^n = \mathbb{C}^m \times \mathbb{C}^{n-m} \to \mathbb{C}^m$ still exhibits $\operatorname{loc} P$ as a finite branched holomorphic cover of pure order r over $\mathbb{C}^m$ for some r.
4. In a finite branched holomorphic cover $\pi : V \to W$ prove that each point where the branching order is one is a regular point of the variety V - that is, a point which has a neighborhood in V biholomorphic to a polydisc in $\mathbb{C}^m$ for some m.
5. Prove that if V is an irreducible germ of a variety at $0 \in \mathbb{C}^n$ then V has a representative in some neighborhood of 0 for which the set of regular points is connected.
6. Prove the Nullstellensatz for the local ring ${}_V\mathcal{H}$ of the germ of a holomorphic variety.

7. Use Theorem 5.22 to prove that if U is an open subset of $\mathbb{C}^n$ and $\pi : U \to \mathbb{C}^n$ is a one to one holomorphic map then $\pi(U)$ is open.
8. Prove that if $\pi : V \to W$ is the germ of a holomorphic map between two varieties, then we may represent V and W by subvarieties of neighborhoods in $\mathbb{C}^n$ and $\mathbb{C}^m$ $(m \leq n)$ in such a way that the projection $\mathbb{C}^n \to \mathbb{C}^m$ induces π.
9. Use Proposition 5.12(a) to prove that if $\pi : V \to W$ is a finite branched holomorphic cover and V' is a subset of V such that $\pi : V' \to W$ is also a finite branched holomorphic cover, then V'_0 is both open and closed in V_0. Here, $V_0 = pi^{-1}(W_0)$ and $V'_0 = V' \cap \pi^{-1}(W_0)$ with W_0 as in Def. 5.11.

6. Dimension

We continue with our study of the local properties of regular and holomorphic functions and algebraic and holomorphic varieties. In this section we will be concerned with a germ V of an algebraic or holomorphic variety and with the local ring ${}_V\mathcal{O}$ or ${}_V\mathcal{H}$. In each case there are three notions of dimension of the local ring: a topological dimension, a geometric dimension and a tangential dimension. We show that the first two agree and that they agree with the third if and only if the variety is regular at the point in question. We shall use this last result to prove that the singular locus of a variety is a subvariety. Here the *regular locus* of a holomorphic variety is the set at which it is locally biholomorphic to a polydisc in complex Euclidean space and the *singular locus* is the complement of the regular locus. The regular and singular locus of an algebraic variety have not yet been defined and, in fact, the usual definition is that the regular locus is the set where the second and third of the three dimensions, referred to above, are equal.

We initially take as our definition of the dimension of a holomorphic variety the topological definition:

6.1 Definition. *If V is a holomorphic variety then the dimension of V is the dimension of the complex manifold that is the regular locus of V.*

Note that if $\lambda \in V$ then the dimensions of smaller and smaller neighborhoods of λ in V eventually stabilize and so it makes sense to talk about the dimension of the germ of V at λ. Also note that if V is not irreducible then its regular locus may decompose into germs of manifolds of different dimensions. In this case, by dimension we mean the maximal dimension that occurs. Thus, if $V = V_1 \cup V_2 \cup \cdots \cup V_k$ is the decomposition of a germ of a variety into its irreducible components, then the dimension of V is the maximum of the dimensions of the V_j.

6.2 Lemma. *If V is the germ of a variety in $\mathbb{C}^n$ and $I = \operatorname{id} V$ is regular in the variables $z_{m+1}, \dots, z_n$ then $\dim V = m$.*

Proof. If V is irreducible then this is an obvious corollary of Theorem 5.17. If V is not irreducible, let $V = V_1 \cup \dots V_r$ be a decomposition of V into irreducibles. By Theorem 5.21, the projection $\pi : \mathbb{C}^n \to \mathbb{C}^m$ determines a finite branched holomorphic cover of each V_i onto a germ of a subvariety of $\mathbb{C}^m$. The dimension of V_i is the same as the dimension of $\pi(V_i)$ and this is less than or equal to m for each i. To complete the proof, we must show that one of these dimensions is equal to m. If not, then, for each i, $\pi(V_i)$ is a germ of a proper subvariety of $\mathbb{C}^m$ and, hence, there is a non-zero element $g_i \in {}_m\mathcal{H}$ which vanishes on $\pi(V_i)$. Then $g = g_1 g_2 \dots g_r \in \operatorname{id} V = \sqrt{I}$ and so some power of g is a non-zero element of $I \cap {}_m\mathcal{H}$ which violates the assumption that the ideal I is regular in $z_{m+1}, \dots, z_n$. This completes the proof.

6.3 Theorem. *If V and W are germs of holomorphic varieties with $V \subset W$ then $\dim V \leq \dim W$ and the two are equal exactly when V and W have a common irreducible component of dimension $\dim W$.*

Proof. Since $V \subset W$, each irreducible component of V is contained in some irreducible component of W. Thus, the theorem can be reduced to the case where V and W are irreducible. In this case we can assume that W is a germ of a subvariety of $\mathbb{C}^n$ and $\operatorname{id} W$ is regular in variables $z_{m+1}, \dots, z_n$ where $m = \dim W$. This means that there are functions $f_j \in {}_j\mathcal{H} \cap \operatorname{id} W$ regular in z_j for $j = m+1, \dots, n$. Note that these functions also belong to $\operatorname{id} V$ since $V \subset W$. Now if ${}_m\mathcal{H} \cap \operatorname{id} V = 0$ then $\operatorname{id} V$ is also regular in the variables $z_{m+1}, \dots, z_n$. If not, then we can choose a non-zero $f_m \in {}_m\mathcal{H} \cap \operatorname{id} V$ and, after a change of coordinates involving only $z_1, \dots, z_m$, assume that it is regular in z_m. Continuing in this way, we may choose coordinates so that $\operatorname{id} V$ is regular in the coordinates $z_{k+1}, \dots, z_n$ for some $k \leq m$ which by the previous lemma must be $\dim V$. This makes it obvious that $\dim V \leq \dim W$ and that if $\dim V = \dim W$ then both V and W have germs that are exhibited as finite holomorphic branched covers of the same polydisc by the projection $\mathbb{C}^n \to \mathbb{C}^m$. If the points that lie above a proper subvariety D of this polydisc are removed then the remaining regular parts V_0 and W_0 of V and W are unbranched holomorphic covers of the same set in $\mathbb{C}^m$. Both are connected since V and W are irreducible. It follows that V_0 and W_0 must coincide (Problem 5.9), from which it follows that their closures V and W coincide.

6.4 Theorem. *If V is the germ of a holomorphic subvariety at the origin in $\mathbb{C}^n$ then $\dim V$ is the smallest integer k so that there is a linear subspace L of dimension $n-k$ such that $L \cap V = 0$.*

Proof. If there is a $n-k$-dimensional linear subspace L with $V \cap L = 0$ then choose coordinates for which $L = \{z \in \mathbb{C}^n : z_1 = \cdots = z_k = 0\}$. By Theorem 5.21 there are elements $f_j \in {}_j\mathcal{H} \cap \operatorname{id} V$ for $j = k+1, \dots n$ such that f_j is regular in z_j. But then after some linear change of variables in the first k coordinates we may assume that $\operatorname{id} V$ is regular in $z_{m+1}, \dots, z_n$ for some $m \leq k$. By lemma 6.2 it follows that m is necessarily $\dim V$. Hence, $\dim V \leq k$.

On the other hand, if $k = \dim V$ let $V = \cup_i V_i$ be the decomposition of V into irreducibles and. After a change of variables, it may be assumed that for each i, $\operatorname{id} V_i$ is regular in the variables $z_{n_i+1}, \dots, z_n$ for some integer n_i. Necessarily, $n_i = \dim V_i$ by lemma 6.2 again. It follows from Theorem 5.21 that $L_i \cap V_i = 0$ where $L_i = \{z \in \mathbb{C}^n : z_1 = \cdots = z_{n_i} = 0\}$. Since $k = \max_i n_i$ it follows that $L \cap V_i = 0$ for each i where $L = \{z \in \mathbb{C}^n : z_1 = \cdots = z_k = 0\}$. Therefore L is an $n-k$ dimensional subspace with $L \cap V = 0$ and this completes the proof.

A germ of a variety has pure dimension r if each of its irreducible components has dimension r.

6.5 Theorem. *A germ of a variety V in $\mathbb{C}^n$ has pure dimension $n-1$ if and only if $\operatorname{id} V$ is a principal ideal.*

Proof. It is easy to see that this is true in general if it is true for irreducible varieties (Problem 6.1). Thus, let V be an irreducible variety such that the prime ideal $P = \operatorname{id} V$ is generated by a single element f. Then f must be irreducible. After a change of variables, we may also assume that f is regular in z_n and, after multiplying by a unit, we may assume it is a Weierstrass polynomial in ${}_{n-1}\mathcal{H}[z_n]$. We also have that $P \cap {}_{n-1}\mathcal{H} = 0$ since otherwise there would be an element $g = fh \in {}_{n-1}\mathcal{H}$. But for each $z' \in \mathbb{C}^{n-1}$ near the origin the polynomial f has at least one root and this implies that $g(z') = 0$ for all z' near the origin, i. e. that g is the zero germ. Now we have that P is regular in the variable z_n and, by Lemma 6.2, $V = \operatorname{loc} P$ has dimension $n-1$.

Conversely, suppose that V is an irreducible germ of a variety of dimension $n-1$. Then we may choose coordinates so that $P = \mathrm{id}\, V$ is strictly regular in $z_{m+1}, \ldots, z_n$ for some $m < n$. However, by Lemma 6.2, m is the dimension of V and, thus, $m = n-1$. In other words, P is regular in z_n. Let p be the minimal polynomial for z_n mod P. Recall that by 5.17, the projection $\mathbb{C}^n \to \mathbb{C}^{n-1}$ induces a finite branched holomorphic cover $\pi : V \to W$ where W is some a connected neighborhood of 0 in $\mathbb{C}^{n-1}$. Furthermore, for z' in a dense open subset W_0 of W the points $\pi^{-1}(z')$ are exactly the roots of $p(z', z_n)$. Now if $f \in P$ then $f = pg + r$ for some $g \in {}_n\mathcal{H}$ and some polynomial $r \in {}_{n-1}\mathcal{H}[z_n]$ of degree less than the degree of p. However, since f and p belong to P so does r and, therefore for each $z' \in W_0$ it vanishes at the points of $\pi^{-1}(z')$, but there are two many of these and so it vanishes on W_0 and, hence, on W. Thus we have proved that p generates P. This completes the proof.

6.6 Lemma. *Suppose X is a compact topological space, D is a domain in $\mathbb{C}^n$ containing 0 and $f_1, \ldots, f_k$ are continuous functions on $D \times X$ which are holomorphic in $z \in D$ for each fixed $x \in X$. Let $V(x)$ be the subvariety of D defined by*

$$V(x) = \{z \in D : f_1(z, x) = \cdots = f_k(z, x)\} = 0$$

Then the dimension of the germ of $V(x)$ at 0 is an upper semicontinuous function of x.

Proof. Let $\dim V(x)$ denote the dimension of the germ of $V(x)$ at 0. Since dimension is integral valued, to show that $\dim V(x)$ is upper semicontinuous we must show that for each $x_0 \in X$ there is a neighborhood of x_0 in which $\dim V(x) \leq \dim V(x_0)$. Let x_0 be such a point and let $m = \dim V(x_0)$. Then by Theorem 6.4 there is a linear subspace of $\mathbb{C}^n$ of dimension $n - m$ such that $L \cap V(x_0) = 0$. Choose a neighborhood W of 0 which has compact closure in D. The boundary of W does not meet $L \cap V(x_0)$. In other words, the function $\sum |f_i(z, x)|$ does not vanish on the compact set $(L \cap \partial W) \times \{x_0\}$. It follows that there is a neighborhood U of x_0 such that this function does not vanish on $(L \cap \partial W) \times U$. Then for $x \in U$ we have that $V(x) \cap L$ does not meet ∂W and so $V(x) \cap L \cap W$ is either empty or is a compact subvariety K of W. The maximum modulus theorem, applied to the restriction to K of each coordinate function in $\mathbb{C}^n$, implies that K is a finite set (Problem 6.2). Thus, the germ of $V(x) \cap L$ at 0 is either empty or 0. In the first case, the germ of $V(x)$ is empty (dimension -1) and in the second case $\dim V(x) \leq m = \dim V(x_0)$ by Theorem 6.4 again. This completes the proof.

6.7 Theorem. *Let V be the germ of a holomorphic variety of pure dimension m and let $f \in {}_V\mathcal{H}$ be a non unit and a non zero divisor. Then the locus of the ideal generated by f in ${}_V\mathcal{H}$ is a subvariety of V of pure dimension $m - 1$.*

Proof. Choose representatives in a neighborhood U of 0 for V, f and holomorphic functions g_i, $i = 1, \ldots, k$, such that

$$V = \{z \in U : g_1(z) = \cdots = g_k(z) = 0\}$$

Set $W = \{z \in V : f(z) = 0\}$. Since f is not a unit, W is not empty. Since f is not a zero divisor, it does not vanish identically on any irreducible component of V and so V and

W have no irreducible components in common. Let Δ be a polydisc centered at 0 small enough that $\Delta + \Delta \subset U$ and consider the family of subvarieties $V(\lambda)$ of Δ given by

$$V(\lambda) = (V - \lambda) \cap \Delta = \{z \in \Delta : g_1(z + \lambda) = \cdots = g_k(z + \lambda) = 0\}$$

for $\lambda \in \Delta$ and the family $W(\lambda)$ given by

$$W(\lambda) = (W - \lambda) \cap \Delta = \{z \in \Delta : f(z + \lambda) = g_1(z + \lambda) = \cdots = g_k(z + \lambda) = 0\}$$

for $\lambda \in \Delta$. Since each $V(\lambda)$ is the intersection with Δ of a translate of V, if Δ is chosen small enough, each $V(\lambda)$ will be a variety of pure dimension m and for an open dense set of $\lambda \in \Delta$, 0 will be a regular point of $V(\lambda)$. For these values of λ, some nieghborhood of 0 in $V(\lambda)$ will be biholomorphic to a polydisc. It follows from Theorem 6.5 that, at such points, $W(\lambda)$ has pure dimension $m - 1$. Then Lemma 6.6 implies that $W = W(0)$ has dimension at least $m-1$. However, the only possiblility other than $m-1$ for this dimension is m and, in this case, V and W would have to have a common irreducible component by Theorem 6.3. We have already pointed out that this cannot happen since f is not a zero divisor. This completes the proof.

6.8 Definition. *The Krull dimension of a local ring A is the largest integer d for which there exists a strict chain $P_0 \subset P_1 \subset \cdots \subset P_d$ of prime ideals of A – that is, a chain in which all the containments are proper (A itself is not considered a prime ideal).*

This is the second notion of dimension we referred to earlier - the geometric one.

6.9 Theorem. *If V is the germ of a holomorphic variety and ${}_V\mathcal{H}$ is its local ring, then the dimension of V is equal to the Krull dimension of ${}_V\mathcal{H}$.*

Proof. Note that it follows from the previous theorem that if $W' \subset W''$ are germs of irreducible subvarieties of V and $\dim W'' - \dim W' \geq 2$ then there is another irreducible subvariety W''' which is properly contained in W'' and properly contains W'. Indeed, there must be a germ $f \in {}_V\mathcal{H}$ which belongs to $\operatorname{id} W'$ but not to $\operatorname{id} W''$. The zero locus in W'' of such a function contains W' and is a subvariety of W'' of pure dimension $\dim W'' - 1$. Thus, some irreducible component of this variety properly contains W' and is properly contained in W'' and will do as our W'''.

It is clear from the above paragraph that if $0 = W_d \subset W_{d-1} \subset \cdots \subset W_0$ is a maximal chain of irreducible subvarieties of W, then successive varieties in the chain differ in dimension by exactly one. It follows that $d = \dim W_0$ for such a chain. Clearly d will be largest possible when W_0 is an irreducible subvariety of V of largest dimension. That is, when W_0 is an irreducible component of V with the same dimension as V. This completes the proof.

We now turn to the third notion of dimension – tangential dimension.

6.10 Definition. *If V is a germ at 0 of a algebraic or holomorphic variety then a tangent vector to V is a derivation at 0 – that is, linear map $t : {}_V\mathcal{H} \to \mathbb{C}$ ($t : {}_V\mathcal{O} \to \mathbb{C}$) such that $t(fg) = f(0)t(g) + g(0)t(f)$. The vector space of all tangent vectors is called the tangent space to V and is denoted $T(V)$. Its dimension is the tangential dimension of V and is denoted* $\operatorname{tdim} V$.

6.11 Theorem. *The vector space $T(V)$ is naturally isomorphic to the dual of M/M^2 where M is the maximal ideal of ${}_V\mathcal{H}$ (${}_V\mathcal{O}$).*

Proof. If $t \in T(V)$ then $t(1) = 2t(1)$ and so t kills constants and is, thus, determined by its restriction to M. However, if t is any linear functional on ${}_V\mathcal{H}$ (${}_V\mathcal{O}$) which kills constants and $f = 1 + f_1, g = 1 + g_1$ with $f_1, g_1 \in M$ then

$$t(fg) = t(f_1) + t(g_1) + t(f_1 g_1) = g(0)t(f) + f(0)t(g) + t(f_1 g_1)$$

from which we conclude that t is a tangent vector if and only if t vanishes on M^2. Thus, restriction to M defines an isomorphism between $T(V)$ and the dual of M/M^2. This completes the proof.

It is clear that if $\mathbb{C}^n$ is considered either a holomorphic or an algebraic variety then $T(\mathbb{C}^n)$ is naturally isomorphic to $\mathbb{C}^n$ where a given point $t = (t_1, \dots, t_n) \in \mathbb{C}^n$ is associated to the derivation on ${}_n\mathcal{H}$ or ${}_n\mathcal{O}$ defined by $t(f) = \sum t_i \dfrac{\partial f}{\partial z_i}(0)$.

If $F : V \to W$ is a germ of a holomorphic map between varieties, then F induces an algebra homomorphism $F^* : {}_W\mathcal{H} \to {}_V\mathcal{H}$ ($F^* : {}_W\mathcal{O} \to {}_V\mathcal{O}$) which, in turn, induces a linear map $dF : T(V) \to T(W)$ by $dF(t)(f) = t(F^*(f))$.

6.12 Theorem. *If V is a germ at 0 of a subvariety of $\mathbb{C}^n$ and $F : V \to \mathbb{C}^n$ is the inclusion, then $dF : T(V) \to T(\mathbb{C}^n)$ is injective and its image is $\{t \in T(\mathbb{C}^n) : t(g) = 0 \quad \text{whenever} \quad g \in \operatorname{id} V\}$.*

Proof. We have that

$$dF(t)(g) = t(F^*(g)) = t(g \circ F) = t(g|_V)$$

and so $dF(t) = 0$ if and only if $t = 0$. Furthermore, a derivation on ${}_n\mathcal{H}$ of the form $dF(t)$ clearly vanishes on $\operatorname{id} V = \{g \in {}_n\mathcal{H} : g|_V = 0\}$. Conversely, if s is a derivation on ${}_n\mathcal{H}$ which vanishes on $\operatorname{id} V$, then s determines a well defined linear functional t on ${}_V\mathcal{H} = {}_n\mathcal{H}/\operatorname{id} V$ such that $t(g|_V) = s(g)$. Since s is a derivation, t clearly is as well. This completes the proof for ${}_n\mathcal{H}$. The proof is the same for ${}_n\mathcal{O}$.

In the above theorem, suppose that $\operatorname{id} V$ is generated by $g_1, \dots, g_m$. Then a derivation in $T(C_n)$ vanishes on $\operatorname{id} V$ if and only if it vanishes at each g_i. Thus, if we identify $T(V)$ with its image under dF then

$$T(V) = \{t \in \mathbb{C}^n : \sum_i t_i \frac{\partial g_j}{\partial z_i}(0) = 0 \quad \text{for} \quad j = 1, \dots, m\}$$

From this it is immediate that:

6.13 Corollary. *If V is a germ of holomorphic (algebraic) subvariety of $\mathbb{C}^n$ and $g_1, \dots, g_m$ generate $\operatorname{id} V$ then*

$$\operatorname{tdim} V = n - \operatorname{rank} J_G(0)$$

where $G : \mathbb{C}^n \to \mathbb{C}^m$ is the holomorphic (algebraic) map with the g_j as coordinate functions and J_G is its Jacobian matrix.

6.14 Theorem. *A germ V of a holomorphic variety can be represented as a subvariety of $\mathbb{C}^n$ if and only if $n \geq \operatorname{tdim} V$.*

Proof. Certainly $\operatorname{tdim} V \leq n$ if V is represented as a germ of a subvariety of $\mathbb{C}^n$. On the other hand if V is a holomorphic subvariety of $\mathbb{C}^m$ and $n = \operatorname{tdim} V$ then $\operatorname{rank} J_G(0) = m-n$ where $g_1, \dots, g_k$ is a set of generators for $\operatorname{id} V$. It follows from the implicit function theorem that we may choose holomorphic coordinates for a neighborhood of zero so that the last $m-n$ coordinate functions belong to the set $\{g_1, \dots, g_k\}$. In other words, we may assume that $z_{n+1}, \dots, z_m \in \operatorname{id} V$. This means that V is contained in an n dimensional subspace of $\mathbb{C}^m$ as required.

A germ V of a subvariety of $\mathbb{C}^n$ is said to be *neatly embedded* if $n = \operatorname{tdim} V$. The above theorem says that every germ of a variety can be neatly embedded.

The following is a form of the implicit function theorem that holds for varieties. It will play a key role in chapter 15.

6.15 Theorem. *If $f : V \to W$ is a holomorphic mapping between germs of holomorphic varieties and if $df : T(V) \to T(W)$ is injective, then f is a biholomorphic mapping between V and a germ of a holomorphic subvariety of W.*

Proof. We may assume that V and W are neatly imbedded germs of subvarieties at the origin in $\mathbb{C}^m$ and $\mathbb{C}^n$, respectively. Then f may be regarded as a germ of a holomorphic map from V into $\mathbb{C}^n$. Furthermore, f extends to a germ of a holomorphic mapping $g : \mathbb{C}^m \to \mathbb{C}^n$ since each of its coordinate functions has a holomorphic extension to a neighborhood of 0 in $\mathbb{C}^m$. Since V and W are neatly imbedded, we have that $T(V) = \mathbb{C}^m$ and $T(W) = \mathbb{C}^n$ with the natural identifications given by Theorem 6.12. It follows that the linear maps $df : T(V) \to T(W)$ and $dg : \mathbb{C}^m \to \mathbb{C}^n$ agree after this identification and, hence that dg is injective. The matrix representing dg is the Jacobian matrix $J_g(0)$ and so the usual inverse mapping theorem implies that g is the germ of a biholomorphic map of a neighborhood of 0 in $\mathbb{C}^m$ onto a germ of an m-dimensional submanifold of $\mathbb{C}^n$. Its restriction to V is then a biholomorphic map of V onto a germ of a subvariety of $\mathbb{C}^n$, as required.

The next theorem is useful for identifying cases where the previous theorem applies.

6.16 Theorem. *If $f : V \to \mathbb{C}^n$ is a holomorphic mapping between germs of holomorphic varieties and if the coordinate functions of f generate the maximal ideal of ${}_V\mathcal{H}$, then $df : T(V) \to T(\mathbb{C}^n)$ is injective.*

Proof. This follows from Corollary 6.13 in the following fashion: Assume V is neatly embedded as a germ at the origin of a subvariety of $\mathbb{C}^m$. Extend f to a map g defined in a neighborhood in $\mathbb{C}^m$, as in the previous theorem. Let $\{k_1, \cdots, k_p\}$ be a set of germs in ${}_m\mathcal{H}$ which generates $\operatorname{id} V$ and let $k : \mathbb{C}^m \to \mathbb{C}^p$ be the map with these as coordinate functions. If $h = g \oplus k : \mathbb{C}^m \to \mathbb{C}^{n+p}$, then the set of coordinate functions of h is the union of the set of coordinate functions of g and those of k and so clearly generates the maximal ideal of ${}_m\mathcal{H}$. Since the variety (0) has tangential dimension 0, it follows from Corollary 6.13 that $\operatorname{rank} J_h = m$ and, hence, that dh is injective. However, k has coordinate functions which generate $\operatorname{id} V$. Since V is neatly embedded, $\operatorname{tdim} V = m$ and so it follows, also by Corollary 6.13, that $\operatorname{rank} J_k(0) = 0$. Thus, dk vanishes at 0 and we conclude that dg is injective. Since df and dg agree after the appropriate identification, the proof is complete.

6.17 Theorem. *If V is the germ of a holomorphic variety then* $\dim V \le \operatorname{tdim} V$ *and V is regular if and only if* $\operatorname{tdim} V = \dim V$.

Proof. That $\dim V \le \operatorname{tdim} V$ clearly follows from Theorem 6.14. A germ of a variety is regular if and only if it is biholomorphic to the germ of a neighborhood of zero in $\mathbb{C}^n$ where $n = \dim V$. By Theorem 6.14 again, this is equivalent to $\operatorname{tdim} V = n$.

6.18 Theorem. *If V is a holomorphic subvariety of a domain in $\mathbb{C}^n$ then the singular locus of V is a holomorphic subvariety of V.*

Proof. We don't yet have all of the machinery necessary to prove this. but it fits naturally into this circle of ideas. Thus, we will present a proof that assumes a result that will be proved later (Theorem 12.8).

The result we are after is a local result and so we may assume V is a subvariety of some polydisc Δ and is the union of finitely many irrreducible components V_j. Then the singular locus of V is the union of the singular loci of the V_j and the sets of intersection $V_i \cap V_j$ for $i \neq j$. Thus, it is enough to prove the theorem in the case where V is irreducible. However, for an irreducible subvariety V of a connected set, the dimension of the germ V_z of the variety at $z \in V$ is a constant m. On the other hand, $\operatorname{tdim} V_z$ is $n - \operatorname{rank} J_G(z)$ where J_G is the Jacobian of a holomorphic map G whose coordinate functions form a set of generators of $\operatorname{id} V_z$. By Theorem 12.8, we may choose Δ small enough that there exists a G such that the coordinate functions of G generate $\operatorname{id} V$ at every point $z \in V$. The set where $\operatorname{rank} J_G(z) < k$ is a subvariety for each k and so the set where $\operatorname{tdim} V_z > m = \dim V$ is a subvariety. By the previous theorem, this is the singular locus of V. This completes the proof.

The dimension of the germ V_λ at λ of an algebraic variety V is defined to be the Krull dimension of the corresponding local ring ${}_V\mathcal{O}_\lambda$. A point λ of an algebraic variety V is said to be a *regular point* if $\operatorname{tdim} V_\lambda = \dim V_\lambda$. The singular locus of V is the set of singular points.

Of course, to each algebraic variety there is associated a holomorphic variety $\tilde{V}$ which is the same point set but with a different topology and a different local ring ${}_{\tilde{V}}\mathcal{H}$ associated to a point. it makes sense to ask whether or not the germ af an algebraic variety at a point has the same dimension (or tangential dimension), as the germ of the corresponding holomorphic variety. It also makes sense to ask if the singular locus of an algebraic variety is a proper subvariety and whether or not it agrees with the singular locus of $\tilde{V}$. The rest of this section is devoted to showing that the answer to all these questions is yes. We will need the following two lemmas from commutative algebra which we will not prove (see Atiyah-Macdonald chapter 11):

6.19 Lemma. *Let $B \subset A$ be integral domains with B integrally closed and A integral over B. Then for each prime ideal M of A, M is maximal in A if and only if $N = M \cap B$ is maximal in B and, in this case, the local rings A_M and B_N have the same Krull dimension.*

6.20 Lemma. *The Krull dimension of ${}_n\mathcal{O}$ is n.*

6.21 Theorem. *If V is an irreducible algebraic subvariety of $\mathbb{C}^n$ then* $\dim V_\lambda = \dim \tilde{V}_\lambda$ *at any point $\lambda \in V$, where $\tilde{V}$ is the holomorphic subvariety determined by V.*

Proof. By proceeding as in section 5, we may choose a coordinate system for $\mathbb{C}^n$ such that $\operatorname{id} V$ is regular in the variables $z_{m+1}, \dots, z_n$, that is, $\operatorname{id} V \cap \mathbb{C}[z_1, \dots, z_m] = 0$ and for $j = m+1, \dots, n$ there is a $p_j \in \operatorname{id} V \cap \mathbb{C}[z_1, \dots, z_j]$ which is regular in z_j. This implies that $\mathcal{O}_V = \mathbb{C}[z_1, \dots, z_n]/\operatorname{id} V$ is an integral ring extension of $\mathbb{C}[z_1, \dots, z_m]$. Thus, Lemma 6.19 applies with $A = \mathcal{O}_V$ and $B = \mathbb{C}[z_1, \dots, z_m]$. It follows from Lemma 6.19 that for any point $\lambda \in V$, $\dim{}_V\mathcal{O}_\lambda = \dim{}_m\mathcal{O}_\mu$ where μ is the image of λ under the projection $\mathbb{C}^n \to \mathbb{C}^m$. However, $\dim{}_m\mathcal{O}_\mu = m$ by Lemma 6.20. Since the projection $\mathbb{C}^n \to \mathbb{C}^m$ induces a finite branched holomorphic cover $\tilde{V} \to \mathbb{C}^m$, m is also the dimension of the holomorphic variety $\tilde{V}$.

6.22 Theorem. *If V is an irreducible algebraic subvariety of $\mathbb{C}^n$ and $\tilde{V}$ the corresponding holomorphic subvariety then $\operatorname{tdim} V_\lambda = \operatorname{tdim} \tilde{V}\lambda$ at each $\lambda \in V$. Thus, the singular locus of V is the same point set as the singular locus of $\tilde{V}$. Furthermore, the singular locus of V is a proper algebraic subvariety of V.*

Proof. We will prove in the next section that a generating set for the ideal $\operatorname{id} V_\lambda \subset {}_n\mathcal{O}_\lambda$ is also a generating set for the ideal $\operatorname{id} \tilde{V}_\lambda \subset {}_n\mathcal{H}_\lambda$ (Theorem 7.13). That the tangential dimensions of the germs $\tilde{V}_\lambda$ and V_λ are the same follows from Theorem 6.13.

As in the holomorphic case, that the singular locus of V is an algebraic subvariety follows from a Jacobian argument and the fact that we know from the preceding theorem that $\dim V_\lambda$ is constant in $\lambda \in V$ for an irreducible algebraic variety. The proof is identical to that of Theorem 6.18 except that the role played by Theorem 12.8 in Theorem 6.18 is played by Theorem 10.19 in the algebraic case. The singular locus is a proper subvariety because it is the same point set as the singular locus of $\tilde{V}$ which is proper.

6. Problems

1. Let V be a germ of a subvariety of $\mathbb{C}^n$. Prove that if $\operatorname{id} V_i$ is a principal ideal in ${}_n\mathcal{H}$ for every irreducible component V_i of V, then $\operatorname{id} V$ is also a principal ideal.
2. Let V be a subvariety of a domain in $\mathbb{C}^n$. Prove that if the modulus of a holomorphic function f on V has a local maximum at $z \in V$ then f is constant on the irreducible component of V containing z. Use this to prove that a compact subvariety of an open set in $\mathbb{C}^n$ must be finite.
3. Prove that if V is a germ of a variety then $\operatorname{tdim} V$ is the minimal number of generators for the maximal ideal of ${}_V\mathcal{H}$.
4. Prove that if V is a germ of a holomorphic variety and $P \subset {}_V\mathcal{H}$ is a prime ideal, then $\operatorname{depth}(P) + \operatorname{height}(P) = \dim V$, where $\operatorname{depth}(P)$ is the maximal length of a strict chain of primes with P at the bottom and $\operatorname{height}(P)$ is the maximal length of a strict chain of primes with P at the top.
5. Prove the first part of Lemma 6.19: If $B \subset A$ are integral domains with A integral over B and if $M \subset A$ is a prime ideal and $N = M \cap B$ then M is maximal if and only if N is maximal.

7. Completion of Local Rings

In this section we begin the study of the passage from an algebraic variety V to the corresponding holomorphic variety. Locally, at a point λ of V, this amounts to studying the relationship between ${}_V\mathcal{O}_\lambda$ and ${}_V\mathcal{H}_\lambda$. Clearly ${}_V\mathcal{O}_\lambda$ is a subalgebra of ${}_V\mathcal{H}_\lambda$. Our main goal in this section is to prove that ${}_V\mathcal{H}_\lambda$ is *faithfully flat* over ${}_V\mathcal{O}_\lambda$. We begin with a brief discussion of this notion.

If $A \subset B$ is a pair consisting of a commutative algebra and a subalgebra containing the identity, then there is a functor $X \to X_B$ from the category of A-modules to the category of B-modules defined by

$$X_B = B \otimes_A X$$

The situation is particularly nice when this functor is both exact and faithful ($X \neq 0$ implies $X_B \neq 0$).

7.1 Definition. *Let B be a commutative ring and $A \subset B$ a subring containing the identity. Then B is said to be faithfully flat over A if B is a flat A-module and for each non-zero A-module X the the B-module $X_B = B \otimes_A X$ is non-zero.*

7.2 Lemma. *The following are equivalent:*

(i) *B is faithfully flat over A;*
(ii) *B is flat over A and $x \to 1 \otimes x : X \to X_B$ is injective for every A-module X;*
(iii) *B is flat over A and $IB \cap A = I$ for every ideal I of A;*
(iv) *B/A is a flat A-module.*

Proof. If B is faithfully flat over A and X is an A-module, let K be the kernel of $x \to 1 \otimes x : X \to X_B$. Then the flatness of B implies that $B \otimes_A K \to B \otimes_A X$ is injective. On the other hand, the composition of the maps

$$K \to B \otimes_A K \to B \otimes_A X$$

is the zero map. This implies that $k \to 1 \otimes k : K \to B \otimes_A K$ is the zero map; but the image of this map generates $B \otimes_A K$ over B and, hence, $B \otimes_A K = 0$. It follows from (i) that $K = 0$. This proves that (i) implies (ii). That (ii) implies (i) is obvious.

Note that $I \subset IB \cap A$ for any ideal $I \subset A$. Also, if B is flat over A and $I \subset A$ is an ideal, then applying $B \otimes_A (\cdot)$ to the short exact sequence

$$0 \to I \to A \to A/I \to 0$$

and noting that $B \otimes_A I = BI$, yields the isomorphism $B \otimes_A A/I \simeq B/BI$. Thus, the condition $IB \cap A = I$ is equivalent to the injectivity of the map $A/I \to B \otimes_A A/I$. This proves that (ii) implies (iii). It also proves (iii) implies (ii) since for any A-module X and any $x \in X$ the submodule $Ax \subset X$ has the form A/I for an ideal I.

Now for an A-module X the short exact sequence

$$0 \to A \to B \to B/A \to 0$$

leads to a long exact sequence

$$\cdots \to \mathrm{tor}_1^A(A,X) \to \mathrm{tor}_1^A(B,X) \to \mathrm{tor}_1^A(B/A,X) \to A \otimes_A X \to B \otimes_A X \to \cdots$$

but since $\mathrm{tor}_1^A(A,X) = 0$ and $A \otimes_A X = X$. This becomes

$$0 \to \mathrm{tor}_1^A(B,X) \to \mathrm{tor}_1^A(B/A,X) \to X \to B \otimes_A X \to \cdots$$

This makes it clear that $\mathrm{tor}_1^A(B/A,X) = 0$ for all A-modules X if and only if $\mathrm{tor}_1^A(B,X) = 0$ and $X \to B \otimes_A X$ is injective for all A-modules X. This proves that (ii) is equivalent to (iv) since an A-module is flat if and only if it has vanishing tor_1^A with respect to every A-module. This completes the proof.

7.3 Theorem. *If $A \subset B \subset C$ are algebras with C faithfully flat over A and C faithfully flat over B then B is faithfully flat over A.*

Proof. Suppose $X \to Y$ is an injective morphism of A-modules and let N be the kernel of $B \otimes_A X \to B \otimes_A X$. Then since C is B-flat we have an exact sequence

$$0 \to C \otimes_B N \to C \otimes_B (B \otimes_A X) \to C \otimes_B (B \otimes_A Y)$$

but by the associativity of tensor product we have that $C \otimes_B (B \otimes_A X) = C \otimes_A X$ and $C \otimes_B (B \otimes_A Y) = C \otimes_A Y$. But since C is A-flat we have that $C \otimes_A X \to C \otimes_A Y$ is injective and, hence, $C \otimes_A N = 0$. But this implies that $N = 0$ since C is faithfully flat over A. Thus, we have proved that B is A-flat.

Now suppose that X is an A-module and consider the maps $X \to B \otimes_A X \to C \otimes_A X$. Since the composition is an injection due to the fact that C is faithfully flat over A it follows that the first map is an injection as well and, hence, that B is faithfully flat over A. This completes the proof.

Our strategy for proving that ${}_V\mathcal{H}$ is faithfully flat over ${}_V\mathcal{O}$ will be to inject both of them into a third algebra – the M-adic completion of ${}_V\mathcal{O}$ with respect to its maximal ideal M and to show that this algebra is faithfully flat over both ${}_V\mathcal{O}$ and ${}_V\mathcal{H}$. Then the previous theorem will give us the desired result. To this end, we need to study the completion $\hat{A}$ of a local ring with respect to its maximal ideal.

To begin with we need a lemma about graded Noetherian rings. A graded ring is is a ring $A = \oplus_{n=0}^{\infty} A_n$ which is the direct sum of subspaces A_n in such a way that $A_n \cdot A_m \subset A_{n+m}$ for all n, m. The elements of A_n are said to be *homogeneous of degree n*. A graded Noetherian ring is a graded ring which is also Noetherian as a ring.

7.4 Lemma. *Let A be a graded Noetherian ring. Then*

(i) *A_0 is a Noetherian ring.*
(ii) *A is a finitely generated A_0-algebra.*

Proof. (i) Put $A_+ = \oplus_{n=1}^{\infty} A_n$. Then A_+ is an ideal in A and $A_0 = A/A_+$.

(ii) A_+ is finitely generated. Let $x_1, x_2, \dots, x_s$ be a set of generators of A_+. Without loss of generality we may assume these generators are homogeneous since, otherwise, we may decompose them into homogeneous components. Let $d_i = \deg x_i$, $1 \le i \le s$. Let B be the A_0-subalgebra generated by $x_1, \ldots, x_s$. We claim that $A_n \subseteq B$, $n \in \mathbb{Z}_+$. Clearly, $A_0 \subseteq B$. Assume that $n > 0$, $A_m \subset B$ for $m < n$ and $y \in A_n$. Then $y \in A_+$ and therefore $y = \sum_{i=1}^{s} y_i x_i$ where $y_i \in A_{n-d_i}$. It follows that the induction assumption applies to y_i, $1 \le i \le s$. This implies that $y \in B$ which completes the proof.

We use the above result to prove the key ingredient in the study of m-adic completions of local rings:

7.5 Theorem (Artin, Rees). *Let A be a Noetherian local ring with maximal ideal M and let Y be a finitely generated A-module and X a submodule of Y. Then there exists $m_0 \in \mathbb{Z}_+$ such that*
$$M^{p+m_0}Y \cap X = M^p(M^{m_0}Y \cap X)$$
for all $p \in \mathbb{Z}_+$.

Proof. Put $A^* = \oplus_{n=0}^{\infty} M^n$. Then A^* has a natural structure of a graded ring. Let $(a_1, \ldots, a_s)$ be a set of elements in M with the property the the images of the a_i in M/M^2 generate it. Then for each n, M^n is generated as an A-module by the monomials of degree n in the a_i. Thus, we have a natural surjective morphism $A[x_1, \ldots, x_s] \to A^*$ determined by $x_{i_1} \cdots x_{i_n} \to a_{i_1} \cdots a_{i_n} \in M^n$ which implies that A^* is a graded Noetherian ring. Let $Y^* = \oplus_{n=0}^{\infty} M^n Y$. Then Y^* is a graded A^*-module. It is clearly generated by $Y_0^* = Y$ as an A^*-module. Since Y is a finitely generated A-module, we conclude that Y^* is a finitely generated A^*-module.

In addition, put $X^* = \oplus_{n=0}^{\infty}(X \cap M^n Y) \subset Y^*$. Then
$$M^p(X \cap M^n Y) \subset M^p X \cap M^{n+p} Y \subset X \cap M^{n+p} Y$$
implies that X^* is an A^*-submodule of Y^*. Since A^* is a Noetherian ring, X^* is finitely generated. There exists $m_0 \in \mathbb{Z}_+$ such that $\oplus_{n=0}^{m_0}(X \cap M^n Y)$ generates X^*. Then for any $p \in \mathbb{Z}_+$,
$$X \cap M^{p+m_0} Y = \sum_{s=0}^{m_0} M^{p+m_0-s}(X \cap M^s Y) \subset M^p(X \cap M^{m_0} Y) \subset X \cap M^{p+m_0} Y.$$
Therefore, the inclusions are equalities and the proof is complete.

Now let A be a Noetherian local ring with maximal ideal M. Given any A-module X, we define a topology on X by declaring a neighborhood base for the topology at $x \in X$ to consist of the sets $x + M^n X$ for $n \in \mathbb{Z}_+$. This is a uniform topology for the additive group structure of X and so we may define a completion $\hat{X}$ of X relative to this topology. The completion consists of equivalence classes of Cauchy sequences from X where a sequence $\{x_k\}$ is Cauchy if for each $n \in \mathbb{Z}_+$ there is a $K \in \mathbb{Z}_+$ such that $x_{k_1} - x_{k_2} \in M^n X$ whenever $k_1, k_2 > K$. In other words, a sequence is Cauchy if and only if it is eventually constant mod $M^n X$ for each $n \in \mathbb{Z}_+$. This description makes it clear that $\hat{X}$ may also be described as the inverse limit
$$\hat{X} = \varprojlim X/M^n X$$

When $X = A$ we obtain a completion $\hat{A}$ for A itself. It is easy to see that $\hat{A}$ is also a ring and for each module X over A the completion $\hat{X}$ is a module over $\hat{A}$. In fact, $X \to \hat{X}$ is a covariant functor from A-modules to $\hat{A}$-modules. This is an exact functor when restricted to finitely generated modules as is shown below:

7.6 Theorem. *Let A be a Noetherian local ring and let*

$$0 \to X \to Y \to Z \to 0$$

be an exact sequence of finitely generated A modules. Then the sequence

$$0 \to \hat{X} \to \hat{Y} \to \hat{Z} \to 0$$

is also exact.

Proof. By Artin-Rees we have that there exists an m_0 such that

$$M^{p+m_0}X \subset M^{p+m_0}Y \cap X \subset M^pX$$

This implies that the completion of X with respect to the topology determined by the filtration $\{M^pX\}$ agrees with that determined by the filtration $\{X \cap M^pY\}$. In other words, if $X_p = X/(X \cap M^pY)$ then $\hat{X} = \lim\limits_{\leftarrow} X_p$. But we have for each p a short exact sequence

$$0 \to X_p \to Y_p \to Z_p \to 0$$

where $Y_p = Y/M^pY$ and $Z_p = Z/M^pZ$. Now limits of inverse sequences preserve left exactness but do not always preserve right exactness. Right exactness is, however, preserved in the case where the left hand sequence $\{X_p\}$ is *surjective* in the sense that each map $X_{p+1} \to X_p$ is surjective (Problem 7.1) as is true in our situation. It follows that

$$0 \to \hat{X} \to \hat{Y} \to \hat{Z} \to 0$$

is exact, as required.

7.7 Theorem. *If A is a Noetherian local ring and X is a finitely generated A-module, then $\hat{A} \otimes_A X \to \hat{X}$ is an isomorphism.*

Proof. It is clear that $X \to \hat{X}$ commutes with taking finite direct sums. Thus, since $\hat{A} \otimes_A A \to \hat{A}$ is an isomorophism, we conclude that $\hat{A} \otimes_A F \to \hat{F}$ is an isomorphism whenever F is a finitely generated free module. Since X is finitely generated, we can find a short exact sequence

$$0 \to K \to F \to X \to 0$$

with F free and finitely generated. This yields a diagram

$$\begin{tikzcd}
 & \hat{A} \otimes_A K \arrow[r] \arrow[d] & \hat{A} \otimes_A F \arrow[r] \arrow[d] & \hat{A} \otimes_A X \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \hat{K} \arrow[r] & \hat{F} \arrow[r] & \hat{X} \arrow[r] & 0
\end{tikzcd}$$

in which the bottom row is exact by Theorem 7.6, the top row is right exact and the middle vertical map is an isomorphism. A simple diagram chase shows that $\hat{A} \otimes_A X \to \hat{X}$ is surjective. However, since A is Noetherian, we also have that K is a finitely generated A-module and, by what we just proved, the map $\hat{A} \otimes_A K \to \hat{K}$ is also surjective. It then follows from another diagram chase that $\hat{A} \otimes_A X \to \hat{X}$ is injective. This completes the proof.

7.8 Theorem. *If A is a Noetherian local ring, then $\hat{A}$ is faithfully flat over A.*

Proof. It is easy to see (Problem 7.2) that an A-module Y is flat if and only if whenever X_1 and X_2 are finitely generated A-modules and $X_1 \to X_2$ is injective, then $Y \otimes_A X_1 \to Y \otimes_A X_2$ is also injective. Since Theorems 7.6 and 7.7 prove that $\hat{A} \otimes_A (\cdot)$ preserves exactness of short exact sequences of finitely generated A-modules, we conclude that $\hat{A}$ is flat as an A-module.

Now suppose that X is a finitely generated A-module. Then the kernel of the map $X \to \hat{X}$ is $E = \cap_n M^n X$. It follows from Artin-Rees applied to $E \subset X$ that $ME = E$. Then Nakayama's lemma implies that $E = 0$. Now by Theorem 7.7 we conclude that the map $X \to \hat{A} \otimes_A X$ is injective whenever X is finitely generated. But this clearly implies that this map is injective in general and, hence, that $\hat{A}$ is faithfully flat over A.

7.9 Theorem. *If A is a Noetherian local ring then*

(i) *the unique maximal ideal of $\hat{A}$ is $\hat{M} = \hat{A}M$;*
(ii) *$M^n = A \cap \hat{M}^n$ for all $n \in \mathbb{Z}_+$; and*
(iii) *$A/M^n \to \hat{A}/\hat{M}^{n+1}$ is an isomorphism for all n.*
(iv) *$\hat{A}$ is complete in the $\hat{M}$-adic topology.*

Proof. Since $M^p(A/M) = 0$ for $p \in \mathbb{Z}_+$ we have that A/M is complete in the M-adic topology. We apply the completion functor to the exact sequence

$$0 \to M \to A \to A/M \to 0$$

and use the fact that this functor is exact to conclude that we have an exact sequence

$$0 \to \hat{M} \to \hat{A} \to A/M \to 0$$

This implies that $\hat{M}$ is a maximal ideal of $\hat{A}$ since A/M is a field. We also have that $\hat{M} = \hat{A} \otimes_A M = \hat{A}M$, which implies that $\hat{M}^n = \hat{A}M^n$. Now since $\hat{A}$ is faithfully flat over A, Lemma 7.2 implies that $M^n\hat{A} \cap A = M^n$ and we conclude that $\hat{M}^n \cap A = M^n$. This proves (ii).

That $A/M^n \to \hat{A}/\hat{M}^{n+1}$ is surjective follows from the fact that a Cauchy sequence in the M-adic topology is eventually constant modulo M^n. That this map is injective follows from (ii). This completes the proof of (iii).

Part (iv) follows immediately from (iii) which shows that the $\hat{M}$-adic completion of $\hat{A}$ is $\hat{A}$.

To complete the proof of (i) we must show that $\hat{M}$ is the only maximal ideal of $\hat{A}$. To do this, we need only show that $1 - a$ is a unit in $\hat{A}$ for every $a \in M$. In fact, the inverse of $1 - a$ for $a \in M$ is $1 + a + a^2 + \cdots + a^n + \dots$ which converges in the $\hat{M}$-adic topology of $\hat{A}$.

7.10 Theorem. *If A is a Noetherian local ring then $\hat{A}$ is also a Noetherian local ring.*

Proof. We have that $\hat{A}$ is a local ring from the previous theorem. Thus, we need only show that $\hat{A}$ is Noetherian. The graded ring $G(A) = \sum_{n=0}^{\infty} M^n/M^{n+1}$ associated to A is a finitely generated algebra over the field A/M and is therefore Noetherian by the Hilbert basis theorem. It follows that $G(\hat{A})$ is also Noetherian since it is isomorphic to $G(A)$ by Theorem 7.9(iii). Suppose I is an ideal of $\hat{A}$. If we give I the filtration $\{\hat{M}^n \cap I\}$ then $G(I)$ embedds as an ideal of $G(\hat{A})$ and, as such, it is finitely generated. Let $\{\bar{a}_i\,; i = 1, \ldots, n\}$ be a set of homogeneous generators of $G(I)$, set $r_i = \deg(\bar{a}_i)$ and let $a_i \in I \cap M^{r_i}$ be a representative of $\bar{a}_i$ for each i. Let J be the ideal in $\hat{A}$ generated by $a_1, \ldots, a_n$.

We will prove that $J = I$. Clearly $G(I) = G(J)$. Suppose $u \in I$. Since $\hat{A}$ is Hausdorff, there exists p such that $u \in \hat{M}^p - \hat{M}^{p+1}$. Then there exist $v_{0i} \in \hat{M}^{p-r_i}$ such that $u - \sum v_{0i}a_i \in I \cap \hat{M}^{p+1}$. By continuing this construction we obtain sequences $\{v_{ji}\,; j \in \mathbb{Z}_+, i = 1, \ldots, n\}$ such that $v_{ji} \in \hat{M}^{p+j-r_i}$ and

$$u - \sum_{i=1}^{n} \sum_{j=0}^{s} v_{ji}a_i \in I \cap \hat{M}^{p+s+1}$$

Since $\hat{A}$ is complete, the series $\sum_{j=0}^{\infty} v_{ji}$ converges to some $v_i \in \hat{A}$ for each i and we have $u = \sum_{i=1}^{n} v_i a_i$. Thus, $u \in J$ and the proof is complete.

We now return to the study of the algebras ${}_V\mathcal{O}$ and ${}_V\mathcal{H}$. Note first that if A is ${}_n\mathcal{O}$ or ${}_n\mathcal{H}$ and M is the maximal ideal of A then A/M^p is just the quotient of the ring of polynomials in n variables modulo the ideal consisting of polynomials all of whose terms are of degree at least p. Thus, the following result is obvious from the definitions:

7.11 Lemma. *The algebras ${}_n\mathcal{O}$ and ${}_n\mathcal{H}$ both have as completion the algebra of formal power series $\mathbb{C}[[z_1, \ldots, z_n]]$.*

The following technical lemma due to Chevalley is the key to showing that ${}_V\mathcal{O}$ and ${}_V\mathcal{H}$ also have the same completion.

7.12 Lemma. *Let V be a germ of an algebraic subvariety of $\mathbb{C}^n$. Then there are no non-zero nilpotent elements of ${}_V\hat{\mathcal{O}}$. That is, ${}_V\hat{\mathcal{O}}$ is reduced.*

Proof. We first reduce to the case where V is irreducible. If V is not irreducible and $V = V_1 \cup \cdots V_k$ is its irreducible decomposition, then consider the map

$${}_V\mathcal{O} \to \oplus_i {}_{V_i}\mathcal{O}$$

induced by restricting to each irreducible component. This is an injection and, hence, it remains an injection if we apply the ${}_V\mathcal{O}$ module completion functor to obtain

$${}_V\hat{\mathcal{O}} \to \oplus_i {}_{V_i}\hat{\mathcal{O}}$$

However, the completion of ${}_{V_i}\hat{\mathcal{O}}$ as an ${}_V\mathcal{O}$ module is the same as its completion as a local ring. Therefore, if the theorem is true for each of the irreducible varieties V_i, then a nilpotent element of ${}_V\hat{\mathcal{O}}$ must map to zero in each ${}_{V_i}\hat{\mathcal{O}}$ and, therefore, must be zero. It follows that the theorem is also true for V.

Thus, we will assume that V is irreducible. By the normalization theorem (the algebraic version of Lemma's 5.5 and 5.6 on a prime ideal being regular in variables $z_{m+1}, \cdots, z_n$), there is an integer m and a subalgebra $A \subset {}_V\mathcal{O}$ isomorphic to ${}_m\mathcal{O}$ such that ${}_V\mathcal{O}$ is integral over A. If K is the field of fractions of A and L is the field of fractions of ${}_V\mathcal{O}$ then L is generated by ${}_V\mathcal{O}$ over K. To see this, let $x = a/b \in L$ with $a, b \in {}_V\mathcal{O}$. Then b satisfies a polynomial equation $b^p + c_{p-1}b^{p-1} \cdots + c_1 b + c_0 = 0$ with $c_0 \neq 0$. Thus, $-c_0^{-1}(b^{p-1} + c_{p-1}b^{p-2} + \cdots + c_1) \in K_V\mathcal{O}$ is the inverse for b and so $x \in K_V\mathcal{O}$. It follows that we may choose elements $q_1, \ldots, q_m \in {}_V\mathcal{O}$ which form a basis for L as a vector space over K. Relative to this basis, L may be represented as an algebra of matrices with entries from K. We introduce a K-valued bilinear form $[,]$ on L by

$$[x, y] = \mathrm{tr}(xy)$$

Since each non-zero element of L has an inverse and since the trace of the identity is $m \neq 0$, this form is non-singular. Let d be the determinant of the invertible matrix $B = ([q_i, q_j])$. Since it belongs to ${}_V\mathcal{O}$, the element $q_i q_j$ is integral over A and it follows from Problem 7.4 that its trace belongs to A. Thus, the matrix B has entries in A. It follows that $d \in A$ and B^{-1} has entries in $d^{-1}A$. By applying B^{-1} to the column vector (q_j), we obtain a vector (a_i) with entries in $d^{-1}{}_V\mathcal{O}$ which forms a dual basis to (q_j) in the sense that $[a_i, q_j] = \delta_{ij}$.

Now suppose $x \in L$ is integral over ${}_V\mathcal{O}$. Then each $q_j x$ is also integral over ${}_V\mathcal{O}$ and, hence, integral over A. It follows from Problem 7.4 that $[q_j, x] = \mathrm{tr}(q_j x) \in K$ is actually in A for all j. This implies that when x is expanded in the dual basis (a_i), its coeficients lie in A and, hence, that $x \in d^{-1}{}_V\mathcal{O}$. Thus, we have proved that every element of L which is integral over ${}_V\mathcal{O}$ actually lies in $d^{-1}{}_V\mathcal{O}$.

At this point we pass to the completion $\hat{A}$, which is a just the ring of formal power series by Theorem 7.11 and, hence, is an integral domain. We denote its field of fractions by $\tilde{K}$ and remark that this is an extension field of K. The completion ${}_V\hat{\mathcal{O}}$ of ${}_V\mathcal{O}$ is $\hat{A} \otimes_A {}_V\mathcal{O}$ by Problem 7.3 . We define a $\tilde{K}$ algebra $\tilde{L}$ to be the result of passing from L to $\hat{A} \otimes_A L$ and then localizing relative to the multiplicative set consisting of the non-zero elements of $\hat{A}$. By Theorem 7.8, $\hat{A}$ is faithfully flat over A and so ${}_V\hat{\mathcal{O}} = \hat{A} \otimes_A {}_V\mathcal{O}$ is embedded as a subalgebra of $\hat{A} \otimes_A L$. We claim that this algebra is, in turn, embedded as a subalgebra of $\tilde{L}$. That is, we must show that nothing is killed when we localize. This means we must show that $am = 0$ for $0 \neq a \in \hat{A}$ and $m \in \hat{A} \otimes_A L$ implies that $m = 0$. Since, $\{q_i\}$ forms a basis for L over K, we may write $m = \sum a_i \otimes q_i/c_i$ with $a_i \in \hat{A}, c_i \in A$. Then $am = 0$ implies that $acm = 0$ where $c = \prod c_i \in A$. This implies that

$$\sum a a_i (c/c_i) \otimes q_i = 0$$

and, hence, that $a a_i (c/c_i) = 0$ for each i. Since $\hat{A}$ is an integral domain, this implies that $a_i = 0$ for each i and, hence, that $m = O$. Thus, we have shown that ${}_V\hat{\mathcal{O}}$ is embedded as a subalgebra of $\tilde{L}$.

It is clear from the construction that we may also describe $\tilde{L}$ as $\tilde{K} \otimes_K L$, that is, as the algebra obtained from L by extending its ground field from K to $\tilde{K}$. This is clearly an algebra over $\tilde{K}$ with basis (q_j) as a $\tilde{K}$ vector space. It has a non-singular bilinear form $[,]$, and dual basis (a_i), determined as above. By Problem 7.5, the algebra $\hat{A}$ is integrally closed. Therefore the result of Problem 7.4 applies and we may argue as in the previous paragraph that any element of $\tilde{L}$ that is integral over ${}_V\hat{\mathcal{O}}$ is actually in $d^{-1}{}_V\hat{\mathcal{O}}$. Now let $x \in {}_V\hat{\mathcal{O}}$ be nilpotent. If $a \in A$ then x/a is also nilpotent. Since a nilpotent element is, in particular, integral, it follows as above that $x/a \in d^{-1}{}_V\hat{\mathcal{O}}$. In particular, this holds when $a = y^p$ for any non-zero element y of M and any positive integer p and from this we conclude that $dx \in M^p{}_V\hat{\mathcal{O}} \subset \hat{M}^p$ for each p. Since it follows from Theorems 7.8 and 7.10 that $\cap_{p=1}^{\infty}\hat{M}^p = 0$ we conclude that $x = 0$. This completes the proof.

When confusion might otherwise result we will denote by $\tilde{V}$ the germ of the holomorphic variety associated to V. When it is clear which is meant we will simply write V for either the algebraic or holomorphic variety.

7.13 Theorem. *Let V be a germ of an algebraic variety in $\mathbb{C}^n$ and $\tilde{V}$ be the corresponding germ of a holomorphic variety. Then $\operatorname{id}\tilde{V} = {}_n\mathcal{H}\cdot\operatorname{id}V$.*

Proof. Let $I = \operatorname{id}V \subset {}_n\mathcal{O}$. If $J = {}_n\mathcal{H}I$ then $\operatorname{loc}J = \tilde{V}$ and so it follows from the Nullstellensatz that $\operatorname{id}\tilde{V} = \sqrt{J}$. Thus, if $f \in \operatorname{id}\tilde{V}$ then $f^m \in J$ for some positive integer m. If we consider f to be a formal power series then we have an element $f \in {}_n\hat{\mathcal{O}}$ such that f^m belongs to the ideal ${}_n\hat{\mathcal{O}}I$ in the completion ${}_n\hat{\mathcal{O}} = \mathbb{C}[z_1, \dots, z_n]$ of ${}_n\mathcal{O}$. By the exactness property of completion for finitely generated modules (Theorem 7.6), we have that the quotient of ${}_n\hat{\mathcal{O}}$ modulo this ideal is just ${}_V\hat{\mathcal{O}}$. We conclude that the image of f in ${}_V\hat{\mathcal{O}}$ is nilpotent and, hence, zero by the previous lemma. Thus, $f \in {}_n\hat{\mathcal{O}}I = {}_n\hat{\mathcal{O}}J$. Since $f \in {}_n\mathcal{H}$, it follows from Lemma 7.2(iii) that $f \in J = {}_n\mathcal{H}I$. We conclude that $id\tilde{V} = {}_n\mathcal{H}\cdot\operatorname{id}V$ as claimed.

7.14 Theorem. *If V is a germ of an algebraic variety then ${}_V\hat{\mathcal{H}} = {}_V\hat{\mathcal{O}}$.*

Proof. Let $\{p_1, \dots, p_r\}$ be a set of generators for $\operatorname{id}V \subset {}_n\mathcal{O}$. Then by the previous theorem we have that it is also a set of generators for $\operatorname{id}\tilde{V} \subset {}_n\mathcal{H}$. Thus, we have a commutative diagram

$$\begin{array}{ccccccc}
{}_n\mathcal{O}^r & \longrightarrow & {}_n\mathcal{O} & \longrightarrow & {}_V\mathcal{O} & \longrightarrow & 0 \\
\downarrow & & \downarrow & & \downarrow & & \\
{}_n\mathcal{H}^r & \longrightarrow & {}_n\mathcal{H} & \longrightarrow & {}_V\mathcal{H} & \longrightarrow & 0
\end{array}$$

with exact rows, where $(f_1, \dots, f_r) \to p_1 f_1 + \cdots + p_r f_r$ defines the maps ${}_n\mathcal{O}^r \to {}_n\mathcal{O}$ and ${}_n\mathcal{H}^r \to {}_n\mathcal{H}$. On applying the completion functor to this diagram, the first two vertical maps become isomorphisms and, hence, the third must be one also. This completes the proof.

We now have ${}_V\mathcal{O} \subset {}_V\mathcal{H} \subset {}_V\hat{\mathcal{O}} = {}_V\hat{\mathcal{H}}$. In view of Theorems 7.3 and 7.8 we have proved the main theorem of this chapter:

7.15 Corollary. *If V is a germ of an algebraic variety, then ${}_V\mathcal{H}$ is faithfully flat over ${}_V\mathcal{O}$.*

7. Problems

1. Prove that an inverse limit of an inverse sequence of short exact sequences is exact provided the left hand inverse sequence is surjective.
2. Prove that an A-module Y is flat if and only if whenever X_1 and X_2 are finitely generated A-modules and $X_1 \to X_2$ is injective, then $Y \otimes_A X_1 \to Y \otimes_A X_2$ is also injective.
3. Prove that if A is a Noetherian local ring with maximal ideal M and B is a local ring which is a finitely generated integral ring extension of A and if N is the maximal ideal of B then the completion of B in the N-adic topology is the same as its completion as an A-module with the M-adic topology.
4. Suppose A is an integrally closed integral domain and K is its field of fractions. Prove that if a matrix with entries in K is integral over A then its trace lies in A.
5. Prove that the formal power series ring $\mathbb{C}[[z_1, \ldots, z_n]]$ is a unique factorization domain. Hint: Use induction on the number of variables, Gauss's Theorem (A is a UFD implies A[z] is a UFD), and an extension of the Weierstrass preparation theorem to formal power series.

8. Sheaves

Sheaf theory provides the formal machinery for passing from local to global solutions for a wide variety of problems as well as for classifying the obstruction to so doing when local solutions do not give rise to global solutions. The following is a list of typical examples of such local to global problems:

1. If X is a compact Hausdorff space and f is a continuous complex valued function on X which never vanishes, then f locally has a continuous logarithm. Does it have a logarithm globally? In other words, is there a continuous function g on X such that $f = \exp g$?
2. If U is a domain in $\mathbb{C}$ and g is a $\mathcal{C}^\infty$ function on U then the equation $\dfrac{\partial f}{\partial \bar{z}} = g$ has a solution locally in a neighborhood of each point. Does it have a global solution on U?
3. If U is a domain in $\mathbb{C}^n$, $V \subset U$ is a holomorphic subvariety and f is holomorphic on V, then for each point $\lambda \in V$ there is a holomorphic function defined in a neighborhood U_λ of λ in $\mathbb{C}^n$ whose restriction to $U_\lambda \cap V$ agrees with that of f. Is there a holomorphic function defined on all of U whose restriction to V is f?
4. If U is a domain in $\mathbb{C}^n$ and $V \subset U$ a holomorphic subvariety then V is locally defined as the set of common zeroes of some set of holomorphic functions. Is there a set of holomorphic functions defined on all of U so that V is its set of common zeroes?

Generally these problems involve classes of functions – continuous, holomorphic, $\mathcal{C}^\infty$, etc. – which make sense on any domain in the underlying space. The notion of sheaf simply abstracts this idea:

8.1 Definition. *Let X be a topological space. We consider the collection of open subsets of X to be a category where the morphisms are the inclusions $U \subset V$. Then a presheaf on X is a contravariant functor from this category to the category of abelian groups. A morphism between two sheaves on X is a morphism of functors.*

Thus, a presheaf S on X assigns to each open set $U \subset X$ an abelian group $S(U)$ and to each inclusion of open sets $U \subset V$ a group homomorphism $\rho_{U,V} : S(V) \to S(U)$, called the restriction map, in such a way that $\rho_{U,U} = id$ for any open set U and $\rho_{U,W} = \rho_{U,V} \circ \rho_{V,W}$ for any triple $U \subset V \subset W$.

A morphism $\phi : S \to T$ between two presheaves on X assigns a morphism $\phi_U : S(U) \to T(U)$ to each open set U in a way which commutes with restriction. Unless the context dictates otherwise, we shall usually drop the subscript from ϕ_U and write simply ϕ.

An example of a presheaf is the assignment to each open subset $U \subset X$ of the algebra of continuous functions $\mathcal{C}(U)$. The restriction map $\rho_{U,V}$ is just restriction of functions in this case. The resulting presheaf $\mathcal{C}$ is called the presheaf of continuous functions. The presheave, $\mathcal{C}^\infty$, of $\mathcal{C}^\infty$-funtions on a $\mathcal{C}^\infty$-manifold and the presheaf, $\mathcal{H}$, of holomorphic functions on $\mathbb{C}^n$ are defined in the same way. If $\mathbb{C}^n$ is given the Zariski topology, then we may define on it the presheaf, $\mathcal{O}$, of regular functions. If X is any topological space and G is a fixed abelian group, then we may define a presheaf called the constant presheaf by assigning G to each non-empty open set and 0 to the empty set. The first four of the above examples are actually presheaves of algebras not just of abelian groups. As we shall see, the existence of additional structure on the objects $S(U)$ for a presheaf S is the typical situation, although the abelian group structure is all that is needed in much of the theory.

If S is a presheaf, then an element $s \in S(U)$ will be called a *section* of S over U. If $U \subset V$ then the image of a section $s \in S(V)$ under the restriction map $\rho_{U,V} : S(V) \to S(U)$ will often be denoted $s|_U$ and called the *restriction* of s to U.

A sheaf is a presheaf which is locally defined in a sense made precise in the following definition:

8.2 Definition. *If S is a presheaf on X then S is called a sheaf if the following conditions are satisfied for each open subset $U \subset X$ and each open cover $U = \bigcup_{i\in I} U_i$ of U:*

(i) *if $s \in S(U)$ is a section such that $s|_{U_i} = 0$ for all $i \in I$, then $s = 0$;*
(ii) *if $\{s_i \in S(U_i)\}_{i\in I}$ is a collection of sections such that $s_i|_{U_i\cap U_j} = s_j|_{U_i\cap U_j}$ for all $i, j \in I$, then there is a section $s \in S(U)$ such that $s|_{U_i} = s_i$ for all $i \in I$.*

Note that since the empty cover is an open cover of the empty set, it follows from 8.2(i) that $S(\emptyset) = 0$ if S is a sheaf.

The presheaves of continuous, $\mathcal{C}^\infty$, holomorphic and regular functions described earlier are obviously sheaves. However, the presheaf which assigns a fixed group G to each non-empty open set is not a sheaf unless the underlying space has the property that every open set is connected. There is, however, a closely related sheaf: the sheaf of locally constant functions with values in G. In fact, for every presheaf there is an associated sheaf, as we shall show in Theorem 8.6.

If S is a presheaf on X then the *stalk*, S_x, of S at $x \in X$ is the group $\varinjlim\{S(U) : x \in U\}$. Given a section $s \in S(U)$ and an $x \in U$, the image of s in S_x is denoted s_x and is called the *germ* of s at x. The stalks contain the local information in a presheaf. If $\phi : S \to T$ is a morphism of presheaves then clearly ϕ induces a morphism $\phi_x : S_x \to T_x$ for each x. The fact that sheaves are presheaves that are locally defined is illustrated by the following two results:

8.3 Theorem. *If S is a sheaf on X and $s \in S(U)$ is a section over U, then $s = 0$ if and only if $s_x = 0$ for each $x \in U$.*

Proof. If each s_x vanishes then for each $x \in U$ there is a neighborhood U_x of x such that the restriction of s to U_x is zero. By 8.2(i), this implies that $s = 0$.

8.4 Theorem. *If $\phi : S \to T$ is a morphism of sheaves on X, then ϕ_U is injective for each open set U if and only if ϕ_x is injective for each $x \in X$. Furthermore, ϕ_U is an isomorphism for each open set U if and only if ϕ_x is an isomorphism for each $x \in X$.*

Proof. That ϕ_U injective (surjective) for each open set U implies ϕ_x is injective (surjective) for each $x \in X$ is obvious from the definition of direct limit.

If each ϕ_x is injective then a section $s \in S(U)$ in the kernel of ϕ_U has vanishing germ at each $x \in U$. By the previous theorem this implies that $s = 0$. Thus, ϕ_U is injective.

If ϕ_x is an isomorphism for each $x \in X$ then ϕ_U is injective for each open set U by the previous paragraph and so to complete the proof we must show it is also surjective. Given a section $t \in T(U)$ there is an $s_x \in S_x$ such that $\phi_x(s_x) = t_x$ for each $x \in U$. But this means that for each $x \in U$ there is a neighborhood U_x of x and a section $s_{U_x} \in S(U_x)$

such that $\phi(s_{U_x}) = t|_{U_x}$. Then for $x, y \in U$, $\phi(s_{U_x}|_{U_x \cap U_y} - s_{U_y}|_{U_x \cap U_y}) = 0$. Since $\phi_{U_x \cap U_y}$ is injective, this implies that $s_{U_x}|_{U_x \cap U_y} = s_{U_y}|_{U_x \cap U_y}$ and, by 8.2(ii), that there exists a section $s \in S(U)$ such that $s|_{U_x} = s_{U_x}$ for each $x \in U$. It follows that $\phi(s)_x = t_x$ for each $x \in U$ and, hence that $\phi(s) = t$. Thus, ϕ_U is surjective and the proof is complete.

If S is a presheaf on X then we may construct a topological space $\mathcal{S} = \bigcup_{x \in X} S_x$ by choosing as a neighborhood base of a point $u \in S_x$ all sets of the form $\{s_y : y \in U\}$ where U is a neighborhood of x, $s \in S(U)$ and $s_x = u$. There is a continuous projection $\pi : \mathcal{S} \to X$, defined by $\pi(S_x) = x$. A section for π over a subset $Y \subset X$ is a continuous map $\sigma : X \to \mathcal{S}$ such that $\pi \circ \sigma = id$. The following theorem has an elementary proof which is left as an exercise (Problem 8.1).

8.5 Theorem. *With S, X, $\mathcal{S}$, and π as above,*

(a) *the projection $\pi : \mathcal{S} \to X$ is a local homeomorphism;*
(b) *if Y is a subset of X then a function $\sigma : Y \to \mathcal{S}$, with $\pi \circ \sigma = id$, is a section if and only if for each $y \in Y$ there is a neighborhood U of y in X and an $s \in S(U)$ such that $\sigma(x) = s_x$ for each $x \in U$.*

If S is a presheaf then we may construct another presheaf $\tilde{S}$ as follows: Let $\mathcal{S}$ and $\pi : \mathcal{S} \to X$ be as constructed above. Then for each open set $U \subset X$ we let $\tilde{S}(U)$ be the group of sections of $\pi : \mathcal{S} \to X$ over U. There is a morphism of presheaves $S \to \tilde{S}$ defined by sending $s \in S(U)$ to the section $x \to s_x : U \to \mathcal{S}$.

8.6 Theorem. *If S is a presheaf and $\tilde{S}$ is defined as above, then*

(a) *$\tilde{S}$ is a sheaf;*
(b) *$S \to \tilde{S}$ induces an isomorphism $S_x \to \tilde{S}_x$ for each $x \in X$;*
(c) *$S \to \tilde{S}$ is an isomorphism of presheaves if and only if S is a sheaf.*

Proof. Part (a) is obvious since continuity is a local condition. Part (c) will follow from Theorem 8.4 if we can establish part (b). However, if $s \in S(U)$ has germ at x which is sent to zero by $S_x \to \tilde{S}_x$ then $s_y = 0$ for all y in some neighborhood $V \subset U$ of x. This implies that $s|_V = 0$, by Theorem 8.3 and, hence, that s has germ 0 at x. Thus, $S_x \to \tilde{S}_x$ is injective. That it is surjective follows immediately from Theorem 8.5(b).

We shall call $\tilde{S}$ the *sheaf of germs* of the presheaf S. It is a simple exercise to see that the functor $S \to \tilde{S}$ is the left adjoint functor of the forgetful functor which assigns to a sheaf its associated presheaf (Problem 8.4).

The preceding two theorems make it clear that the definition of sheaf we have adopted is equivalent to the definition given below which is the one most often encountered in the literature:

8.7 Alternate Definition of Sheaf. *A sheaf over X is a topological space $\mathcal{S}$ with a map $\pi : \mathcal{S} \to X$ such that:*

(i) *π is a surjective local homeomorphism;*
(ii) *$\mathcal{S}_x = \pi^{-1}(x)$ has the structure of an abelian group for each $x \in X$;*
(iii) *the group operation is continuous in the topology of $\mathcal{S}$.*

Note that condition (i) above implies that the images of local sections of $\mathcal{S}$ form a neighborhood base for the topology of $\mathcal{S}$ and, in view of this, (iii) is equivalent to the requirement that for each open set $U \subset X$, the sum of two sections of $\mathcal{S}$ over U is again a section. With the above definition, a morphism $\phi : \mathcal{S} \to \mathcal{T}$ of sheaves over X is a continuous map which commutes with projection and is a group homomorphism $\mathcal{S}_x \to \mathcal{T}_x$ between stalks. The equivalence between the category of sheaves in our previous sense and the category of sheaves in this sense is given by the constructions $S \to \mathcal{S}$ and $\mathcal{S} \to \tilde{S}$ of Theorems 8.5 and 8.6. We will normally stick with our original definition, but on occasion it will be useful to use the fact that it has the above alternate formulation.

One of the things the alternate definition of sheaf allows us to do easily is define the notion of a section of a sheaf $\pi : \mathcal{S} \to X$ over a subset $Y \subset X$ which is not necessarily open. A section $\sigma : Y \to \mathcal{S}$ is a continuous map such that $\pi \circ \sigma = id$. The group of all sections over Y will be denoted by $\Gamma(Y, \mathcal{S})$. We shall have more to say about this later (Theorem 8.12). Of course, for an open set $U \subset X$, $\Gamma(U; S) = S(U)$.

An additive category is a category such that for each pair A, B $\hom(A, B)$ has an abelian group structure satisfying a distributive law relative to composition, direct sums are defined and there is a zero object. If $\phi : A \to B$ is a morphism in an additive category, then a kernel $\ker \phi$ for ϕ is an object with a morphism $\alpha : \ker \phi \to A$ such that $\phi \circ \alpha = 0$ and such that any morphism $\beta : C \to A$ with $\phi \circ \beta = 0$ factors through α. Similarly, a cokernel $\operatorname{coker} \phi$ for ϕ is an object with a morphism $\gamma : B \to \operatorname{coker} \phi$ such that $\gamma \circ \phi = 0$ and such that any morphism $\delta : B \to D$ with $\delta \circ \phi = 0$ factors through γ. In general, a morphism need not have a kernel or a cokernel. When they do exist they are unique up to isomorphism.

An abelian category is an additive category such that every morphism $\phi : A \to B$ has both a kernel and a cokernel and the natural map

$$\operatorname{coim} \phi = \operatorname{coker}(\ker \phi \to A) \to \operatorname{im} \phi = \ker(B \to \operatorname{coker} \phi)$$

is an isomorphism.

A morphism $\phi : S \to T$ in the category of presheaves has both a kernel and a cokernel and these are the obvious presheaves: $U \to \ker \phi_U$ and $U \to \operatorname{coker} \phi_U$. It is easy to see that the presheaves on a given space X form an abelian category. However, we are not really interested in this category. We are interested in the category of sheaves.

Suppose $\phi : S \to T$ is a morphism of sheaves. It is easy to see that the kernel of ϕ as a morphism of presheaves (the presheaf $U \to \ker \phi_U$) is, in fact, a sheaf and is the category theoretic kernel of ϕ as a sheaf morphism (Problem 8.2). However, the analogous statement is not true in general for the cokernel (Problem 8.3). In other words, $U \to \operatorname{coker} \phi_U$ need not be a sheaf. However, we have the following:

8.8 Theorem. *Suppose $\phi : S \to T$ is a morphism of sheaves. Then the sheaf of germs of the presheaf $U \to \operatorname{coker} \phi_U$ is a cokernel for ϕ.*

Proof. If C is the presheaf $U \to \operatorname{coker} \phi_U$ and $\tilde{C}$ is its sheaf of germs, then the composition of the presheaf morphisms $T \to C$ and $C \to \tilde{C}$ is a sheaf morphism γ such that $\gamma \circ \phi = 0$. Any sheaf morphism $\delta : T \to D$ such that $\delta \circ \phi = 0$ must factor through $T \to C$ as a presheaf morphism but since it is a sheaf morphism it must actually factor through γ by Problem 8.4. This completes the proof.

We shall denote the sheaf of germs of the presheaf $U \to \operatorname{coker} \phi_U$ by $\operatorname{coker} \phi$. This might seem ambiguous since one might use the same notation for $U \to \operatorname{coker} \phi_U$ itself since it is the cokernel of ϕ as a presheaf. However, we will never do this, since our focus will be on the category of sheaves.

Note that if U is an open set, then the space of sections $(\operatorname{coker} \phi)(U)$ is not the obvious candidate, $\operatorname{coker} \phi_U$, since $U \to \operatorname{coker} \phi_U$ is not a sheaf; however, it is true that we get the obvious thing at the level of stalks: that is, $(\operatorname{coker} \phi)_x = \operatorname{coker} \phi_x$. This follows from the fact that a presheaf and its sheaf of germs have the same stalks (Theorem 8.6(b)). It follows that $\operatorname{im} \phi = \ker(T \to \operatorname{coker} \phi)$ is characterized by

$$(\operatorname{im} \phi)(U) = \{t \in T(U) : t_x \in \operatorname{im} \phi_x \quad \text{for all} \quad x \in U\}$$

In particular, ϕ is an epimorphism ($\operatorname{coker} \phi = 0$) if and only if each ϕ_x is surjective. We shall call ϕ surjective in this case.

Note that we have defined things in such a way that, for a morphism of sheaves, the notions of kernel, image, and cokernel as well as epimorphism and monomorphism are all local, that is, are all defined stalkwise. In particular, a sequence $A \to B \to C$ is exact ($\operatorname{im}(A \to B) = \ker(B \to C)$) if and only if the corresponding sequence of stalks $A_x \to B_x \to C_x$ is exact for each $x \in X$. Using this fact, it is easy to see that the category of sheaves on X is an abelian category (Problem 8.5).

What we have gained in passing from presheaves to sheaves is that sheaves and their morphisms are defined locally. What we have lost in passing from presheaves to sheaves is that the functor which assigns to a sheaf S its group of sections $S(U)$ over an open set U is no longer exact. Studying this loss of exactness is the central theme of sheaf theory.

Recall that $\Gamma(U; \cdot)$ is the functor which assigns to a sheaf its group of sections over a set U. For U open, $\Gamma(U; S) = S(U)$. In case $U = X$, we will sometimes write simply $\Gamma(S)$ for $\Gamma(X; S)$.

8.9 Theorem. *For each open set U the functor $\Gamma(U; \cdot)$ is left exact.*

Proof. Let $0 \to A \xrightarrow{\alpha} B \xrightarrow{\beta} C$ be an exact sequence of sheaves on X. This means that $0 \to A_x \xrightarrow{\alpha} B_x \xrightarrow{\beta} C_x$ is exact for each $x \in X$. Let $U \subset X$ be open. Then $\alpha_U : A(U) \to B(U)$ is injective by Theorem 8.4. Now suppose $b \in B(U)$ is in the kernel of β_U. For each $x \in U$ there is a neighborhood V_x of x and a section $a_x \in A(V_x)$ such that $\alpha(a_x) = b|_{V_x}$. Then $\alpha(a_x)|_{V_x \cap V_y} = \alpha(a_y)|_{V_x \cap V_y}$ for each pair $x, y \in U$ and it follows from the injectivity of $\alpha_{V_x \cap V_y}$ that $a_x|_{V_x \cap V_y} = a_y|_{V_x \cap V_y}$. From the definition of a sheaf it now follows that there exists $a \in A(U)$ such that $a|_{V_x} = a_x$ for each x and from this that $\alpha(a) = b$. This completes the proof.

We shall show that the category of sheaves has enough injectives to construct injective resolutions for each sheaf. From this it follows that we have naturally defined right derived functors for every left exact functor – in particular, for the functors $\Gamma(U; \cdot)$. The resulting functors are those of sheaf cohomology and are the subject of the next section.

We end this section with a discussion of the operations on sheaves that are induced by a continuous map $f : Y \to X$.

8.10 Definition. *Let $f : Y \to X$ be a continuous map of topological spaces. Then,*

(a) *if T is a sheaf on Y then the direct image, f_*T, of T under f is the sheaf on X defined by $U \to T(f^{-1}(U))$;*

(b) *if S is a sheaf on X then the inverse image, $f^{-1}(S)$, of S under f is the sheaf of germs associated to the presheaf on Y defined by $V \to \varinjlim\{S(U) : U \text{open}, f(V) \subset U\}$.*

Part (b) of the above definition is a case where it would be more instructive to use the alternate definition of sheaf given in 8.7. In fact if $\pi : \mathcal{S} \to X$ is a sheaf in this sense, then $f^{-1}\mathcal{S}$ is just the pullback of $\pi : \mathcal{S} \to X$ via $f : Y \to X$. This is the topological space

$$f^{-1}\mathcal{S} = \{(s, y) \in \mathcal{S} \times Y : \pi(s) = f(y)\}$$

with projection $f^{-1}\mathcal{S} \to Y$ given by projection on the second coordinate. This description of $f^{-1}\mathcal{S}$ makes it clear that it is a sheaf on Y which has stalk at $y \in Y$ equal to the stalk of S at $f(y)$. Form this it follows easily that f^{-1} is an exact functor from sheaves on X to sheaves on Y.

The direct image functor f_* from sheaves on Y to sheaves on X is not exact in general. To see this, consider the map $f : Y \to pt$; in this case, f_*T is the sheaf which assigns to pt the group $\Gamma(Y;T)$ and we know that Γ is not always exact (Problem 8.6). On the other hand, an argument like the one in Theorem 8.9 shows that f_* is alway left exact. Thus, it is another functor for which we expect to be able to construct right derived functors.

A special case of the inverse image functor is the restriction functor. Here, $Y \subset X$ is a subset and $i : Y \to X$ is the inclusion. For a sheaf S on X, $i^{-1}S$ is denoted $S|_Y$ and is called the *restriction* of S to Y. Using the alternate definition of sheaf, the restriction of $\pi : \mathcal{S} \to X$ to a subset Y is the space $\mathcal{S}_Y = \pi^{-1}(Y)$ with topology and projection $\mathcal{S}_Y \to Y$ inherited from S and π. The group of global sections of the restricted sheaf, $\Gamma(Y, S|_Y)$, is the same as the group $\Gamma(Y;S)$ of continuous sections of S over Y. It follows from Theorem 8.9 that

8.11 Corollary. *For any subset $Y \subset X$ and any sheaf S on X, $\Gamma(Y;\cdot)$ is a left exact functor.*

Thus $\Gamma(Y;\cdot)$ is also a functor which we expect to have right derived functors if we can show that there are enough injective sheaves.

If S is a sheaf on X, Y is a subspace of X and U is an open set containing Y, then restriction defines a morphism $S(U) = \Gamma(U;S) \to \Gamma(Y;S)$. Thus, restriction defines a morphism $\psi : \varinjlim\{\Gamma(U;S) : Y \subset U\} \to \Gamma(Y;S)$. This is often, but not always an isomorphism:

8.12 Theorem. *The morphism*

$$\psi : \varinjlim\{\Gamma(U;S) : Y \subset U\} \to \Gamma(Y;S)$$

is injective. It is an isomorphism if X is Hausdorff and Y is compact or if X is paracompact and Y is closed.

Proof. If $Y \subset U$ and $s \in \Gamma(U;S)$ is such that s has image zero in $\Gamma(Y;S)$ then $s_x = 0$ for each $x \in Y$. This implies that s has restriction zero in an open set V containing Y

and, hence, that s determines the zero element of $\psi : \varinjlim\{\Gamma(U;S) : Y \subset U\}$. Thus, ψ is injective.

Now suppose $s \in \Gamma(Y;S)$. Then there is an open cover $\{U_i\}_{i\in I}$ of X and a family of sections $\{s_i \in \Gamma(U_i;S)\}_{i\in I}$ such that $s|_{V_i\cap Y} = s_i|_{V_i\cap Y}$ for each i. If X is Hausdorff and Y compact or if X is paracompact and Y closed then we may assume that the cover $\{U_i\}$ is a locally finite cover of X and, furthermore, that there is locally finite collection of open sets $\{V_i\}_{i\in I}$ with $\bar{V}_i \subset U_i$ for each i and $Y \subset \bigcup_{i\in I} V_i$. Now for each $x \in X$ set $I(x) = \{i \in I : x \in \bar{V}_i\}$ and

$$W = \{x \in \bigcup V_i : s_{ix} = s_{jx} \quad \text{for all} \quad i, j \in I(x)\}$$

Then $I(x)$ is a finite set. For a given x let U be a neighborhood of x which meets $\bar{V}_i$ for i in only a finite set, $J \subset I$. Then the set

$$W_x = U - \bigcup_{i\in I-I(x)} \bar{V}_i$$

is an open set containing x. In fact, W_x is the set of $y \in U$ such that $I(y) \subset I(x)$. Now if $x \in W$ choose N_x to be a neighborhood of x contained in $\cap_{i\in I(x)} U_i \cap W_x$ on which $s_i|_{N_x} = s_j|_{N_x}$ for $i, j \in I(x)$. Then the conditions for membership in W will also be satisfied for any point $y \in N_x$. In other words, W is an open subset of X. The fact that the s_i fit together to define a section s on Y means that $Y \subset W$. By construction, the s_i fit together on W to define a section in $\Gamma(W;S)$ which restricts to s on Y.

The functors indroduced in the following definition are also left exact functors on sheaves (Problem 8.7):

8.13 Definition. *If S is a sheaf on X and $s \in \Gamma(X;S)$ is a section, then the support of s is the (necessarily closed) set $K = \{x \in X : s_x \neq 0\}$. If $Y \subset X$ is any subset of X then $\Gamma_Y(X;S)$ is the group of sections $s \in \Gamma(X;S)$ with support contained in Y. If ϕ is a family of subsets of X which is closed under finite unions, then $\Gamma_\phi(X;S)$ is the group of sections $s \in \Gamma(X;S)$ with support contained in some member of ϕ.*

A family ϕ as above is called a family of supports. A common situation in which $\Gamma_\phi(X;S)$ is useful is when X is locally compact and the family of supports is ϕ is the family of compact subsets of X.

8. Problems

1. Prove Theorem 8.5.
2. Prove that if ϕ is a morphism of sheaves then the presheaf $U \to \ker \phi_U$ is a sheaf and is a kernel for ϕ in the category of sheaves.
3. Give an example of a morphism of presheaves ϕ such that the presheaf $U \to \operatorname{coker} \phi_U$ is not a sheaf.
4. Prove that if S is a presheaf then each presheaf morphism $S \to T$, where T is a sheaf, factors uniquely through the morphism $S \to \tilde{S}$ of S to its sheaf of germs $\tilde{S}$. Prove that

this means that $S \to \tilde{S}$ is a left adjoint functor for the forgetful functor which regards a sheaf as just a presheaf.

5. Prove that the category of sheaves is an abelian category.
6. Give an example to show that the functor $\Gamma(X; \cdot)$ is not necessarily exact.
7. Prove that the functors $\Gamma_Y(X, \cdot)$ and $\Gamma_\phi(X, \cdot)$ of Definition 8.11 are left exact.

9. Sheaf Cohomology

Recall that in any abelian category an object A is called *injective* if the functor $\hom(\cdot, A)$ is exact and is called *projective* if the functor $\hom(A, \cdot)$ is exact. In this section we shall be concerned with injective objects. The following summarizes the most elementary properties of injectives (Problem 9.1):

9.1 Theorem. *In any abelian category:*

(1) *An object A is injective if and only if for every monomorphism $i : B \to C$, each morphism $\beta : B \to A$ extends to a morphism $\gamma : C \to A$ such that $\beta = \gamma \circ i$.*
(2) *Every monomorphism $i : A \to B$, with A injective, splits (has a left inverse).*
(3) *If $0 \to A \to B \to C \to 0$ is an exact sequence and A is injective, then B is injective if and only if C is injective.*

Our abelian category is said to have *enough* injectives if for every object A there is an injective object I and a monomorphism $A \to I$. This, and the fact that every morphism has a cokernel, allows one to construct injective resolutions

$$0 \to A \to I^0 \to I^1 \to \cdots \to I^n \to \cdots$$

for each object A. There is another way of thinking about such resolutions which yields both economy of notation and additional insight and so is worth introducing. We will let I denote the complex

$$0 \to I^0 \to I^1 \to \cdots \to I^n \to \cdots$$

and identify A with the complex

$$0 \to A \to 0$$

where A appears in degree zero. Then an injective resolution of A may be thought of as a morphism of complexes

$$i : A \to I$$

where I is a complex of injective objects (zero in negative degrees) and i induces an isomorphism on cohomology (both complexes have vanishing cohomology in all degrees except zero where the cohomology is A). A morphism of complexes which induces an isomorphism of cohomology is called a *quasi-isomorphism*. Thus, an injective resolution of an object A is a quasi-isomorphism $A \to I$ where I is a complex of injectives which vanishes in negative degrees. Actually, insisting that I vanish in negative degrees is equivalent, for the purposes of this theory, to insisting that it be bounded on the left – that is, vanish for sufficiently high negative degrees.

A morphism $\alpha : X \to Y$ of complexes is homotopic to zero if there are morphisms $h^n : X^n \to Y^{n-1}$ such that $h^{n+1} \circ \delta_X^n + \delta_Y^{n-1} \circ h^n = \alpha^n$ for each n. Here, $\delta_X^n : X^n \to X^{n+1}$ and $\delta_Y^{n-1} : Y^{n-1} \to Y^n$ are the differentials in the complexes X and Y. Two morphisms $\alpha, \beta : X \to Y$ of complexes are said to be homotopic if their difference, $\alpha - \beta$ is homotopic to zero. A key result is the following (Problem 9.2):

9.2 Theoem. *If $\alpha : A \to B$ is a morphism in an abelian category and $A \to I$ and $B \to J$ are injective resolutions of A and B, then there is a morphism of complexes $\tilde{\alpha} : I \to J$ such that the diagram*

$$\begin{array}{ccc} A & \longrightarrow & I \\ {\scriptstyle\alpha}\Big\downarrow & & \Big\downarrow{\scriptstyle\tilde{\alpha}} \\ B & \longrightarrow & J \end{array}$$

is commutative. Furthermore, any two maps $I \to J$ with this property are homotopic.

When Theorem 9.2 is applied to two different injective resolutions of A and the identity morphism from A to A, it implies the following:

9.3 Corollary. *If $i : A \to I$ and $j : A \to J$ are two injective resolutions of A then there are morphisms $\alpha : I \to J$ and $\beta : J \to I$ such that*

$$\begin{array}{ccccc} I & \longleftarrow & A & \longrightarrow & I \\ \Big\downarrow{\scriptstyle\alpha} & & \Big\| & & \Big\uparrow{\scriptstyle\beta} \\ J & \longleftarrow & A & \longrightarrow & J \end{array}$$

is a commutative diagram and both $\alpha \circ \beta$ and $\beta \circ \alpha$ are homotopic to the identity.

If $X = \{X^n, \delta^n\}$ is a complex, then its *cohomology* is the graded group $\{H^n(X)\}$ where $H^n(X) = \ker \delta^n / \operatorname{im} \delta^{n-1}$. A morphism $\alpha : X \to Y$ of complexes induces a morphism $H(\alpha) : H(X) \to H(Y)$ of cohomology. Two morphisms which are homotopic induce the same morphism of cohomology. Also, if two morphisms α and β are homotopic and F is a functor into another abelian category, then $F(\alpha)$ and $F(\beta)$ are also homotopic. These facts and Theorem 9.2 imply the following:

9.4 Corollary. *If $\alpha : A \to B$ is a morphism, $i : A \to I$ and $j : B \to J$ are injective resolutions of A and B and F is a functor from our category to another abelian category, then the morphisms $\tilde{\alpha} : I \to J$ of Theorem 9.2 induce a single, well defined morphism $\alpha^* : H(F(I)) \to H(F(J))$. This is an isomorphism if α is an isomorphism.*

In any reasonable category with enough injectives there is a way of assigning to each object A a particular injective resolution $I(A)$. In a small category this can be done using the axiom of choice. We shall show how to do this in the case of the category of sheaves. In many cases this can be done in such a fashion that $A \to I(A)$ is a functor – ideally an exact functor. This is nice when it can be done (and it can be for sheaves) but it is not necessary for the developement of the theory. In any case, assuming that we are working in a category with enough injectives and with some way of assigning an injective resolution to an object, we may construct the higher derived functors of a left exact functor F as follows: the n-th derived functor of A is

$$R^n F(A) = H^n(F(I(A)))$$

the n-th cohomology of the complex obtained by applying the functor F to the complex of injectives $I(A)$. Of course, by Corollary 9.4, different choices of resolutions $I(A)$ will yield isomorphic objects $R^nF(A)$ but, in order that the R^nF be functors, it is important that a specific way of making the assignment $A \to I(A)$ be available. That R^nF is a functor then follows from Corollary 9.4.

The fact that F is left exact means that if $A \to I^0 \to I^1 \to \dots$ is an injective resolution, then on applying F exactness is preserved at the first two terms; that is,

$$0 \to F(A) \to F(I^0) \to F(I^1)$$

is exact. From this it follows that:

9.5 Theorem. *If F is a left exact functor from an abelian category with enough injectives to an abelian category, then there is an isomorphism of functors $F \to R^0F$.*

If (as is the case for the category of sheaves) $A \to I(A)$ is an exact functor then an exact sequence of objects

$$0 \to A \to B \to C \to 0$$

yields an exact sequence of complexes

$$0 \to I(A) \to I(B) \to I(C) \to 0$$

because the objects in these complexes are injective, this short exact sequence splits and, hence, remains exact when we apply a left exact functor F. This yields a short exact sequence of complexes

$$0 \to F(I(A)) \to F(I(B)) \to F(I(C)) \to 0$$

which, in turn, yields a long exact sequence of cohomology

$$0 \to F(A) \to F(B) \to F(C) \to R^1F(A) \to R^1F(B) \to R^1F(C) \to R^2F(A) \to \cdots$$

However, to get this result it is not neccessary to assume that $A \to I(A)$ is exact or is even a functor. In fact, given injective resolutions $A \to I(A)$ and $C \to I(C)$ one may construct an injective resolution $B \to J$ along with morphisms $I(A) \to J$ and $J \to I(C)$ such that

$$0 \to I(A) \to J \to I(C) \to 0$$

is an exact sequence of complexes. To do this, we set $J^n = I^n(A) \oplus I^n(C)$ and define morphisms $j : B \to J^0$ and $\delta^n : J^n \to J^{n+1}$ as follows: Using the fact that $I^0(A)$ is injective, we extend $i_A : A \to I^0(A)$ to a morphism $j_1 : B \to I^0(A)$. We let $j_2 : B \to I^0(C)$ be the composition of $B \to C$ with $i_C : C \to I^0(C)$. Then $j = j_1 \oplus j_2$. Clearly j is a monomorphism of B into J^0 and the diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \\
 & & {\scriptstyle i_A}\big\downarrow & & {\scriptstyle j}\big\downarrow & & {\scriptstyle i_C}\big\downarrow & & \\
0 & \longrightarrow & I^0(A) & \longrightarrow & J^0 & \longrightarrow & I^0(C) & \longrightarrow & 0
\end{array}$$

is commutative. We repeat this argument with $i_A : A \to I^0(A)$ and $i_C : C \to I^0(C)$ replaced by $\operatorname{coker} i_A \to I^1(A)$ and $\operatorname{coker} i_C :\to I^1(C)$, respectively, and obtain a commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & A \arrow[r] \arrow[d, "i_A"'] & B \arrow[r] \arrow[d, "j"'] & C \arrow[r] \arrow[d, "i_C"'] & 0 \\
0 \arrow[r] & I^0(A) \arrow[r] \arrow[d, "\delta_A^0"'] & J^0 \arrow[r] \arrow[d, "\delta^0"'] & I^0(C) \arrow[r] \arrow[d, "\delta_C^0"'] & 0 \\
0 \arrow[r] & I^1(A) \arrow[r] & J^1 \arrow[r] & I^1(C) \arrow[r] & 0
\end{tikzcd}$$

Continuing in this way, we construct an injective resolution $B \to J$ of B and morphisms of complexes $I(A) \to J$ and $J \to I(C)$ for which the diagram

$$\begin{tikzcd}
0 \arrow[r] & A \arrow[r] \arrow[d, "i_A"'] & B \arrow[r] \arrow[d, "j"'] & C \arrow[r] \arrow[d, "i_C"'] & 0 \\
0 \arrow[r] & I(A) \arrow[r] & J \arrow[r] & I(C) \arrow[r] & 0
\end{tikzcd}$$

is commutative with exact rows. On applying any left exact functor F we conclude, without the assumption that $I(\cdot)$ is exact or even a functor, that

9.6 Theorem. *If F is a left exact functor from an abelian category with enough injectives to an abelian category, and if*

$$0 \to A \to B \to C \to 0$$

is a short exact sequence in the first category, then there is a long exact sequence

$$0 \to F(A) \to F(B) \to F(C) \to R^1F(A) \to R^1F(B) \to R^1F(C) \to R^2F(A) \to \cdots$$

In this result, the morphisms $R^nF(A) \to R^nF(B)$ and $R^nF(B) \to R^nF(C)$ are just those induced by $A \to B$ and $B \to C$ as in Corollary 9.4; i. e. they are the images of these morphisms under the functor R^nF. The connecting morphisms $R^nF(C) \to R^{n+1}F(A)$ a priori depend on the choices made in the construction of J. In fact, they do not depend on these choices. They are well defined and depend functorially on the short exact sequence $0 \to A \to B \to C \to 0$. To prove this requires some diagram chasing which we choose not to do here. It is done in any number of books on homological algebra.

Of course, one does not really compute the objects $R^nF(A)$ using injective resolutions. In practice, one uses the long exact sequence to reduce the computation of $R^nF(A)$ for complicated objects A to that for simpler objects or one uses Theorem 9.8 below which often allows one to compute $R^nF(A)$ using much simpler resolutions of A.

An object C is said to be *acyclic* for the functor F if $R^nF(C) = 0$ for $n > 0$.

9.7 Theorem. *Let F be as in Theorem 9.6. Then injective objects are F-acyclic.*

Proof. If I is injective, then $0 \to I \to I \to 0$ is an injective resolution of I. It follows from Corollary 9.4 that $R^nF(I)$ may be computed by applying F to this resolution and taking cohomology. Thus, $R^nF(I) = 0$ for $n > 0$ and I is acyclic.

9.8 Theorem. *Let F be as in Theorem 9.6 and suppose that $A \to J$ is a resolution of A by a complex of F-acyclic objects. Then there is an isomorphism $R^nF(A) \to H^n(F(J))$ for each n.*

Proof. Let $K^n = \ker\{J^n \to J^{n+1}\}$ and consider the long exact sequences for $\{R^nF\}$ determined by the short exact sequences

$$0 \to K^p \to J^p \to K^{p+1} \to 0, \quad p \geq 0$$

where $K^0 = A$. Using the fact that $R^nF(J^p) = 0$ for $n > 0$, we conclude from these long exact sequences that

$$R^qF(K^p) \simeq R^{q-1}F(K^{p+1}), \quad p \geq 0,\ q > 1$$

and

$$R^1F(K^{n-1}) \simeq F(K^n)/\operatorname{im}\{F(J^{n-1}) \to F(K^n)\} \simeq H^n(J)$$

An induction argument then shows that $R^nF(A) = R^nF(K^0) \simeq H^n(J)$, as required.

At this point we leave general homological algebra and return to the study of sheaves. It is convenient to suppose we have a fixed sheaf of rings $\mathcal{R}$ on a space X and then to study the abelian category consisting of all sheaves of $\mathcal{R}$-modules. A sheaf of $\mathcal{R}$-modules is a sheaf $\mathcal{M}$ on X such that for each open set $U \subset X$, $\mathcal{M}(U)$ is a $\mathcal{R}(U)$-module in such a way that the module action $\mathcal{R}(U) \times \mathcal{M}(U) \to \mathcal{M}(U)$ commutes with restriction. An abelian group may be regarded as a module over the ring of integers. Therefore, an ordinary sheaf of abelian groups is a sheaf of modules over a sheaf of rings if for the sheaf of rings we choose the sheaf of germs of the constant presheaf $U \to \mathbb{Z}$.

It is well known that the category of modules over a ring has enough injectives and, in fact, that there is a functor which assigns to each module a monomorphism into an injective module. We give a brief description of the construction:

An abelian group is injective if and only if it is divisible (for every element g and every integer n there is an element h so that $nh = g$). If G is an abelian group, let $F(G)$ be the free abelian group generated by G as a set. Thus, an element of $F(G)$ is a function $f : G \to \mathbb{Z}$ which is zero except at finitely many points and the group operation is pointwise addition of functions. There is a group homomorphism $F(G) \to G$ defined by $f \to \sum f(g)g$. Let $K(G)$ be its kernel. Then let $Q(G)$ be the free module over the rationals generated by G (finitely non-zero funtions from G to Q) and note that $Q(G)$ contains $F(G)$ and, hence, $K(G)$ as subgroups. Finally, we set $D(G) = Q(G)/K(G)$. Then $D(G)$ is a divisible group containing a copy of G as a subgroup. Now if R is any ring and M is a module over R then we consider M as an abelian group and construct $D(M)$. Then the R-module

$$\hom_Z(R, D(M))$$

is injective (problem 9.3) and there is a monomorphism $M \to \hom_Z(R, D(M))$ defined by $m \to (r \to rm)$ where M is regarded as a subgroup of $D(M)$. Thus, we have defined a functorial way to assign to each R-module M a monomorphism of M into an injective module.

If R happens to be an algebra over a field k, then $M \to \hom_k(R, M)$ is a simpler way of embedding each R-module into an injective R-module.

Now if we have a sheaf of modules $\mathcal{M}$ over a sheaf of rings $\mathcal{R}$ on X it is a simple matter to embed $\mathcal{M}$ in an injective sheaf of $\mathcal{R}$-modules. For each stalk $\mathcal{M}_x$ of $\mathcal{M}$ let $\tilde{\mathcal{M}}_x$ be the injective $\mathcal{R}_x$ module containing $\mathcal{M}_x$ constructed using one of the methods described above. Then $\mathcal{I}^0(\mathcal{M})$ is the sheaf defined as follows: $\mathcal{I}^0(\mathcal{M})(U)$ is the $\mathcal{R}(U)$ module consisting of all functions which assign to each $x \in U$ an element of $\tilde{\mathcal{M}}_x$. Note that there is no requirement that these functions be continuous. Clearly $\mathcal{M}$ is embedded in $\mathcal{I}^0(\mathcal{M})$ as those functions which are continuous and have values in the modules $\mathcal{M}_x$.

9.9 Theorem. *The sheaf $\mathcal{I}^0(\mathcal{M})$ is an injective object in the category of sheaves of $\mathcal{R}$-modules.*

Proof. Suppose $\alpha : \mathcal{S} \to \mathcal{T}$ is a monomorphism of sheaves of $\mathcal{R}$ modules and $\beta : \mathcal{S} \to \mathcal{I}^0(\mathcal{M})$ is a morphism of sheaves of $\mathcal{R}$-modules. Then β determines a function ϕ which assigns to each $x \in X$ an element $\phi_x \in \hom_{\mathcal{R}_x}(\mathcal{S}_x, \tilde{\mathcal{M}}_x)$ as follows: $\phi_x(s_x)$ is the value at x of the germ $\beta_x(s_x) \in \mathcal{I}^0(\mathcal{M})$. Then, for an open set U and a section $s \in S(U)$ we have that $\beta(s) \in \mathcal{I}^0(\mathcal{M})$ is the function on U defined by $\beta(s)(x) = \phi_x(s_x)$. By the injectivity of $\tilde{\mathcal{M}}_x$ the morphism $\phi_x : \mathcal{S}_x \to \tilde{\mathcal{M}}$ has an extension to a morphism $\psi_x : \mathcal{T}_x \to \tilde{\mathcal{M}}$ for each x (we fix one such extension for each x using the axion of choice). If we set $\gamma(t)(x) = \psi_x(t_x)$ for U open and $t \in \mathcal{T}(U)$ then we have defined a morphism $\gamma : \mathcal{T} \to \mathcal{I}^0(\mathcal{M})$ which is an extension of β, that is $\gamma \circ \alpha = \beta$. Thus, $\mathcal{I}^0(\mathcal{M})$ is injective in the category of sheaves of $\mathcal{R}$-modules.

9.10 Corollary. *If $\mathcal{R}$ is a sheaf of rings on X then the category of sheaves of $\mathcal{R}$-modules on X has enough injectives.*

In fact, it is clear that the above construction can be used to define a functor which assigns to each $\mathcal{R}$-module $\mathcal{M}$ an injective resolution $\mathcal{M} \to \mathcal{I}(\mathcal{M})$.

Since we have enough injectives, we know that for any left exact functor from the category sheaves of $\mathcal{R}$-modules to an abelian category there are corresponding higher derived functors for which Theorems 9.5, 9.6, 9.7 and 9.8 hold.

9.11 Definition. *If $\mathcal{M}$ is a sheaf of $\mathcal{R}$-modules on X, $Y \subset X$ is a subset of X and ϕ a family of closed subsets of X closed under finite union, then we set*

(i) $H^p(Y; \mathcal{M}) = R^p\Gamma(Y; \mathcal{M})$;
(ii) $H^p_Y(X; \mathcal{M}) = R^p\Gamma_Y(X; \mathcal{M})$;
(iii) $H^p_\phi(X; \mathcal{M}) = R^p\Gamma_\phi(X; \mathcal{M})$

for each p. These groups are called the p th sheaf cohomology group of $\mathcal{M}$ on Y, of $\mathcal{M}$ on X with support in Y and of $\mathcal{M}$ on X with supports in ϕ, repectively.

If $f : Y \to X$ is a continuous map then there are also higher derived functors, $R^p f_*(\cdot)$, for the direct image functor, though we won't bother to give them special names. Here, if $\mathcal{R}$ is a sheaf of rings on X then $f^{-1}\mathcal{R}$ is a sheaf of rings on Y and the functor f_* from sheaves on Y to sheaves on X maps a sheaf of $f^{-1}\mathcal{R}$-modules to a sheaf of $\mathcal{R}$-modules.

One of the most common uses of the long exact sequence theorem is to derive a relation between the cohomology of a sheaf on X and on a closed subset $Y \subset X$. Thus, let $\mathcal{M}$

be a sheaf of $\mathcal{R}$-modules on X, let $i : Y \to X$ be the inclusion of a closed subset and let $U = X - Y$. Consider the sheaf $\mathcal{M}_Y = i_* i^{-1}\mathcal{M}$ on X. This is the sheaf $V \to \mathcal{M}(Y \cap V)$. Thus, an intuitive description of the construction of $\mathcal{M}_Y$ is "restrict $\mathcal{M}$ to Y and then extend by zero to X". There is an epimorphism from $\mathcal{M}$ to $\mathcal{M}_Y$ given by $s \to s|_{Y\cap V}$: $\mathcal{M}(V) \to \mathcal{M}_Y(V) = \mathcal{M}(Y \cap V)$. In fact $\mathcal{M}_Y$ is the unique quotient of $\mathcal{M}$ whose stalks at points of Y agree with those of $\mathcal{M}$ and whose stalks at points not in Y are zero. We will call the kernel of this map $\mathcal{M}_U$ since its intuitive description is also "restrict $\mathcal{M}$ to U and then extend by zero to all of X. However, its precise description is quite different from that of $\mathcal{M}_Y$. It is the sheaf of germs of the presheaf that sends V to $\mathcal{M}(V)$ if $V \subset U$ and to zero otherwise. It is the unique subsheaf of $\mathcal{M}$ which has stalks equal to those of M at points of U and zero at points not in U.

Since we have a short exact sequence

$$0 \to \mathcal{M}_U \to \mathcal{M} \to \mathcal{M}_Y \to 0$$

we have a corresponding long exact sequence of cohomology.

$$\cdots \to H^p(X;\mathcal{M}_U) \to H^p(X;\mathcal{M}) \to H^p(X;\mathcal{M}_Y) \to H^{p+1}(X;\mathcal{M}_U) \to \cdots$$

It remains to interpret the meaning of $H^p(X;\mathcal{M}_U)$ and $H^p(X;\mathcal{M}_Y)$. These can be interpreted as the pth derived functors, applied to $\mathcal{M}$ of the functors $\mathcal{M} \to \Gamma(X;M_U)$ and $\mathcal{M} \to \Gamma(X;M_Y)$, repectively. Then it is easy to see that (Problem 9.4)

$$H^p(X;\mathcal{M}_Y) = H^p(Y;\mathcal{M})$$

and

$$H^p(X;\mathcal{M}_U) = H^p_\phi(U;\mathcal{M})$$

where ϕ is the family of supports in U consisting of subsets of U which are closed in X. We thus have the following:

9.12 Theorem. *Let $\mathcal{M}$ be a sheaf of $\mathcal{R}$-modules on X, Y a closed subset of X and $U = X - Y$. Then there is a long exact sequence:*

$$\cdots \to H^p_\phi(U;\mathcal{M}) \to H^p(X;\mathcal{M}) \to H^p(Y;\mathcal{M}) \to H^{p+1}_\phi(U;\mathcal{M}_U) \to \cdots$$

where ϕ is the family of supports consisting of subsets of U which are closed in X.

From Theorem 9.8 we know that we can compute the sheaf cohomology groups from any resolution by sheaves that are acyclic for the given functor $(\Gamma, \Gamma_Y, \Gamma_\phi)$. Thus, we will devote the remainder of this chapter to describing various classes of acyclic sheaves and dicussing some examples.

9.13 Definition. *Let $\mathcal{M}$ be a sheaf of $\mathcal{R}$-modules on X then*

(a) *$\mathcal{M}$ is called flabby if for every open set $U \subset X$ the restriction map $\Gamma(X;\mathcal{M}) \to \Gamma(U;\mathcal{M})$ is surjective;*
(b) *$\mathcal{M}$ is called soft if for every closed set $Y \subset X$ the restriction map $\Gamma(X;\mathcal{M}) \to \Gamma(Y;\mathcal{M})$ is surjective;*
(c) *$\mathcal{M}$ is called fine if for every locally finite open cover $\{U_i\}_{i\in I}$ of X there is a family of morphisms $\{\phi_i : \mathcal{M} \to \mathcal{M}\}_{i\in I}$ such that ϕ_i is supported in U_i for each i and $\sum_i \phi_i = id$. The family $\{\phi_i\}_{i\in I}$ is called a partition of unity for the sheaf $\mathcal{M}$ subordinate to the cover $\{U_i\}$.*

The definition of fine sheaf needs some comment. The statement that ϕ_i is supported in U_i means that there is a closed subset K_i of U_i such that $\phi_x = 0$ for all $x \in X - K_i$. The sum $\sum_i \phi_i$ makes sense as a morphism from $\mathcal{M}$ to $\mathcal{M}$ because the open cover $\{U_i\}$ is locally finite. This and the fact that ϕ_i is supported in U_i implies that, in a sufficiently small neighborhood of each point, only finitely many terms of $\sum_i \phi_i$ are non zero and, hence, the sum makes sense in such a neighborhood; but then these local morphisms fit together to define a morphism of sheaves globally.

In what follows, the term Γ-*acyclic* will mean acyclic for the functor $\Gamma(X;\cdot)$.

9.14 Theorem. *If X is any topological space, then in the category of sheaves of $\mathcal{R}$-modules on X*

(i) *if the first two terms of a short exact sequence are flabby then so is the third;*
(ii) *injective $\Rightarrow$ flabby $\Rightarrow$ Γ-acyclic;*

Proof. For U open and $\mathcal{M}$ a sheaf of $\mathcal{R}$-modules on X, let $\mathcal{R}_U$ be the restriction of $\mathcal{R}$ to U followed by extension by zero to all of X. Then we have an inclusion $\mathcal{R}_U \to \mathcal{R}$. Thus, if $\mathcal{M}$ is injective then each morphism from $\mathcal{R}_U$ to $\mathcal{M}$ extends to a morphism from $\mathcal{R}$ to $\mathcal{M}$. But a morphism $\mathcal{R}_U \to \mathcal{M}$ is just a section of $\mathcal{M}$ over U and a morphism $\mathcal{R} \to \mathcal{M}$ is just a section of $\mathcal{M}$ over X (Problem 9.5). Thus, sections of $\mathcal{M}$ over U extend to sections of $\mathcal{M}$ over X. Thus, injective $\Rightarrow$ flabby.

Now suppose that $\mathcal{A}$ is a flabby sheaf and

$$0 \longrightarrow \mathcal{A} \xrightarrow{\alpha} \mathcal{B} \xrightarrow{\beta} \mathcal{C} \longrightarrow 0$$

is an exact sequence of sheaves of $\mathcal{R}$-modules. We wish to prove that

$$0 \longrightarrow \Gamma(X;\mathcal{A}) \xrightarrow{\alpha} \Gamma(X;\mathcal{B}) \xrightarrow{\beta} \Gamma(X;\mathcal{C}) \longrightarrow 0$$

is also exact. Of course, we need only prove that $\beta : \Gamma(X;\mathcal{B}) \to \Gamma(X;\mathcal{C})$ is surjective. To this end, we let c be an element of $\Gamma(X;\mathcal{C})$ and consider the class of pairs (U, b) where U is an open subset of X, $b \in \Gamma(U;\mathcal{B})$ and $\beta(b) = c|_U$. This class is non-empty since c is locally in the image of β and it is partially ordered under the relation: $(U_1, b_1) < (U_2, b_2)$ if $U_1 \subset U_2$ and $g_1 = g_2|_{U_1}$. It also has the property that a maximal totally ordered subset has a maximal element (by taking union). Thus, it follows from Zorn's Lemma that there is a maximal element (U, b) in this class. If $U = X$ we are done. If not then there is an $x \in X - U$. We may choose a neighborhood V of x and an element $b^0 \in \Gamma(V;\mathcal{B})$ such that $\beta(b_0) = c|_V$. We then have $\beta(b|_{U\cap V} - b_0|_{U\cap V}) = 0$ and so there is an element $a_0 \in \Gamma(U \cap V;\mathcal{A})$ such that $\alpha(a) = b_{U\cap V} - b_0|_{U\cap V}$. We use the fact that $\mathcal{A}$ is flabby to extend a_0 to a section a of $\mathcal{A}$ on all of V. Then b and $b_0 - \alpha(a)$ agree when restricted to $U \cap V$ and, hence, define a section $b' \in \Gamma(U';\mathcal{B})$ where $U' = U \cup V$. But then (U', b') is in our class and is larger than (U, b), contradicting the maximality. This proves the exactness of the above sequence.

We may now prove part(i) of the Theorem. If $\mathcal{A}$ and $\mathcal{B}$ are flabby in the above short exact sequence, then their restrictions to an open set U are also flabby. It follows from the above paragraph and the fact that $\mathcal{A}|_U$ is flabby that any section c of $\mathcal{C}$ over U is $\beta(b)$ for some section of $\mathcal{B}$ over U. But, since $\mathcal{B}$ is flabby, the section b is the restriction to U of a global section b' of $\mathcal{B}$. Then $\beta(b')$ provides an extension of c to a global section of $\mathcal{C}$.

Now to prove that a flabby sheaf $\mathcal{A}$ is Γ-acyclic we embed it in an injective $\mathcal{I}$ and use the long exact sequence of cohomology for the short exact sequence

$$0 \longrightarrow \mathcal{A} \xrightarrow{\alpha} \mathcal{I} \xrightarrow{\beta} \mathcal{C} \longrightarrow 0$$

where α is the inclusion and C is its cokernel. Since $\mathcal{I}$ is Γ-acyclic we have that

$$H^p(X;\mathcal{A}) \simeq H^{p-1}(X;\mathcal{C}) \quad p > 1$$

and, by what we proved in the first part of the argument, $H^1(X;\mathcal{A}) = 0$. Also, $\mathcal{C}$ is flabby because $\mathcal{A}$ and $\mathcal{I}$ are flabby. Thus, $H^2(X;\mathcal{A}) \simeq H^1(X;\mathcal{C}) = 0$. By iterating this argument we conclude that $H^p(X;\mathcal{A}) = 0$ for all $p > 0$. Thus, we have proved that flabby $\Rightarrow$ acyclic.

There is a similar result with a similar proof for soft sheaves. It is slightly more complicated and requires that the space be paracompact.

9.15 Theorem. *If X is paracompact, then for sheaves of $\mathcal{R}$-modules on X*

(i) *if the first two terms of a short exact sequence are soft then so is the third;*
(ii) *flabby $\Rightarrow$ soft $\Rightarrow$ Γ-acyclic;*
(iii) *fine $\Rightarrow$ soft $\Rightarrow$ Γ-acyclic.*

Proof. Let $U \subset X$ be open and let $\mathcal{M}$ be a sheaf of $\mathcal{R}$-modules on X. Since X is paracompact, By Theorem 8.12, every section of a sheaf on a closed set Y extends to an open set containing Y and, hence, extends to all of X if the sheaf is flabby. Thus, flabby $\Rightarrow$ soft.

Now suppose that $\mathcal{A}$ is a soft sheaf and

$$0 \longrightarrow \mathcal{A} \xrightarrow{\alpha} \mathcal{B} \xrightarrow{\beta} \mathcal{C} \longrightarrow 0$$

is an exact sequence of sheaves of $\mathcal{R}$-modules. We wish to prove that

$$0 \longrightarrow \Gamma(X;\mathcal{A}) \xrightarrow{\alpha} \Gamma(X;\mathcal{B}) \xrightarrow{\beta} \Gamma(X;\mathcal{C}) \longrightarrow 0$$

is also exact. Of course, we need only prove that $\beta : \Gamma(X;\mathcal{B}) \to \Gamma(X;\mathcal{C})$ is surjective. To this end, let c be a section of $\mathcal{C}$ over X and note that c is locally in the image of β thus, there is an open cover $\{U_i\}$ of X, which we may assume locally finite since X is paracompact, and elements $b_i \in \Gamma(X;\mathcal{U}_i)$ such that $\beta(b_i) = c|_{U_i}$. Since paracompact spaces are normal, we may choose an open cover V_i of X which has the property that $\bar{V}_i \subset U_i$. We then consider the class of pairs (Y, b) where Y is a subset of X which is a union of some of the sets in the collection $\bar{V}_i$, $b \in \Gamma(Y;\mathcal{B})$ and $\beta(b) = c|_Y$. Any set Y from such a pair is closed due to the fact that the collection $\{\bar{V}_i\}$ is locally finite. This class is non-empty since it contains each pair $(\bar{V}_i, b_i)$ and it is partially ordered under the relation: $(Y_1, b_1) < (Y_2, b_2)$ if $Y_1 \subset Y_2$ and $g_1 = g_2|_{Y_1}$. It also has the property that a maximal totally ordered subset has a maximal element (by taking union). Thus, it follows from Zorn's Lemma that there is a maximal element (Y, b) in this class. If $Y = X$ we are done. If not then there is an

i so that $\bar{V}_i$ is not contained in Y. We then have $\beta(b|_{Y\cap\bar{V}_i} - b_i|_{Y\cap\bar{V}_i}) = 0$ and so there is an element $a_0 \in \Gamma(Y \cap \bar{V}_i; \mathcal{A})$ such that $\alpha(a_0) = b|_{Y\cap\bar{V}_i} - b_i|_{Y\cap\bar{V}_i}$. We use the fact that $\mathcal{A}$ is soft to extend a_0 to a section a of $\mathcal{A}$ on all of $\bar{V}_i$. Then b and $b_i - \alpha(a)$ agree when restricted to $Y \cap \bar{V}_i$ and, hence, define a section $b' \in \Gamma(Y'; \mathcal{B})$ where $Y' = Y \cup \bar{V}_i$. But then (Y', b') is in our class and is larger than (Y, b), contradicting the maximality. This proves the exactness of the above sequence. This result may be used to prove part (i) (Problem 9.6). Then to prove that $\mathcal{A}$ is acyclic we embed it in an injective $\mathcal{I}$ and use the long exact sequence of cohomology for the short exact sequence

$$0 \longrightarrow \mathcal{A} \xrightarrow{\alpha} \mathcal{I} \xrightarrow{\beta} \mathcal{C} \longrightarrow 0$$

where α is the inclusion and C is its cokernel. Since $\mathcal{I}$ is acyclic we have that

$$H^p(X; \mathcal{A}) \simeq H^{p-1}(X; \mathcal{C}) \quad p > 1$$

and, by what we proved in the first part of the argument, $H^1(X; \mathcal{A}) = 0$. Also, $\mathcal{C}$ is soft because $\mathcal{A}$ and $\mathcal{I}$ are soft. Thus, $H^2(X; \mathcal{A}) \simeq H^1(X; \mathcal{C}) = 0$. By iterating this argument we conclude that $H^p(X; \mathcal{A}) = 0$ for all $p > 0$. Thus, we have proved that soft $\Rightarrow$ acyclic.

It remains to prove that fine $\Rightarrow$ soft. Thus, let $Y \subset X$ be closed and let $\mathcal{M}$ be a fine sheaf on X. If $s \in \Gamma(Y; \mathcal{M})$ then for each $x \in Y$ the germ of s at x is represented by a section defined in a neighborhood of x which agrees with s when restricted to that neighborhood intersected with Y. Thus, we may choose an open cover $\{U_i\}$ of X and elements $s_i \in \Gamma(U_i; \mathcal{M})$ such that $s|_{U_i\cap Y} = s_i|_{U_i\cap Y}$ for each i - one of these open sets will be the complement of Y and will have the zero section assigned to it and the others will be neighborhoods of points of Y. Because X is paracompact, we may assume that $\{U_i\}$ is locally finite. Now because $\mathcal{M}$ is fine, we may choose for each i an endomorphism $\phi_i : \mathcal{M} \to \mathcal{M}$ supported in U_i in such a way that $\sum_i \phi_i = id$. For each i we interpret $\phi_i s_i$ to be a section on all of X by extending it to be zero on the complement of U_i. We then set $s' = \sum_i \phi_i s_i \in \Gamma(X, \mathcal{M})$. That this sum makes sense follows from the local finiteness of the open cover which means that in a neighborhood of any point we are summing only finitely many non-zero terms. We also have that $s'_x = s_x$ at each point $x \in Y$ so that s' is an extension of s to all of X. It follows that $\mathcal{M}$ is soft. Thus, we have proved that fine $\Rightarrow$ soft and completed the proof of the Theorem.

We conclude this chapter with some examples of acyclic resolutions which show that certain classical cohomology theories are just examples of sheaf cohomology.

By $\mathcal{C}^0$ we shall mean the sheaf of continuous functions on X. If X is the appropriate kind of differentiable manifold, we shall denote by $\mathcal{C}^p$ and $\mathcal{C}^\infty$ the sheaves of functions with continuous partial derivatives up to order p and functions with continuous partial derivatives of all orders, respectively.

9.16 Theorem. *Each of $\mathcal{C}^0$, $\mathcal{C}^p$ and $\mathcal{C}^\infty$ is a fine sheaf if X is a paracompact space and, in the case of $\mathcal{C}^p$ and $\mathcal{C}^\infty$, X has the appropriate differentiable manifold structure.*

Proof. Paracompact implies normal which implies that Urysohn's Lemma holds. Urysohn's Lemma can be used to constuct continuous partitions of unity subordinate to any locally finite open cover. A continuous function with support inside a given open set defines, by multiplication, an endomorphism of $\mathcal{C}^0$ with support in the open set. Thus, a partition of unity in the algebra of continuous functions on X defines a partition of unity for the sheaf $\mathcal{C}^0$ in the sense of definition 9.13(c). This proves that $\mathcal{C}^0$ is a fine sheaf. The same result for $\mathcal{C}^p$ and $\mathcal{C}^\infty$ follows from the fact that on a $\mathcal{C}^p$ or $\mathcal{C}^\infty$ manifold partitions of unity may be chosen to be comprised of $\mathcal{C}^p$ or $\mathcal{C}^\infty$ functions.

One of the nice things about the notion of fine sheaf is that if a sheaf of rings $\mathcal{R}$ is fine then so is any sheaf of modules over $\mathcal{R}$ (Problem 9.7). In particular, any sheaf of modules over $\mathcal{C}^p, \quad (p = 0, 1, \cdots, \infty)$ is fine and, hence, Γ-acyclic. In particular, the sheaf of continuous sections of a vector bundle on a topological space is a fine sheaf as is the sheaf of $\mathcal{C}^\infty$ differential forms of degree p on a $\mathcal{C}^\infty$-manifold.

Let X be a $\mathcal{C}^\infty$-manifold and denote by $\mathcal{E}^P$ the sheaf of $\mathcal{C}^\infty$ differential forms of degree p on X. The de Rham complex of sheaves is the complex

$$0 \longrightarrow \mathbb{C} \longrightarrow \mathcal{E}^0 \xrightarrow{d^0} \mathcal{E}^1 \xrightarrow{d^1} \cdots \xrightarrow{d^{p-1}} \mathcal{E}^p \xrightarrow{d^p} \cdots$$

where d^p is exterior differentiation and $\mathbb{C}$ here stands for the constant sheaf with stalks $\mathbb{C}$. The Poincaré Lemma says that if X is any open ball in $\mathbb{R}^n$ then the corresponding sequence of sections is exact (this is proved by constructing an explicit homotopy between the identity and zero using integration along lines from the center of the ball). Since a $\mathcal{C}^\infty$ manifold looks locally like a ball in $\mathbb{R}^n$ it follows that the de Rham complex is exact as a complex of sheaves. Thus, it defines a resolution $\mathbb{C} \to \mathcal{E}$ of the constant sheaf $\mathbb{C}$ by a complex $\mathcal{E}$ of fine sheaves. On passing to global sections of $\mathcal{E}$, we obtain the classical de Rham complex $\mathcal{E}(X)$ of differential forms on X. The cohomology of this complex is called the de Rham cohomology (with coefficients in $\mathbb{C}$) of X. By theorem 9.8 we have:

9.17 Theorem. *There is a natural isomorphism between the de Rham cohomology of a $\mathcal{C}^\infty$ manifold X and the sheaf cohomology $H(X;\mathbb{C})$ for the constant sheaf $\mathbb{C}$ on X.*

We end this chapter with a discussion of Čech cohomology and its relation to sheaf cohomology. Let $\mathcal{S}$ be a sheaf of $\mathcal{R}$-modules on X and let $\mathcal{U} = \{U_i\}_{i\in I}$ be an open cover of X. If $\alpha = (i_0, \cdots, i_p) \in I^{p+1}$ is a multi-index then we set $U_\alpha = U_{i_0} \cap \cdots \cap U_{i_p}$. Then for each open set $U \subset X$, the module of *Čech p-cochains* on U, for the cover $\mathcal{U}$, is the direct product $\prod_{\alpha\in I^{p+1}} \Gamma(U_\alpha \cap U, \mathcal{S})$ and is denoted $\mathcal{C}^p(\mathcal{U}, \mathcal{S})(U)$. In other words, a p-cochain for $\mathcal{U}$ on U is a function f which assigns to each $\alpha \in I^{p+1}$ an element $f(\alpha) \in \Gamma(U_\alpha \cap U; \mathcal{S})$. If $V \subset U$ then restriction clearly defines a morphism $\mathcal{C}^p(\mathcal{U}, \mathcal{S})(U) \to \mathcal{C}^p(\mathcal{U}, \mathcal{S})(V)$ and, thus, $U \to \mathcal{C}^p(\mathcal{U}, \mathcal{S})(U)$ is a presheaf. In fact, it is clearly a sheaf. We denote it by $\mathcal{C}^p(\mathcal{U}, \mathcal{S})$. The module of global sections of this sheaf is $\mathcal{C}^p(\mathcal{U}, \mathcal{S})(X)$. This is the classical space of Čech cochains for $\mathcal{S}$ and the cover $\mathcal{U}$ and will be denoted $C^p(\mathcal{U}, \mathcal{S})$. We next define a coboundary mapping

$$\delta^p : \mathcal{C}^p(\mathcal{U}, \mathcal{S}) \to \mathcal{C}^{p+1}(\mathcal{U}, \mathcal{S})$$

by

$$\delta^p f(\alpha) = \sum_{j=0}^{p+1} (-1)^j f(\alpha_j)|_{U_\alpha \cap U}$$

where $f \in \mathcal{C}^p(\mathcal{U}, \mathcal{S})(U)$ and $\alpha_j \in I^{p+1}$ is obtained from $\alpha \in I^{p+2}$ by deleting its j th entry.

9.18 Theorem. *We have $\delta^{p+1} \circ \delta^p = 0$ so that*

$$0 \longrightarrow \mathcal{C}^0(\mathcal{U},\mathcal{S}) \xrightarrow{\delta^0} \mathcal{C}^1(\mathcal{U},\mathcal{S}) \xrightarrow{\delta^1} \cdots \xrightarrow{\delta^{p-1}} \mathcal{C}^p(\mathcal{U},\mathcal{S}) \xrightarrow{\delta^p} \cdots$$

is a complex of sheaves and sheaf morphisms.

Proof. For $\alpha \in I^{p+3}$ let $\alpha_{j,k}$ denote the result of deleting both the j th and the k th entries from α. Then

$$(\alpha_j)_k = \alpha_{j,k} \quad \text{if} \quad 0 \le k < j \le p+2$$

and

$$(\alpha_j)_k = \alpha_{j,k+1} \quad \text{if} \quad 0 \le j < k \le p+1$$

It follows from this that for $f \in \mathbb{C}^p(\mathcal{U},\mathcal{S})(U)$ we have

$$\begin{aligned}
\delta^{p+1} \circ \delta^p f(\alpha) &= \sum_{j=0}^{p+2} (-1)^j \left[\sum_{k=0}^{p+1} (-1)^k f((\alpha_j)_k)|_{U_{\alpha_j} \cap U} \right] |_{U_\alpha \cap U} \\
&= \sum_{k<j} (-1)^{j+k} f(\alpha_{j,k})|_{U_\alpha \cap U} + \sum_{k \ge j} (-1)^{j+k} f(\alpha_{j,k+1})|_{U_\alpha \cap U} \\
&= 0
\end{aligned}$$

due to the fact that, in the middle line above, the second term is equal to the negative of the first, which is evident from the change of variables $j \to k, \quad k+1 \to j$ applied to the second term and the observation that $\alpha_{j,k} = \alpha_{k,j}$.

The restriction maps $\mathcal{S}(U) \to \mathcal{S}(U_i \cap U)$ define a sheaf morphism $\epsilon : \mathcal{S} \to \mathcal{C}^0(\mathcal{U},\mathcal{S})$ whose composition with δ^0 is 0. In fact, much more is true:

Theorem 9.19. *If $\mathcal{U}$ is an open cover of X and $\mathcal{S}$ a sheaf on X then the complex*

$$0 \longrightarrow \mathcal{S} \xrightarrow{\epsilon} \mathcal{C}^0(\mathcal{U},\mathcal{S}) \xrightarrow{\delta^0} \mathcal{C}^1(\mathcal{U},\mathcal{S}) \xrightarrow{\delta^1} \cdots \xrightarrow{\delta^{p-1}} \mathcal{C}^p(\mathcal{U},\mathcal{S}) \xrightarrow{\delta^p} \cdots$$

is an exact sequence of sheaves.

Proof. Fix an $x \in X$. Let U be any neighborhood of x which is contained in a member of $\mathcal{U}$, say $U \subset U_k$. Now suppose $f \in \mathcal{C}^p(\mathcal{U},\mathcal{S})(U)$ and $\delta^p f = 0$. Define $g \in \mathcal{C}^{p-1}(\mathcal{U},\mathcal{S})(U)$ by

$$g(\beta) = f((k,\beta))$$

where we set $(k,\beta) = (k, i_0, \dots, i_{p-1}) \in I^{p+1}$ for $\beta = (i_0, \dots, i_{p-1}) \in I^p$. Note that $g \in \mathcal{C}^{p-1}(\mathcal{U},\mathcal{S})(U)$ due to the fact that $U \subset U_k$ which implies that $U_k \cap U_\beta \cap U = U_\beta \cap U$. Then for $\alpha \in I^{p+1}$,

$$\begin{aligned}
0 = \delta^p f((k,\alpha)) = f(\alpha)|_{U_k \cap U_\alpha \cap U} - \sum_{j=0}^{p} (-1)^j f((k,\alpha_j))|_{U_k \cap U_\alpha \cap U} \\
= f(\alpha) - \delta^{p-1} g(\alpha)
\end{aligned}$$

again due to the fact that $U_k \cap U_\alpha \cap U = U_\alpha \cap U$. This shows that, locally, the kernel of δ^p is the image of δ^{p-1} for $p > 0$. The same argument also works for $p = 0$ with δ^{-1} replaced by ϵ. This proves the exactness of the above sequence.

The complex of Theorem 9.19 is called the Čech complex of sheaves and is denoted $\mathcal{C}(\mathcal{U}, \mathcal{S})$. Its complex of global sections, $\Gamma(\mathcal{C}(\mathcal{U}, \mathcal{S})) = C(\mathcal{U}, \mathcal{S})$, is what is classically called the Čech complex and we shall call the *global Čech complex* for the sheaf $\mathcal{S}$ and the open cover $\mathcal{U}$.

9.20 Definition. *The Čech cohomology $\{\check{H}^p(\mathcal{U}, \mathcal{S})\}_{p\geq 0}$ of the sheaf $\mathcal{S}$ for the cover $\mathcal{U}$ is the cohomology of the global Čech complex $C(\mathcal{U}, \mathcal{S})$.*

9.21 Theorem. *For any sheaf $\mathcal{S}$ and any open cover $\mathcal{U}$ we have*
(i) *$\check{H}^0(\mathcal{U}, \mathcal{S}) \simeq \Gamma(X, \mathcal{S})$; and*
(ii) *if $\mathcal{S}$ is flabby, then $\check{H}^p(\mathcal{U}, \mathcal{S}) = 0$ for $p > 0$.*

Proof. By Theorem 9.19 we have an exact sequence of sheaves

$$0 \longrightarrow \mathcal{S} \xrightarrow{\epsilon} \mathcal{C}^0(\mathcal{U}, \mathcal{S}) \xrightarrow{\delta^0} \mathcal{C}^1(\mathcal{U}, \mathcal{S}),$$

which implies that

$$0 \longrightarrow \Gamma(X, \mathcal{S}) \xrightarrow{\epsilon} C^0(\mathcal{U}, \mathcal{S}) \xrightarrow{\delta^0} C^1(\mathcal{U}, \mathcal{S}),$$

is exact and this implies statement (i).

The sheaf $\mathcal{C}^p(\mathcal{U}, \mathcal{S})$of Čech p-cochains is the direct product of the sheaves $(i_\alpha)_* i_\alpha^{-1}\mathcal{S}$ where $i_\alpha : U_\alpha \to X$ is the inclusion map. If $\mathcal{S}$ is flabby then so is its restriction, $i_\alpha^{-1}\mathcal{S}$, to U and the direct image of a flabby sheaf is flabby by Problem 9.8. It follows that $(i_\alpha)_* i_\alpha^{-1}\mathcal{S}$ is flabby for each α and, hence, that $\mathcal{C}^p(\mathcal{U}, \mathcal{S})$ is flabby for each p. Thus, the complex in Theorem 9.19 is an exact sequence of flabby sheaves. Then Theorem 9.14 implies that the complex obtained from it by applying $\Gamma(X, \cdot)$ is also exact. Part (ii) follows.

Let $\mathcal{S}$ be a sheaf on X, $\mathcal{U}$ an open cover of X and $\mathcal{S} \to \mathcal{C}(\mathcal{U}, \mathcal{S})$ the corresponding Čech resolution as in Theorem 9.19. Let $\mathcal{S} \to \mathcal{I}$ be an injective resolution of $\mathcal{S}$. Then the injectivity of the terms in $\mathcal{I}$ can be used, as in Theorem 9.2, to inductively construct a morphism of complexes $\mathcal{C}(\mathcal{U}, \mathcal{S}) \to \mathcal{I}$ for which the diagram

$$\begin{tikzcd} \mathcal{S} \arrow[r] \arrow[d, equal] & \mathcal{C}(\mathcal{U}, \mathcal{S}) \arrow[d] \\ \mathcal{S} \arrow[r] & \mathcal{I} \end{tikzcd}$$

is commutative. This morphism is unique up to homotopy and, hence, after we apply $\Gamma(X, \cdot)$ it determines a well defined morphism $\check{H}^p(\mathcal{U}, \mathcal{S}) \to H^p(X, \mathcal{S})$.

The open cover $\mathcal{U}$ is called a *Leray covering* for the sheaf $\mathcal{S}$ if for each multi-index α the sheaf $\mathcal{S}$ is acyclic on U_α. For example, it follows from the Poincaré lemma that the deRham cohomology of any convex open set in $\mathbb{R}^n$ vanishes in positive degrees. From this and Theorem 9.17 we conclude that the constant sheaf $\mathbb{C}$ is acyclic on any convex open set in $\mathbb{R}^n$. This implies that any open cover of $\mathbb{R}^n$ by open convex sets is a Leray cover for the constant sheaf $\mathbb{C}$.

9.22 Theorem. *If $\mathcal{U}$ is a Leray cover for the sheaf $\mathcal{S}$ then the natural morphism*

$$\check{H}^p(\mathcal{U},\mathcal{S}) \to H^p(X,\mathcal{S})$$

is an isomorphism.

Proof. Choose an embedding $\mathcal{S} \to \mathcal{F}$ of $\mathcal{S}$ in a flabby sheaf (an injective, for example). Let $\mathcal{G}$ be the cokernel and consider the short exact sequence:

$$0 \longrightarrow \mathcal{S} \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow 0$$

Then from the long exact sequence and the fact that $\mathcal{S}$ is acyclic on $\mathcal{U}_\alpha$ we conclude that

$$0 \longrightarrow \mathcal{S}(U_\alpha) \longrightarrow \mathcal{F}(U_\alpha) \longrightarrow \mathcal{G}(U_\alpha) \longrightarrow 0$$

is exact for each multi-index α and from this that we have an exact sequence

$$0 \longrightarrow C(\mathcal{U},\mathcal{S}) \longrightarrow C(\mathcal{U},\mathcal{F}) \longrightarrow C(\mathcal{U},\mathcal{G}) \longrightarrow 0$$

of global Čech complexes. This, in turn, gives us a long exact sequence of Čech cohomology, which, along with the long exact sequence for sheaf cohomology, the induced morphisms from Čech to sheaf cohomology discussed above and the fact that $\mathcal{F}$ is flabby, gives us the following commutative diagram with exact rows:

$$\begin{tikzcd}
0 \arrow[r] & \check{H}^0(\mathcal{U},\mathcal{S}) \arrow[r] \arrow[d] & \check{H}^0(\mathcal{U},\mathcal{F}) \arrow[r] \arrow[d] & \check{H}^0(\mathcal{U},\mathcal{G}) \arrow[r] \arrow[d] & \check{H}^0(\mathcal{U},\mathcal{S}) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \Gamma(X,\mathcal{S}) \arrow[r] & \Gamma(X,\mathcal{F}) \arrow[r] & \Gamma(X,\mathcal{G}) \arrow[r] & H^1(X,\mathcal{S}) \arrow[r] & 0
\end{tikzcd}$$

The first three vertical arrows are isomorphisms and, hence, so is the fourth. This establishes our result in the case $p = 1$. We prove the general case using the following commutative diagram, which is also part of the the diagram which comes from the long exact sequences associated to the above short exact sequence:

$$\begin{tikzcd}
0 \arrow[r] & \check{H}^p(\mathcal{U},\mathcal{G}) \arrow[r] \arrow[d] & \check{H}^{p+1}(\mathcal{U},\mathcal{S}) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & H^p(X,\mathcal{G}) \arrow[r] & H^{p+1}(X,\mathcal{S}) \arrow[r] & 0
\end{tikzcd}$$

for $p > 0$. This has exact rows and so if the first vertical morphism is an isomophism so is the second. Consider the class of all sheaves $\mathcal{T}$ such that the open cover $\mathcal{U}$ is a Leray cover for $\mathcal{T}$. From the long exact sequence for sheaf cohomology it follows that if the first two sheaves in a short exact sequence belong to this class then so does the third. By hypothesis, $\mathcal{S}$ belongs to this class and $\mathcal{F}$ does since it is flabby. It follows that $\mathcal{G}$ belongs as well. Thus, if we have proved the theorem for all members of this class and for all degrees less thatn or equal to p then the above diagram shows that it is true for degree $p+1$ as well. By induction, the proof is complete.

What we have been discussing so far is Čech cohomology for an open cover and a sheaf. Čech cohomology for a sheaf without reference to a cover is obtained by passing to a limit over covers. More precisely, if $\mathcal{V}$ is an open cover which is a refinement of the open cover $\mathcal{U}$ then the restriction maps define a morphism of complexes of sheaves $\mathcal{C}(\mathcal{U},\mathcal{S}) \to \mathcal{C}(\mathcal{V},\mathcal{S})$. By passing to the limit over the directed set of open covers of X we obtain a complex of sheaves $\mathcal{C}(\mathcal{S})$, which we shall call the limit Čech complex of sheaves, and a quasi-isomorphism $\mathcal{S} \to C(\mathcal{S})$. The complex of global sections of $\mathcal{C}(\mathcal{S})$ is $C(\mathcal{S}) = \mathcal{C}(\mathcal{S})(X) = \varinjlim C(\mathcal{U},\mathcal{S})$ and is called the limit global Čech complex. Its cohomology is the Čech cohomology of $\mathcal{S}$ on X and is denoted $\{\check{H}^p(X,\mathcal{S})\}$. Clearly the morphisms $\check{H}^p(\mathcal{U},\mathcal{S}) \to H^p(X,\mathcal{S})$ induce a morphism $\check{H}^p(X,\mathcal{S}) \to H^p(X,\mathcal{S})$.

9.23 Theorem. *If X is paracompact then the natural morphism*

$$\check{H}^p(X,\mathcal{S}) \to H^p(X,\mathcal{S})$$

is an isomorphism.

Proof. . We shall prove that $\mathcal{S} \to C(\mathcal{S})$ is a resolution of $\mathcal{S}$ by soft sheaves. Since soft sheaves are acyclic, the result will then follow from Theorem 9.8.

We have already remarked that $\mathcal{S} \to C(\mathcal{S})$ is a quasi-isomorphism – in other words, that the exactness of the sequences in Theorem 9.19 is preserved on passing to the direct limit – since direct limits generally preserve exactness. It remains to prove that each $\mathcal{C}^p(\mathcal{S})$ is a soft sheaf. Thus, let $Y \subset X$ be closed and suppose f is a section of $\mathcal{C}^p(\mathcal{S})$ over Y. Since X is paracompact and Y is closed, we have that f may be represented by a section in a neighborhood U of Y. This section may, in turn, be represented by an element $f' \in \mathcal{C}^p(\mathcal{U},\mathcal{S})(U)$ for some open cover $\mathcal{U}$ of X. We may choose a locally finite refinement $\mathcal{V}$ of $\mathcal{U}$ with the property that each set in $\mathcal{V}$ is either contained in U or is contained in $X - Y$. We define a section g' of $\mathcal{C}^p(\mathcal{V},\mathcal{S})$ on X by setting g' equal to the image of f' under the refinement map on those multi-indices for which all the correspondings sets in $\mathcal{V}$ are contained in U and setting it equal to zero otherwise. On passing to the image g of g' in the space of limit Čech p-cochains, we obtain a global section of $\mathbb{C}^p(\mathcal{S})$ which has f as its restriction to Y. Thus, $\mathbb{C}^p(\mathcal{S})$ is soft and the proof is complete.

When $\mathcal{S}$ is a constant sheaf G. The Čech cohomology in the sense of this chapter is just the classical Čech cohomology with coeficients in the group G from the theory of algebraic topology.

We end this section with an example which shows how to solve one of the local to global problems posed in chapter 8. This is the problem of finding a global logarithm for a non-vanishing continuous function on X. We assume X is paracompact. Let $\mathcal{C}$ denote the sheaf of continuous functions on X with addition as group operation and $\mathcal{C}^{-1}$ the sheaf of invertible (non-vanishing) continuous functions with multiplication as group operation. Then, due to the fact that a non-vanishing continuous function has a logarithm locally in a neighborhood of each point, the sequence of sheaves:

$$0 \longrightarrow 2\pi i\mathbb{Z} \longrightarrow \mathcal{C} \xrightarrow{\exp} \mathcal{C}^{-1} \longrightarrow 0$$

is exact. Since $\mathcal{C}$ is a fine sheaf and hence acyclic, we conclude from the long exact sequence of cohomology for this short exact sequence that

$$\mathcal{C}^{-1}(X)/\exp(\mathcal{C}(X)) \simeq H^1(X,\mathbb{Z}) \simeq \check{H}^1(X,\mathbb{Z})$$

where, of course, $\mathbb{Z}$ stands for the constant sheaf $\mathbb{Z}$ and we use the fact that $2\pi i\mathbb{Z} \simeq \mathbb{Z}$. Thus, every non-vanishing continuous function on X has a global logarithm if and only if the first Čech cohomology of X with integral coeficients vanishes. More generally, there is a epimorphism $f \to [f]$ from the group of non-vanishing continuous function on X to the Čech cohomology group $\check{H}^1(X,\mathbb{Z})$ with the property that f has a global logarithm if and only if $[f]$ vanishes. This is an elementary but very instructive example of the use of sheaf theory to analyze a problem involving passing from local to global solutions.

9. Problems

1. Prove Theorem 9.1.
2. Prove Theorem 9.2.
3. Prove that $\hom_Z(R,D)$ is an injective R-module if D is a divisible abelian group.
4. With $\mathcal{M}_U$ and $\mathcal{M}_Y$ as defined in the discussion preceding Theorem 9.12, prove that
$$H^p(X;\mathcal{M}_Y) = H^p(Y;\mathcal{M})$$
and
$$H^p(X;\mathcal{M}_U) = H^p_\phi(U;\mathcal{M})$$
where ϕ is the family of supports in U consisting of subsets of U which are closed in X.
5. Prove that if U is an open subset of X, $\mathcal{M}$ a sheaf of $\mathcal{R}$-modules on X and $\mathcal{R}|_U$ is the extension by zero of the restriction of $\mathcal{R}$ to U, then $\hom(\mathcal{R}_U,\mathcal{M}) \simeq \Gamma(U;\mathcal{M})$.
6. Finish the proof of Theorem 9.15 by proving that if the first two terms of a short exact sequence of sheaves are soft then so is the third. Hint: you may use the fact, proved in 9.15, that if $\mathcal{A}$ is soft and $0 \to \mathcal{A} \to \mathcal{B} \to \mathcal{C} \to 0$ is a short exact sequence of sheaves on X then $\Gamma(X;\mathcal{B}) \to \Gamma(X;\mathcal{C})$ is surjective.
7. Prove that if a sheaf of rings is fine then so is every sheaf of modules over this sheaf of rings.
8. Prove that if $f : Y \to X$ is any continuous map between topological spaces then $f_*\mathcal{S}$ is a flabby sheaf on X if $\mathcal{S}$ is a flabby sheaf on Y.
9. Prove that $H^2(X,\mathbb{Z})$ is isomorphic to the group, under tensor product, formed by equivalence classes of complex line bundles over X.

10. Coherent Algebraic Sheaves

Up to this point we have only worked with algebraic or holomorphic subvarieties of open sets in $\mathbb{C}^n$. It is time to define the general notions of algebraic and holomorphic varieties and to introduce and study coherent sheaves on such objects.

10.1 Definition. *A ringed space is a pair $(X, \mathcal{R})$ consisting of a topological space X and a sheaf of rings $\mathcal{R}$ on X. A morphism of ringed spaces is a pair $(f, f^\#) : (X, \mathcal{R}) \to (Y, \mathcal{S})$ consisting of a continuous map $f : X \to Y$ together with a morphism of sheaves of rings $f^\# : \mathcal{S} \to f_*\mathcal{R}$. The sheaf of rings $\mathcal{R}$ is called the structure sheaf of the ringed space $(X, \mathcal{R})$.*

It is easy to see that a morphism $(f, f^\#) : (X, \mathcal{R}) \to (Y, \mathcal{S})$ of ringed spaces is an isomorphism (has a two sided inverse) if and only if f is a homeomorphism and $f^\#$ is an isomorphism of sheaves of rings.

If V is an algebraic (holomorphic) subvariety of a domain in $\mathbb{C}^n$ then, henceforth, ${}_V\mathcal{O}$ (${}_V\mathcal{H}$) will denote the sheaf of regular (holomorphic) functions on V. That is, for each open subset $U \subset V$ we let ${}_V\mathcal{O}(U)$ (${}_V\mathcal{H}(U)$) be the space of regular (holomorphic) functions on U.

A subvariety V (algebraic or holomorphic) of an open set (Zariski or Euclidean) in $\mathbb{C}^n$ is a ringed space with the obvious structure sheaf (${}_V\mathcal{O}$ or ${}_V\mathcal{H}$).

10.2 Definition. *A holomorphic variety is a ringed space which is locally isomorphic to a holomorphic subvariety of an open set in $\mathbb{C}^n$ and whose topological space is Hausdorff and second countable. If a holomorphic variety is locally isomorphic, as a ringed space, to an open set in $\mathbb{C}^n$ then it is called a complex manifold.*

The definition of algebraic variety is somewhat more complicated due to the need for a condition which replaces Hausdorff. We begin by defining the notion of affine variety.

10.3 Definition. *An affine variety is a ringed space which is isomorphic to an algebraic subvariety of $\mathbb{C}^n$.*

In most of the algebraic geometry literature an affine variety is defined to be an *irreducible* subvariety of $\mathbb{C}^n$, but this is too restrictive for our purposes.

10.4 Theorem. *If V is an affine variety and f a regular function on V, we set $V_f = \{z \in V : f(z) \neq 0\}$. Then the open subset V_f is also an affine variety.*

Proof. Suppose V is an algebraic subvariety of $\mathbb{C}^n$ and consider the algebraic subvariety W of $\mathbb{C}^{n+1}$ defined by

$$W = \{(z, z_{n+1}) \in \mathbb{C}^{n+1} : 1 - f(z) \cdot z_{n+1} = 0, \quad z \in V\}$$

Then the map $z \to (z, f(z)^{-1}) : V_f \to W$ is a morphism of ringed spaces (Problem 10.4) with inverse $(z, z_{n+1}) \to z : W \to V_f$. Thus, V_f is affine since it is isomorphic to W.

An *algebraic prevariety* is an ringed space which has a finite cover by open sets which are affine varieties. Note that by Theorem 10.4 each open subset of an affine variety V is, itself, a finite union of open sets which are affine varieties (sets of the form V_f). Thus, an open subset of a prevariety is a prevariety. A closed subset of an affine variety is also clearly an affine variety and so a closed subset of a algebraic prevariety is also an algebraic prevariety.

We need one more condition to define the class of algebraic varieties. This is to eliminate pathological examples like the following: Consider the space V which is two copies of the complex plane glued together by the identity map everywhere except at 0. Under the Zariski topology – the topology in which open sets are complements of finite sets – this is a union of two open sets each of which is a copy of $\mathbb{C}$ with the Zariski topology. Thus, it is an algebraic prevariety. However, it is a strange space since it contains two copies of the origin and just one copy of every other point of the plane. Note that if $f, g : \mathbb{C} \to V$ are the embeddings of the two copies of $\mathbb{C}$ into V then each is a morphism of ringed spaces but $\{z \in \mathbb{C} : f(z) = g(z)\} = \mathbb{C} - 0$ is not a closed subset of $\mathbb{C}$. Thus, this space V will not be an algebraic variety if we use the following definition:

10.5 Definition. *An algebraic variety V is an algebraic prevariety with the property that for any algebraic prevariety W and any pair of morphisms $f : W \to V$ and $g : W \to V$ it is true that $\{w \in W : f(w) = g(w)\}$ is a closed set.*

It turns out that the above condition is equivalent to the condition that the diagonal in $V \times V$ is a closed set. However, one has to be careful how one defines the product of two prevarieties V and W – as a pointset it is $V \times W$ but it does not have the cartesian product topology. One defines the product of two affine varieties to be the pointset $V \times W$ with the Zariski topology determined by the tensor product $\mathcal{O}(V) \otimes \mathcal{O}(W)$ of the corresponding rings of regular functions. This ring can be regarded as a ring of functions on $V \times W$ via the map $f \otimes g \to ((z, w) \to f(z)g(w))$. Its localizations to Zariski open sets of $V \times W$ define a presheaf whose sheaf of germs is the sheaf of regular functions for $V \times W$. It is easy to see that the product in this sense of two affine varieties is an affine variety. Once the product of affine varieties is defined then one defines the topology on the product of general algebraic prevarieties by using the cover by products of open affine subvarieties. We leave the problem of proving that an algebraic prevariety V is an algebraic variety if and only if the diagonal is closed in $V \times V$ as a problem (Problem 10.2).

Using the coordinate functions and subtraction in $\mathbb{C}^n$, it is easy to see that affine varieties are algebraic varieties. In fact, the following theorem shows that there are lots of algebraic varieties.

10.6 Theorem. *If V is an algebraic prevariety with the property that for any two points x and y in V there is an open affine subvariety of V containing both x and y, then V is an algebraic variety.*

Proof. Let $f, g : W \to V$ be two morphisms from a prevariety W to V and let $Z = \{w \in W : f(w) = g(w)\}$. For $w \in W - Z$ choose $x = f(w)$, $y = g(w)$ and let U be an open affine subvariety of V containing both x and y. Then $Q = f^{-1}(U) \cap g^{-1}(U)$ is an open set in W containing w. However, f and g both map Q into the affine variety U and so the subset of Q on which they agree, $Z \cap Q$, is closed in Q. Its complement in Q is, thus, an open set containing w and missing Z and, hence, Z is closed.

We will denote the structure sheaf of any algebraic variety V by ${}_V\mathcal{O}$ and of any holomorphic variety V by ${}_V\mathcal{H}$. The corresponding algebras of global sections will be denoted

$\mathcal{O}(V)$ and $\mathcal{H}(V)$.

In what follows we will need a precise description of the global regular functions on an affine variety. Recall that, by definition, a regular function f on a subvariety V of an open set in $\mathbb{C}^n$ is a function with the property that for each point $z \in V$ there is a neighborhood U of z in $\mathbb{C}^n$ and a rational function on $\mathbb{C}^n$ with denominator non-vanishing on U such that f agrees with this rational function on $U \cap V$. However, it is a theorem that regular functions on a subvariety V of $\mathbb{C}^n$ are actually restrictions to V of polynomials on $\mathbb{C}^n$:

Theorem 10.7. *If V is a subvariety of $\mathbb{C}^n$ and $A = \mathbb{C}[z_1, \dots, z_n]/\operatorname{id}(V)$ is the algebra of restrictions to V of polynomials on $\mathbb{C}^n$, then*

(i) *as a topological space, V is the space of maximal ideals of A with its Zariski topology;*
(ii) *if $f \in A$ and $V_f = \{v \in V : f(v) \neq 0\}$ then the algebra ${}_V\mathcal{O}(V_f)$ of regular functions on V_f is the localization A_f of A relative to the multiplicative set $\{f^n : n \geq 0\}$;*
(iii) *in particular, the algebra, $\mathcal{O}(V)$, of regular functions on V is the algebra, A, of restrictions to V of polynomials on $\mathbb{C}^n$;*

Proof. Each point v of V determines a maximal ideal $M_v = \{f \in A : f(v) = 0\}$. On the other hand, if M is a maximal ideal of A, then M is generated by a finite set $f_1, \cdots, f_m/inA$. If there is a point $v \in V$ where the functions f_i all vanish, then $M \subset M_v$ and, hence, $M = M_v$ since M is maximal. Thus, to prove (1) it suffices to show that if $f_1, \dots, f_m \in A$ have no common zeroes on V then there are elements $g_1, \dots, g_m \in A$ such that $\sum f_i g_i = 1$; in other words, a finite set of functions with no common zero lies in no maximal ideal.

To see that this condition holds, choose $h_1, \cdots, h_m \in \mathbb{C}[z_1, \cdots, z_n]$ such that $h_j|V = f_j$ for $j = 1, \cdots, n$ and choose $h_{m+1}, \cdots, h_k \in \mathbb{C}[z_1, \cdots, z_n]$ to be a set of polynomials for which V is exactly its set of common zeroes. Then the set $h_1, \cdots, h_k$ has no common zeroes on $\mathbb{C}^n$; in other words, the ideal, I, which it generates has $\operatorname{loc}(I) = \emptyset$. It follows from the Nullstellensatz that $\sqrt{I} = \mathbb{C}[z_1, \cdots, z_n]$. But this implies that the identity is in $\sqrt{I}$ and, hence, in I. Thus, the equation $\sum p_j h_j = 1$ has a solution. When we restrict the h_j's to V, the last $k - m$ restrict to zero. Thus, the restriction to V of $p_1, \cdots, p_m$ gives the required solution $g_1, \cdots, g_m$ to the equation $\sum g_i f_i = 1$. This shows that $v \to M_v$ is a bijection from points of V to maximal ideals of A. The topology of V is the relative topology inherited from $\mathbb{C}^n$. This is the topology in which closed sets are common zero sets of families of polynomials. Thus, the closed subsets of V are common zero sets of families of polynomials restricted to V. This is the Zariski topology induced on V by A. Thus, we have proved part (i).

One direction of (ii) is trivial: Restriction to V_f defines a natural map $A_f \to {}_V\mathcal{O}(V_f)$. Suppose $g/f^k \in A_f$ $(g \in A)$ determines the zero section of ${}_V\mathcal{O}$ over V_f. This means that $g = 0$ on V_f which implies that $fg = 0$ in A which implies that $g/f^k = 0$ in A_f. Thus, $A_f \to {}_V\mathcal{O}(V_f)$ is injective. The surjectivity will be proved after we prove (iii).

We now prove (iii). Certainly the restriction to V of a polynomial on $\mathbb{C}^n$ is a regular function on V and so $A \subset \mathcal{O}(V)$. Thus, we must show that the restriction map is surjective. Suppose now that $h \in \mathcal{O}(V)$. We can cover V with a collection $\{U_i\}$ of open sets such that $f|_{U_i}$ has the form $p_i/q_i, \in \mathcal{O}(U_i)$ where p_i and q_i are elements of A and q_i does not vanish on U_i. Since A is Noetherian we may assume the collection $\{U_i\}$ is finite. Furthermore, we may assume each set U_i is of the form V_g for some $g \in A$ since these sets form a base for the topology. In fact we may actually assume that $U_i = V_{q_i} = \{v \in V : q_i(V) \neq 0\}$ since if this is not true it can be achieved by replacing p_i and q_i by rp_i and rq_i for an appropriate $r \in A$. Then the condition that the p_i/q_i define a global section of ${}_V\mathcal{O}$ is that $p_i/q_i = p_j/q_j$ on $V_{q_i} \cap V_{q_j}$ for each pair i, j. This means that $(p_iq_j - p_jq_i) = 0$ on $V_{q_i} \cap V_{q_j}$ which, means, as in the previous paragraph, that its product with q_iq_j is the zero element of A. Thus, if we set $p'_i = p_iq_i$ and $q'_i = q_i^2$ for each i, then $p'_i/q'_i = p_i/q_i$ in A_{q_i} for each i and for each pair i, j the equation $p'_iq'_j = p'_jq'_i$ holds in A. Since the q'_j have no common zeroes on V, by part (i) we may choose a set $\{g_i\} \subset A$ such that $\sum g_iq'_i = 1$. Then

$$q_j \sum g_ip'_i = p'_j \sum g_iq'_i = p'_j$$

or $q'_jh' = p'_j$ where $h' = \sum g_ip'_i \in A$. Then clearly the image of h' in $\mathcal{O}(V)$ is h. Thus, $A \to \mathcal{O}(V)$ is surjective and (iii) is proved.

We showed above that $A_f \to {}_V\mathcal{O}(V_f)$ is injective. To prove the surjectivity in (ii) we simply apply (iii) to the image of V_f under the map $z \to (z, f(z)^{-1}) : V_f \to \mathbb{C}^{n+1}$. As we showed in Theorem 10.4, this map is an isomorphism of V_f onto a subvariety W of $\mathbb{C}^{n+1}$. The composition of this map with a polynomial on $\mathbb{C}^{n+1}$ is the restriction to V_f of a function of the form g/f^k, where g is a polynomial on $\mathbb{C}^n$. This shows that $A_f \to {}_V\mathcal{O}(V_f)$ is surjective and finishes the proof of (ii).

There are many operations on sheaves which yield new sheaves. If $\mathcal{M}$ and $\mathcal{N}$ are sheaves of modules over a sheaf of rings $\mathcal{A}$, then $U \to \hom(\mathcal{M}|_U, \mathcal{N}|_U)$ is a presheaf which yields a sheaf $hom_{\mathcal{A}}(\mathcal{M}, \mathcal{N})$ when we pass to its sheaf of germs. One can define the tensor product $\mathcal{M} \otimes_{\mathcal{A}} \mathcal{N}$ of two sheaves of modules over a sheaf of rings $\mathcal{A}$ or the tensor product $\mathcal{M} \otimes_A \mathcal{N}$ of two sheaves of modules over a fixed ring A. These sheaves are defined as the sheaves of germs of the obvious presheaves. One such construction, localization, is the key tool of algebraic geometry. We shall describe it in some detail. Suppose that $\mathcal{A}$ is a sheaf of rings on a space X, $A = \Gamma(X; \mathcal{A})$ and M is an A-module. Then one can *localize* M on X by constructing the sheaf $\mathcal{A} \otimes_A M$. This is the sheaf of germs of the presheaf $U \to \mathcal{A}(U) \otimes_A M$. Clearly, $M \to \mathcal{A} \otimes_A M$ is a functor from the category of A-modules to the category of sheaves of $\mathcal{A}$-modules. While there is no reason to expect this functor to have nice properties in general, there are many situations where it is very nice indeed. What properties would we like it to have? We would like to be able to recover M from $\mathcal{A} \otimes_A M$ as $\Gamma(X; \mathcal{A} \otimes_A M)$. It would be nice if $M \to \mathcal{A} \otimes_A M$ were an exact functor or, even better, an equivalence of categories with $\Gamma(X; \cdot)$ as inverse. This is exactly what happens for the localization functor of algebraic geometry. The version that we describe below is a special case of a much more general theory.

Suppose V is an affine variety and M is a module over the algebra $\mathcal{O}(V)$. Then we set

$$\tilde{M} = {}_V\mathcal{O} \otimes_{\mathcal{O}(V)} M$$

As noted above, this is the sheaf of germs of the presheaf $U \to M(U) = {}_V\mathcal{O}(U) \otimes_{\mathcal{O}(V)} M$. Now, a basis of neighborhoods for the topology of V is given by the sets of the form $V_f = \{v \in V : f(v) \neq 0\}$ for $f \in \mathcal{O}(V)$. Thus, the sheaf $\tilde{M}$ is determined by the restriction of the presheaf $U \to M(U)$ to sets of the form $U = V_f$.

10.8 Theorem. *Let V be an affine variety, M an $\mathcal{O}(V)$-module and f an element of $\mathcal{O}(V)$. Then there is a natural isomorphism $M(V_f) \to M_f$, where M_f is the localization of M relative to the multiplicative set $\{f^k : k \geq 0\}$.*

Proof. The localization, M_f, of M relative to the multiplicative set $\{f^k : k \geq 0\}$ consists of equivalence classes of expressions of the form m/f^k where $m \in M$, with the obvious module operations. Two such expressions m/f^k and n/f^j are equivalent if $f^p(f^j m - f^k n) = 0$ for some $p \geq 0$. Note that this is the same as the localization of M relative to the multiplicative set consisting of functions in $\mathcal{O}(V)$ which are non-vanishing on V_f. This is due to the fact that if $g \in \mathcal{O}(V)$ is non-vanishing on V_f then, by the Nullstellensatz, f is in the radical of the ideal generated by g in $\mathcal{O}(V)$, from which it follows that $f^k = hg$ for some $k \geq 0$ and some $h \in \mathcal{O}(V)$. This implies that the two multiplicative sets $\{f^k : k \geq 0\} \subset \{g \in \mathcal{O}(V) : g(v) \neq 0 \quad \forall v \in V_f\}$ yield the same localization.

As noted in Theorem 10.7 the ring ${}_V\mathcal{O}(V_f)$ is the localization of $\mathcal{O}(V)$ relative to $\{f^k : k \geq 0\}$. We define the morphism

$$ {}_V\mathcal{O}(V_f) \otimes_{\mathcal{O}(V)} M = M(V_f) \to M_f $$

by

$$ g/f^k \otimes m \to gm/f^k $$

Its inverse is the morphism determined by

$$ m/f^k \to 1/f^k \otimes m : M_f \to M(V_f) $$

It is clear that both morphisms are well defined and they are inverses of one another.

10.9 Theorem. *Let V be an affine variety. Then*

(i) *for each open set $U \subset V$ of the form $U = V_f$, the functor $M \to M(U) = {}_V\mathcal{O}(U) \otimes_{\mathcal{O}(V)} M$ is exact; that is, ${}_V\mathcal{O}(U)$ is a flat $\mathcal{O}(V)$-module;*

(ii) *the localization functor $M \to \tilde{M}$ is exact.*

Proof. Using the description of $M(U)$ as ${}_V\mathcal{O}(U) \otimes_{\mathcal{O}(V)} M$ makes it clear that this functor is right exact. Thus, we need only prove that if $N \subset M$ is a submodule then $N(U) \to M(U)$ is injective. To prove this we use the fact, proved in the previous theorem, that $M(U) = M_f$ and $N(U) = N_f$ if $U = V_f$. Thus, let n/f^k represent an element of N_f and suppose that it determines the zero element of M_f. This means that $f^p \cdot n = 0$ for some p. But if this equation holds in M it also holds in N and, hence, n/f represents the zero element of N as well. This proves (i); however, (ii) is an immediate consequence of (i) and the fact that sets of the form V_f form a basis for the topology of V

10.10 Theorem. *If V is an affine variety and M is an $\mathcal{O}(V)$-module then*

(i) *on an open set of the form V_f the natural morphism $M_f = M(V_f) \to \Gamma(V_f; \tilde{M})$ is an isomorphism;*

(ii) *in particular, $M \to \Gamma(V; \tilde{M})$ is an isomorphism.*

Proof. Statements (i) and (ii) are actually equivalent since, if V is affine, so is each V_f. Thus, we will just prove (ii).

This is almost the same as the proof of Theorem 10.7 but there are a few differences. Suppose $m \in M(V)$ determines the zero section of $\tilde{M}$ over V. This means that we may cover V with finitely many open sets U_i such that $m|_{U_i} = 0$ in $\tilde{M}(U_i)$ for each i. Without loss of generality we may assume that these sets are of the form $U_i = V_{q_i}$ where $q_i \in \mathcal{O}(V)$. Then $m|_{U_i} = 0$ means that for each i there is an integer n_i such that $q_i^{n_i} m = 0$. Since the sets U_i cover V, the collection $\{q_i\}$ has no common zero on V. Since V is affine, Theorem 10.7 (i) implies there exists a set $\{g_i\}$, such that $\sum g_i q_i^{n_i} = 1$ in $\tilde{M}(V)$. However, this implies that $m = 0$ since m is killed by $q_i^{n_i}$ for every i. We conclude that $M(V) \to \Gamma(V; \tilde{M})$ is injective.

Now suppose that $s \in \Gamma(V; \tilde{M})$. We can cover V with a finite collection $\{U_i = V_{q_i}\}$ of basic open sets such that $s|_{U_i}$ is the image in $\Gamma(U_i; \tilde{M})$ of an element $m_i/q_i^{n_i}, \in \tilde{M}(U_i)$, where $m_i \in M$. In fact by relabeling each $q_i^{n_i}$ as q_i we may assume that $s|_{U_i}$ is the image in $\Gamma(U_i; \tilde{M})$ of an element of the form m_i/q_i. Since these elements fit together to form a section over V, we may assume (after refining the cover if necessary) that $(m_i/q_i)|_{U_i \cap U_j} - (m_j/q_j)|_{U_i \cap U_j} = 0$. However, $U_i \cap U_j = V_{q_i q_j}$ and so this equality means that there is a positive integer n so that

$$(q_i q_j)^n (q_j m_i - q_i m_j) = 0$$

in M for each pair i, j. If we simply relabel $q_i^n m_i$ by m_i and q_i^{n+1} by q_i, then the fractions m_i/q_i don't change but the above equality becomes simply

$$q_j m_i - q_i m_j = 0$$

Since V is affine and the sets U_i cover V, we may find elements $g_i \in {}_V\mathcal{O}(V_f)$ such that $\sum g_i q_i = 1$ on V. Then, the equation

$$q_j \left(\sum g_i m_i\right) = \left(\sum g_i q_i\right) m_j = m_j$$

holds in M. It says that $m_j = q_j m$ in M where $m = \sum g_i m_i$. Then m is an element of M which restricts to m_j/q_j on U_j for each j. In other words, s is the image of m under the map $M \to \Gamma(V; \tilde{M})$.

10.11 Definition. *A sheaf $\mathcal{M}$ of ${}_X\mathcal{O}$-modules on an algebraic variety X is called a quasi-coherent sheaf if each point of X is contained in an affine neighborhood V such that $\mathcal{M}|_V$ is isomorphic to $\tilde{M}$ for some $\mathcal{O}(V)$-module M. A quasi-coherent sheaf is called coherent if for each point of X this can be achieved with a module M which is finitely generated over $\mathcal{O}(V)$.*

Note that the structure sheaf ${}_X\mathcal{O}$ of an algebraic variety X is a coherent sheaf, since on any affine open set V it is the localization to V of the ring $\mathcal{O}(X)$. It follows that direct sums of copies of the structure sheaf are quasi-coherent and finite direct sums are coherent.

10.12 Lemma. *If V is an affine variety, $f \in \mathcal{O}(V)$, and $\mathcal{S}$ a quasi-coherent sheaf on V, then*

(i) *if $s \in \Gamma(V;\mathcal{S})$ and $s|_{V_f} = 0$ then there exists a positive integer n so that $f^n s = 0$;*
(ii) *if $t \in \Gamma(V_f;\mathcal{S})$ then there exist a positive integer n such that $f^n t$ is the restriction to V_f of a section in $\Gamma(V;\mathcal{S})$;*
(iii) *restriction defines a natural isomorphism $\Gamma(V;\mathcal{S})_f \to \Gamma(V_f;\mathcal{S})$.*

Proof. We may choose a finite collection of basic open sets $\{V_{g_i}\}$ such that, $\mathcal{S}|_{V_{g_i}} = \tilde{M}_i$ for a $\mathcal{O}(V_{g_i})$-module M_i for each i. If $s \in \Gamma(V;\mathcal{S})$ then for each i we have that $s_i = s|_{V_{g_i}} \in \Gamma(V_{g_i}, \tilde{M}_i)$ may be regarded as as element of M_i by Theorem 10.10. If $s|_{V_f} = 0$ then $(s_i)|_{V_f \cap V_{g_i}} = 0$ for each i. Since $V_f \cap V_{g_i} = V_{fg_i}$ this implies that the image of s_i in $(M_i)_f$ is zero by Theorem 10.8. It follows that $f^n s_i = 0$ in M_i for some n and each i. We may choose n independent of i since the open cover $\{V_{g_i}\}$ is finite. Then $f^n s$ is a global section of $\mathcal{S}$ which restricts to zero on each set in this cover and, hence, is the zero section. This proves part (i).

Now suppose that $t \in \Gamma(V_f;\mathcal{S})$. Then for each i, $t|_{V_f \cap V_{g_i}}$ may be regarded as an element of $(M_i)_f$ using Theorems 10.10 and 10.8 again. Thus, for each i, $t|_{V_{fg_i}}$ is a fraction with the numerator the restriction of an element $t_i \in M_i$ and the denominator a power of f. That is, there is an n so that $t_i \in M_i$ and $f^n t \in \Gamma(V,\mathcal{S})$ agree when restricted to V_{fg_i}. The integer n may be chosen independent of i. Now on $V_{g_i} \cap V_{g_j} = V_{g_i g_j}$ we have two sections of $\mathcal{S}$ which agree on $V_{fg_i g_j}$ since they both agree with the restriction of $f^n t$. By part(i) there is an integer m such that $f^m(t_i - t_j) = 0$ on $V_{g_i g_j}$. Again, we may choose m large enough to work for all i, j. But this means that the sections $f^m t_i$ on V_{g_i} agree on intersections and, thus, define a global section s of $\mathcal{S}$. Clearly, the restriction of s to V_f is $f^{n+m}t$. This completes the proof of (ii).

Part(iii) means exactly that (i) and (ii) hold.

10.13 Theorem. *If V is an affine variety, $\mathcal{S}$ is a quasi-coherent sheaf on V and $M = \Gamma(V,\mathcal{S})$, then there is a natural isomorphism*

$$\tilde{M} \to \mathcal{S}$$

of sheaves of ${}_V\mathcal{O}$-modules. Furthermore, $\mathcal{S}$ is coherent if and only if M is finitely generated.

Proof. Suppose that V_f is a basic affine open set set for which $\mathcal{S}|_{V_f}$ is the localization of a ${}_V\mathcal{O}(V_f)$-module. Then necessarily it is the localization of $\mathcal{S}(V_f) = \Gamma(V_f;\mathcal{S})$ by Theorem 10.10. The map $\tilde{M} \to \mathcal{S}$ is defined on such an open set as follows: We have natural isomorphisms $\tilde{M}(V_f) \to M_f$ by Theorem 10.10 and $M_f = \mathcal{S}(V)_f \to \mathcal{S}(V_f)$ by Theorem 10.12. The composition gives us our isomorphism on sets of the form V_f. It clearly commutes with restriction from one set of this form to a subset of this form and, since open sets of this kind, form a basis for the topology of V, this defines an isomorphism of sheaves of ${}_V\mathcal{O}$-modules from $\tilde{M} \to \mathcal{S}$.

If M is finitely generated, then $\mathcal{S}$ is coherent by definition since it is then the localization of a finitely generated module. On the other hand, if $\mathcal{S}$ is coherent then it is locally the localization of a finitely generated module. That is, we may cover V with basic open sets

V_{f_i}, such that for each i, $\mathcal{S}|_{V_{f_i}}$ is the localization of a finitely generated module necessarily isomorphic to M_{f_i}. In other words, there is a finite set $f_i \subset \mathcal{O}(V)$ with no common zeroes on V and with the property that M_{f_i} is finitely generated for each i. This implies that M is finitely generated (Problem 10.5). This completes the proof.

10.14 Corollary. *If V is an affine variety, then the functor $M \to \tilde{M}$ is an equivalence of categories from the category of $\mathcal{O}(V)$-modules to the category of quasi-coherent sheaves of ${}_V\mathcal{O}$-modules on V. The functor $\Gamma(V;\cdot)$ is its inverse functor.*

Proof. By Theorems 10.10 and 10.13, the composition of $M \to \tilde{M}$ and $\mathcal{S} \to \Gamma(V;\mathcal{S})$ in either order is naturally isomorphic to the identity. Thus, each is an equivalence of categories and they are inverses of one another.

Our next major result is that quasi-coherent sheaves on an affine variety are Γ-acyclic. This will follow easily from Theorem 10.16. But first we need to prove the following technical lemma:

10.15 Lemma. *If A is a Noetherian ring, I an injective A-module and K an ideal of A, then the submodule $J \subset I$ defined by $J = \{x \in I : K^n x = 0 \quad \text{for some} \quad n\}$ is also injective.*

Proof. To prove that J is injective, it suffices to prove that if $N \subset M$ are A-modules with M finitely generated, then every morphism $N \to J$ extends to a morphism $M \to J$ (Problem 10.6). Thus, suppose that $\phi : N \to J$ is such a morphism. Then since $\phi(N) \subset J$ and N is also finitely generated, we may choose a fixed n such that $\phi(K^n N) = K^n \phi(N) = 0$. By Krull's Theorem (Problem 10.7), there is an integer m such that $K^m M \cap N \subset K^n N$. Thus, ϕ factors through $N \to N/(K^m M \cap N)$. Since I is injective, the map $N/(K^m M \cap N) \to J \subset I$ induced by ϕ extends to a morphism $\psi : M/K^m M \to I$. However, the image of ψ lies in J since it is, necessarily, killed by K^n. Then the composition of ψ with $M \to M/K^n M$ is the required extension of ϕ. This completes the proof.

10.16 Theorem. *Let V be an affine variety. If I is an injective module over the ring $\mathcal{O}(V)$, then $\tilde{I}$ is a flabby sheaf on V.*

Proof. We first show that for any $f \in \mathcal{O}(V)$ the natural map $I \to I_f$ is surjective. To this end, let x/f^n be an element of I_f, with $x \in I$ and n a non-negative integer. We consider the morphism $f^{n+1}g \to fgx : f^{n+1}\mathcal{O}(V) \to I$. This is well defined since if $f^{n+1}g$ is zero then so is fg and, hence, fgx. Since I is injective, this morphism extends to a morphism $\phi : \mathcal{O}(V) \to I$ such that $\phi(f^{n+1}g) = fgx$. Then $f^{n+1}y = fx$ if $y = \phi(1)$. However, this implies that x/f^n is the image of y under the localization map $I \to I_f$. Thus, $I \to I_f$ is surjective. This proves that the restriction map $\Gamma(V;\tilde{I}) \to \Gamma(U;\tilde{I})$ is surjective in the case where $U \subset V$ is an open subset of the form V_f.

To complete the proof we must show that $\Gamma(V;\tilde{I}) \to \Gamma(U;\tilde{I})$ is surjective if U is an arbitrary open subset of V. Let Y be the subvariety of V which is the support of $\tilde{I}$. If $Y \cap U = \emptyset$ then we are through since the only section of $\tilde{I}$ over U is then zero. Suppose $Y \cap U \neq \emptyset$. Then there is an open set of the form $V_f \subset U$ such that $Y \cap V_f \neq \emptyset$. If $s \in \Gamma(U;\tilde{I})$, then by the first paragraph, the restriction of s to V_f is also the restriction to V_f of a global section t. Then $s = t|_U + r$ where $r \in \Gamma(U;\tilde{I})$ has its support in

$Z = V - V_f$ which is the zero set of f. That is, $r \in \Gamma_Z(U;\tilde{I})$. Suppose we can show that $\Gamma_Z(V,\tilde{I}) \to \Gamma_Z(U;\tilde{I})$ is surjective. Then $r = r'|_U$ for a section $r' \in \Gamma_Z(V;\tilde{I})$ and $s = (t + r')|_U$ for the global section $t + r' \in \Gamma(V;\tilde{I})$.

Now if $J = \{x \in I : f^n x = 0 \quad \text{for some} \quad n\}$, then $\tilde{J}$ is the subsheaf of $\tilde{I}$ consisting of sections killed by some power of f and this is exactly the subsheaf of sections with support in Z. Thus, $\Gamma_Z(V;\tilde{I}) = \Gamma(V;\tilde{J})$ and $\Gamma_Z(U;\tilde{I}) = \Gamma(U;\tilde{J})$. The support of $\tilde{J}$ is contained in $Y \cap Z = Y - V_f$ which is a proper subvariety of Y. Furthermore, J is also an injective $\mathcal{O}(V)$-module by Lemma 10.15. Thus we have reduced our problem to the same problem for a different injective module - one with smaller support. Since subvarieties of an affine variety satisfy the decending chain condition, we may reduce the problem, by induction, to the case where the support is a single point. However, any sheaf supported on a single point is flabby (Problem 10.8). This completes the proof.

10.17 Theorem. *Let V be an affine algebraic variety. Then $H^p(V;\mathcal{S}) = 0$ for $p > 0$ and for all quasi-coherent sheaves $\mathcal{S}$ on V.*

Proof. We set $M = \Gamma(V;\mathcal{S})$ and choose an injective resolution $M \to I$ of M. On localizing this, and using Corollary 10.14 and Theorem 10.16, we obtain a resolution $\mathcal{S} = \tilde{M} \to \tilde{I}$ of $\mathcal{S}$ by flabby sheaves. Thus, we obtain the cohomology of $\mathcal{S}$ by taking global sections of $\mathcal{I}$ and then taking cohomology of the resulting complex. However, by Theorem 10.14, we simply get back the resolution $M \to I$ when we apply global sections to $\mathcal{S} \to \tilde{I}$. Thus, $H^p(V : \mathcal{S}) = 0$ for all $p > 0$.

10.18 Definition. *On an algebraic variety, a sheaf of submodules (ideals) of the structure sheaf ${}_V\mathcal{O}$ is called an ideal sheaf. If an ideal sheaf is coherent as a sheaf of modules then it is called a coherent ideal sheaf. Any subvariety $Y \subset V$ determines an ideal sheaf – the sheaf of sections of ${}_V\mathcal{O}$ which vanish on Y. This is called the ideal sheaf of Y and is often denoted $\mathcal{I}_Y$.*

10.19 Theorem. *Let X be an algebraic variety and $Y \subset X$ a closed subvariety. Then the ideal sheaf $\mathcal{I}_Y$ is a coherent ideal sheaf.*

Proof. We may cover X with affine open sets V. For each such V, the algebra ${}_X\mathcal{O}(V)$ is the quotient of a polynomial algebra and, hence, Noetherian. Thus, the ideal

$$I = \mathcal{I}_Y(V) = \{g \in {}_X\mathcal{O}(V) : g(z) = 0 \quad \forall\, z \in Y \cap V\}$$

is finitely generated. We claim that $\mathcal{I}_Y(V_f) = I_f$ for any basic open set $V_f \subset V$ with $f \in {}_X\mathcal{O}(V)$. In fact, an element of $\mathcal{I}_Y(V_f)$ is a function on V_f of the form g/f^n, with $g \in {}_X\mathcal{O}(V)$, which vanishes on $Y \cap V_f$. Then $fg \in {}_X\mathcal{O}(V)$ vanishes on $Y \cap V$ and, hence, belongs to I. Then, on V_f we have

$$g/f^n = gf/f^{n+1} \in I_f.$$

This proves that, on V, the ideal sheaf $\mathcal{I}$ is the localization $\tilde{I}$ of the finitely generated ideal I. Since X may be covered by such sets V, we have proved that the ideal sheaf is coherent.

Obviously, Corollary 10.14 implies that on an affine variety a coherent sheaf of ideals $\mathcal{I}$ is the localization of an ideal, $\Gamma(V;\mathcal{I})$, of $\mathcal{O}(V)$ and, conversely, an ideal $I \subset \mathcal{O}(V)$ is the ideal of global sections of a coherent sheaf of ideals, $\tilde{I}$.

10.20 Theorem. *Let V be an algebraic variety. Then the following statements are equivalent:*

(i) *V is affine;*
(ii) *$H^p(V;\mathcal{S}) = 0$ for $p > 0$ and for all quasi-coherent sheaves $\mathcal{S}$ on V;*
(iii) *$H^1(V;\mathcal{S}) = 0$ for all coherent $\mathcal{S}$ sheaves on V.*

Proof. That (i) implies (ii) is Theorem 10.17 and the implication (ii) implies (iii) is trivial; thus, to complete the proof we must prove that (iii) implies (i).

Assume (iii) holds. The strategy is to prove that V is the space of maximal ideals of the algebra $\mathcal{O}(V)$ with the Zariski tolology and with its structure sheaf given by localization of $\mathcal{O}(V)$. Then we will prove that $\mathcal{O}(V)$ is finitely generated and, hence, is the algebra $\mathcal{O}(W)$ for some affine variety. Necessarily, then, V is isomorphic to the affine variety W by Theorem 10.7.

For $x \in V$ let U be an affine open set containing x and set $Y = V - U$. Consider the exact sequence of sheaves

$$0 \to \mathcal{I}_{Y\cup\{x\}} \to \mathcal{I}_Y \to \mathbb{C}_{\{x\}} \to 0$$

where $\mathcal{I}_{Y\cup\{x\}}$ and $\mathcal{I}_Y$ are the ideal sheaves in ${}_V\mathcal{O}$ of the subvarieties $Y \cup \{x\}$ and Y, respectively and $\mathbb{C}_{\{x\}}$ is the skyscraper sheaf which is $\mathbb{C}$ at x and zero at all other points. Since the ideal sheaf $\mathcal{I}_{Y\cup\{x\}}$ is coherent, $H^1(V;\mathcal{I}_{Y\cup\{x\}}) = 0$ and the long exact sequence of cohomology implies that

$$0 \to \Gamma(V;\mathcal{I}_{Y\cup\{x\}}) \to \Gamma(V;\mathcal{I}_Y) \to \Gamma(V;\mathbb{C}_{\{x\}}) \to 0$$

is also exact. This just means that there is a function $f \in \mathcal{O}(V)$ which vanishes on Y and does not vanish at x. In other words, every affine neighborhood U of x contains a neighborhood of x of the form V_f where f is a global section in $\mathcal{O}(V)$. This implies, in particular, that the functions in $\mathcal{O}(V)$ separate points in V. It also implies that the open sets V_f form a basis for the topology of V and, hence, that V has the Zariski topology determined by the algebra $\mathcal{O}(V)$ – that is, the topology in which the closed sets are the common zero sets of finite sets of functions from $\mathcal{O}(V)$. It also implies that we may cover V with a finite collection of affine open sets $\{V_{f_i}\}_{i=1}^n$ with $f_i \in \mathcal{O}(V)$ for each i.

We claim that every maximal ideal of the algebra $\mathcal{O}(V)$ has the form $\mathrm{id}(\{x\}) = \{f \in \mathcal{O}(V) : f(x) = 0\}$ for some point x of V. In fact if M is a maximal ideal which does not have this form then there is no point of V at which all functions in M vanish. This means the the collection on open sets $\{V_f : f \in M\}$ covers V and this, in turn, implies that some finite subcollection covers V by Problem 10.3. In other words, there is a finite set $\{f_i\} \subset M$ with no common zero on V. Given any such set $\{f_i\}$ consider the sheaf morphism

$$\oplus^n {}_V\mathcal{O} \to {}_V\mathcal{O}$$

defined by $(g_1, \dots, g_n) \to \sum_i g_i f_i$. This map is surjective, since at each x in V some f_i is non-vanishing and, hence, invertible in some neighborhood. The kernel of this map is coherent (Problem 10.9) and, thus, has vanishing first cohomology. It follows from the long exact sequence of cohomology that the map induced on global sections is also surjective and, hence, that we may solve for $g_1, \dots, g_n \in \mathcal{O}(V)$ such that $\sum g_i f_i = 1$. This contradicts the assumption that $\{f_i\}$ is contained in a maximal ideal and establishes the claim.

Since we have proved that the functions in $\mathcal{O}(V)$ separate points of V, we now have that V is exactly the set of maximal ideas of $\mathcal{O}(V)$ – as a point set and as a topological space if the maximal ideal space is given the Zariski topology.

Next, we show that for any affine open set of the form V_f we have ${}_V\mathcal{O}(V_f) \simeq \mathcal{O}(V)_f$. Suppose $g \in \mathcal{O}(V)$ and $g|_{V_f} = 0$. Then g restricts to zero in each of the open sets $V_f \cap V_{f_i}$. Since the V_{f_i} are affine, it follows from Theorem 10.7 that there is an integer n such that $f^n g$ restricts to be zero in V_{f_i} for each i. However, this means that $f^n g = 0$ in $\mathcal{O}(V)$. Now suppose that $h \in {}_V\mathcal{O}(V_f)$ and h_i is its restriction to $V_f \cap V_{f_i}$. Then, also by Theorem 10.7, there is an integer n so that, for each i, $f^n h_i = g_i$ for some section $g_i \in {}_V\mathcal{O}(V_{f_i})$. Then g_i and g_j agree on the intersection $V_{f_i} \cap V_{f_j}$ for each pair i, j and so the g_i define a global section g. Clearly $h = g/f^n$ on V_f. Thus, ${}_V\mathcal{O}(V_f)$ is the localized algebra $\mathcal{O}(V)_f$. This proves that the structure sheaf ${}_V\mathcal{O}$ of V is just the sheaf obtained by localizing the algebra $\mathcal{O}(V)$.

Next, we show that the algebra $\mathcal{O}(V)$ is finitely generated. Let $\{V_{f_i}\}$ be an affine open cover, as above, with $f_i \in \mathcal{O}(V)$ for each i. We know that each algebra ${}_V\mathcal{O}(V_{f_i}) = \mathcal{O}(V)_{f_i}$ is finitely generated since it is a quotient of a polynomial algebra by Theorem 10.7 and the fact that each V_{f_i} is affine. Thus, we may choose an integer k and elements $h_{ij} \in \mathcal{O}(V)$ such that for each i, the set $\{h_{ij}/f_i^k\}_j$ generates ${}_V\mathcal{O}(V_{f_i})$. We may also choose elements $g_i \in \mathcal{O}(V)$ such that $\sum f_i g_i = 1$. Let A be the subalgebra of $\mathcal{O}(V)$ generated by $\{f_i\} \cup \{g_i\} \cup \{h_{ij}\}$. Then, for each i we have $A_{f_i} = {}_V\mathcal{O}(V_{f_i})$. Thus, if $g \in \mathcal{O}(V)$ then we may choose an integer n and elements $p_i \in A$ such that $f_i^n g = p_i$ in $\mathcal{O}(V)$. However, the fact that the equation $\sum f_i g_i = 1$ holds in A implies that the set $\{f_i\}$ is contained in no maximal ideal of A and this implies that the set $\{f_i^n\}$ is contained in no maximal ideal of A. Hence, we may solve the equation $\sum g_i' f_i^n = 1$ for $g_i' \in A$. Then

$$g = g \sum g_i' f_i^n = \sum g_i' p_i \in A$$

Thus, $A = \mathcal{O}(V)$ and $\mathcal{O}(V)$ is finitely generated.

Finally, since $\mathcal{O}(V)$ is finitely generated, it is the quotient of a polynomial algebra and, hence, is isomorphic to $\mathcal{O}(W)$, where W is an affine variety. Now V is the maximal ideal space of $\mathcal{O}(V)$ with the Zariski topology and the structure sheaf obtained by localization of $\mathcal{O}(V)$. But W is the maximal ideal space of $\mathcal{O}(W)$ with the Zariski topology and the structure sheaf obtained by localization of $\mathcal{O}(W)$. Since the two algebras are isomorphic, V and W are isomorphic as algebraic varieties.

10. Problems

1. Prove that the product $V \times W$ of two prevarieties and the projection morphisms $\pi_V : V \times W \to V$ and $\pi_V : V \times W \to V$ have the following universal property: For each prevariety Q and pair of morphisms $f : Q \to V$ and $f : Q \to W$ there is a morphism $h : Q \to V \times W$ such that $f = \pi_V \circ h$ and $g = \pi_W \circ h$.

2. Prove that an algebraic prevariety V is an algebraic variety if and only if the diagonal is closed in $V \times V$.
3. Prove that every open cover of an algebraic variety has a finite refinement.
4. Prove that if V is a subvariety of an open set in $\mathbb{C}^n$, W is a subvariety of an open set in $\mathbb{C}^m$ and $f : V \to W$ is algebraic (has coordinate functions which are regular) then f is a morphism of ringed spaces.
5. Prove that if V is an affine variety, $\{f_i\}$ a finite set of elements of $\mathcal{O}(V)$ which have no common zero on V and M is an $\mathcal{O}(V)$-module such that M_{f_i} is finitely generated for each i, then M is finitely generated.
6. Use a trascendental induction argument to prove that I is injective if and only if for every singly generated module M and every submodule $N \subset M$, any morphism $N \to I$ has an extension to M.
7. Prove Krull's Theorem: If A is a Noetherian ring, K an ideal of A, M a finitely generated A-module and $N \subset M$ a submodule, then for each positive integer n there is a positive integer m such that $K^m M \cap N \subset K^n N$.
8. A sheaf supported on a single point is called a skyscraper sheaf. Prove that every skyscraper sheaf is flabby.
9. Prove that the kernel, image and cokernel of a morphism between quasicoherent (coherent) sheaves on an algebraic variety is also quasi-coherent (coherent).

11. Dolbeault Cohomology

In sheaf theory, a *vanishing theorem* is a theorem which asserts that cohomology vanishes for some class of sheaves in some range of degrees p (usually $p > 0$). Vanishing theorems generally insure the existence of global solutions to certain locally solvable problems. Theorem 10.17, which says that cohomology vanishes in degree greater than zero for all quasi-coherent sheaves on an affine variety, is an example of a vanishing theorem. Its proof is fairly elementary. We shall find it much more difficult to prove the analogous vanishing theorem for sheaves of modules over a sheaf of rings of holomorphic functions. Ultimately, we want to define the class of coherent analytic sheaves and the class of Stein spaces and prove that the cohomology of a coherent analytic sheaf on a Stein space vanishes in positive degrees. Stein spaces are the holomorphic analogues of affine varieties. However, it is considerably more difficult than it was in the algebraic case just to define and develop the elementary properties of coherent analytic sheaves. This is, in part, due to the fact that, although the local ring $\mathcal{H}_\lambda$ is Noetherian, for every open set U the ring $\mathcal{H}(U)$ fails to be Noetherian. We will carry out this development in the next chapter.

In this chapter we prove the one vanishing theorem for holomorphic sheaves that can be proved without a great deal of work – Dolbeault's Theorem. This is a vanishing theorem for the structure sheaf ${}_n\mathcal{H}$ of $\mathbb{C}^n$ as a complex manifold and it involves constructing a resolution of ${}_n\mathcal{H}$ by a complex of sheaves of differential forms which is analogous to the deRham complex.

On a smooth ($\mathcal{C}^\infty$) manifold X, a (complex) differential p form is a smooth section of the vector bundle which assigns to each $x \in X$ the space of skew symmetric complex p-multilinear forms on the (complexified) tangent space of X at x. On an open set U in $\mathbb{R}^n$, where we may choose bases $\left\{\dfrac{\partial}{\partial x_i}\right\}$ for the tangent space and $\{dx_i\}$ for the cotangent space which correspond to a basis $\{x_i\}$ for $\mathbb{R}^n$, the typical differential p-form may be written as

$$\phi = \sum \phi_{i_1 \ldots i_p} dx_{i_1} \wedge \cdots \wedge dx_{i_p}$$

where $\phi_{i_1\ldots i_p} \in \mathcal{C}^\infty(U)$. We denote the $\mathcal{C}^\infty$-module of differential p-forms on U by $\mathcal{E}^p(U)$. The correspondence $U \to \mathcal{E}^p(U)$ is a sheaf, $\mathcal{E}^p$, of $\mathcal{C}^\infty$ modules. Exterior differentiation $d : \mathcal{E}^p \to \mathcal{E}^{p+1}$ is defined by

$$d(\phi_{i_1\ldots i_p} dx_{i_1} \wedge \cdots \wedge dx_{i_p}) = \sum_i \frac{\partial \phi_{i_1\ldots i_p}}{\partial x_i} dx_i \wedge dx_{i_1} \wedge \cdots \wedge dx_{i_p}$$

It follows from the Poincaré Lemma that $\mathcal{E} = \{\mathcal{E}^P, d^p\}$ is a complex which provides a fine resolution $\mathbb{C} \to \mathcal{E}$ of the constant sheaf $\mathbb{C}$.

If U is an open set in $\mathbb{C}^n$, then we consider it as an open subset of the real vector space $\mathbb{R}^{2n}$ with basis $\{x_1, \ldots, x_n, y_1, \ldots, y_n\}$. However, instead of the usual bases for the complexified tangent and cotangent spaces we use the basis consisting of

$$dz_i = dx_i + i\, dy_i, \quad d\bar{z}_i = dx_i - i\, dy_i, \quad i = 1, \ldots, n$$

for the complexified cotangent space and

$$\frac{\partial}{\partial z_i} = 1/2\left(\frac{\partial}{\partial x_i} - i\,\frac{\partial}{\partial y_i}\right), \quad \frac{\partial}{\partial \bar{z}_i} = 1/2\left(\frac{\partial}{\partial x_i} + i\,\frac{\partial}{\partial y_i}\right), \quad i = 1, \dots, n$$

for the complexified tangent space. Note that these are dual bases to one-another. In terms of the above basis for the cotangent space, a differential form in $\mathcal{E}^r(U)$ may be written as

$$\sum_{p+q=r} \phi_{j_i \dots j_p k_1 \dots k_q} dz_{j_1} \wedge \dots \wedge dz_{j_p} \wedge d\bar{z}_{k_1} \wedge \dots \wedge d\bar{z}_{k_q}$$

with coeficients $\phi_{j_i \dots j_p k_1 \dots k_q} \in \mathcal{C}^\infty(U)$. If we let $\mathcal{E}^{p,q}(U)$ denote the differential forms in $\mathcal{E}^{p+q}(U)$ which are of degree p in the dz_i and of degree q in the $d\bar{z}_i$, then we have a direct sum decomposition:

$$\mathcal{E}^r(U) = \sum_{p+q=r} \mathcal{E}^{p,q}(U).$$

Forms in the space $\mathcal{E}^{p,q}(U)$ are said to have *bidegree* (p, q) and *total degree* $p + q$.

When restricted to $\mathcal{E}^{p,q}$, exterior differentiation defines a map

$$d : \mathcal{E}^{p,q} \to \mathcal{E}^{p+1,q} + \mathcal{E}^{p,q+1}.$$

If we define ∂ and $\bar{\partial}$ to be the operators which, on forms of bidegree (p, q), act as d followed by projection on $\mathcal{E}^{p+1,q}$ and $\mathcal{E}^{p,q+1}$, respectively, then

$$\partial : \mathcal{E}^{p,q} \to \mathcal{E}^{p+1,q}, \quad \bar{\partial} : \mathcal{E}^{p,q} \to \mathcal{E}^{p,q+1}$$

and

$$d = \partial + \bar{\partial}.$$

Note that $0 = d^2 = \partial^2 + \partial\bar{\partial} + \bar{\partial}\partial + \bar{\partial}^2$. After sorting out terms of different bidegree, it follows that

$$\partial^2 = 0, \quad \partial\bar{\partial} + \bar{\partial}\partial = 0, \quad \bar{\partial}^2 = 0.$$

Also, since

$$d(\phi \wedge \psi) = d\phi \wedge \psi + (-1)^r \phi \wedge d\psi,$$

it follows that

$$\partial(\phi \wedge \psi) = \partial\phi \wedge \psi + (-1)^r \phi \wedge \partial\psi, \quad \bar{\partial}(\phi \wedge \psi) = \bar{\partial}\phi \wedge \psi + (-1)^r \phi \wedge \bar{\partial}\psi$$

for $\phi \in \mathcal{E}^r(U)$ and ψ an arbitrary form.

The operator $\bar{\partial}$ is given explicitly by

$$\begin{aligned} &\bar{\partial}(\phi_{j_i \dots j_p k_1 \dots k_q} dz_{j_1} \wedge \dots \wedge dz_{j_p} \wedge d\bar{z}_{k_1} \wedge \dots \wedge d\bar{z}_{k_q}) \\ &= \sum_i \frac{\partial \phi_{j_i \dots j_p k_1 \dots k_q}}{\partial \bar{z}_i} d\bar{z}_i \wedge dz_{j_1} \wedge \dots \wedge dz_{j_p} \wedge d\bar{z}_{k_1} \wedge \dots \wedge d\bar{z}_{k_q}. \end{aligned}$$

In particular, for $f \in \mathcal{E}^{0,0}(U) = \mathcal{C}^\infty(U)$,

$$\bar{\partial} f = \sum_i \frac{\partial f}{\partial \bar{z}_i} d\bar{z}_i$$

and so the Cauchy-Riemann equations assert that $\bar{\partial} f = 0$ if and only if f is holomorphic in U.

11.1 Definition. *For a domain U in $\mathbb{C}^n$ the p^{th} Dolbeault complex is the complex*

$$0 \longrightarrow \mathcal{E}^{p,0}(U) \xrightarrow{\bar\partial} \mathcal{E}^{p,1}(U) \xrightarrow{\bar\partial} \cdots \xrightarrow{\bar\partial} \mathcal{E}^{p,n}(U) \longrightarrow 0$$

Its q th cohomology group is called the (p,q)-Dolbeault cohomology of U and is denoted $\mathcal{H}^{p,q}(U)$.

Note that the $(p,0)$-Dolbeault cohomology is just the space of p-forms in $dz_1, \dots, dz_n$ with holomorphic coeficients. This is called the space of *holomorphic p-forms* and is denoted $\mathcal{H}^p(U)$. The correspondence $U \to \mathcal{H}^p(U)$ is clearly a sheaf and will be denoted $\mathcal{H}^p$. Note that $\mathcal{H}^0$ is the sheaf of holomorphic functions $\mathcal{H}$.

Our main objective in this chapter is to prove that Dolbeault cohomology vanishes for $q > 0$ if U is a polydisc. Among other things, this will imply that for each p, as a complex of sheaves, the Dolbeault complex is exact except in degree zero and provides a fine resolution of the sheaf of holomorphic p-forms. In particular, when $p = 0$ the Dolbeault complex provides a fine resolution of the sheaf of holomorphic functions and, hence, can be used to compute its sheaf cohomology.

The first step is to prove the generalized Cauchy integral theorem:

11.2 Theorem. *Let U be an open subset of $\mathbb{C}$ bounded by a simple closed rectifiable curve γ. If f is a $\mathcal{C}^\infty$ function in a neighborhood of $\bar U$ and $z \in U$, then*

$$f(z) = \frac{1}{2\pi i}\int_\gamma f(\zeta)\frac{d\zeta}{\zeta - z} + \frac{1}{2\pi i}\iint\limits_U \frac{\partial f(\zeta)}{\partial\bar\zeta}\frac{d\zeta\wedge d\bar\zeta}{\zeta - z}.$$

Proof. Note that

$$d\left(f(\zeta)\frac{d\zeta}{\zeta - z}\right) = \frac{\partial}{\partial\bar\zeta}\left(\frac{f(\zeta)}{\zeta - z}\right) d\bar\zeta\wedge d\zeta = \frac{\partial f(\zeta)}{\partial\bar\zeta}\frac{d\bar\zeta\wedge d\zeta}{\zeta - z}$$

Thus, if γ_r is the boundary of the disc $D(z,r)$ and if r is chosen small enough that this disc is contained in U, then Stokes' Theorem implies that

$$\iint\limits_{U_r} \frac{\partial f(\zeta)}{\partial\bar\zeta}\frac{d\bar\zeta\wedge d\zeta}{\zeta - z} = \int_\gamma f(\zeta)\frac{d\zeta}{\zeta - z} - \int_{\gamma_r} f(\zeta)\frac{d\zeta}{\zeta - z}$$

where $U_r = U - D(z,r)$. Now $(\zeta - z)^{-1}$ is integrable on any bounded region of the plane, as is easily seen by integrating its absolute value using polar coordinates centered at z. Thus,

$$\lim_{r\to 0}\iint\limits_{U_r} \frac{\partial f(\zeta)}{\partial\bar\zeta}\frac{d\bar\zeta\wedge d\zeta}{\zeta - z} = \iint\limits_U \frac{\partial f(\zeta)}{\partial\bar\zeta}\frac{d\bar\zeta\wedge d\zeta}{\zeta - z}.$$

Also,

$$\lim_{r\to 0}\int_{\gamma_r} f(\zeta)\frac{d\zeta}{\zeta - z} = \lim_{r\to 0}\int_0^{2\pi} f(z + re^{it})\, idt = 2\pi i f(z)$$

The result follows.

11.3 Theorem. *If $f \in \mathcal{C}^\infty(U)$ for an open set $U \subset \mathbb{C}$ containing a compact set K then there exists a neighborhood V of K with $V \subset U$ and a $g \in \mathcal{C}^\infty(V)$ such that $\partial g/\partial \bar{z} = f$ in V.*

Proof. We modify f so that it is actually $\mathcal{C}^\infty$ on all of $\mathbb{C}$ with compact support in U by multiplying it by a $\mathcal{C}^\infty$ function which is one in a neighborhood V of K and has compact support in U and then extending the resulting function to be zero on the complement of U. Then the integral

$$g(z) = \iint\limits_{\mathbb{C}} f(\zeta)\frac{d\zeta \wedge d\bar{\zeta}}{\zeta - z}$$

is defined for all $z \in \mathbb{C}$ and defines a function $g \in \mathcal{C}^\infty(\mathbb{C})$. We calulate the derivative $\partial g/\partial \bar{z}$ of g using the change of variables $\zeta \to \zeta + z$:

$$\begin{aligned}\frac{\partial}{\partial \bar{z}} g(z) &= \frac{1}{2\pi i}\frac{\partial}{\partial \bar{z}}\iint f(\zeta + z)\frac{d\zeta \wedge d\bar{\zeta}}{\zeta} \\ &= \frac{1}{2\pi i}\iint \frac{\partial f(\zeta + z)}{\partial \bar{z}}\frac{d\zeta \wedge d\bar{\zeta}}{\zeta} = \frac{1}{2\pi i}\iint \frac{\partial f(\zeta + z)}{\partial \bar{\zeta}}\frac{d\zeta \wedge d\bar{\zeta}}{\zeta} \\ &= \frac{1}{2\pi i}\iint \frac{\partial f(\zeta)}{\partial \bar{\zeta}}\frac{d\zeta \wedge d\bar{\zeta}}{\zeta - z} = f(z)\end{aligned}$$

where the last line follows from reversing the change of variables and using the generalized Cauchy integral theorem on U (recall that f vanishes on the complement of a compact subset of U and so the line integral in Theorem 10.2 vanishes). Thus, $\partial g/\partial \bar{z} = f$ on all of $\mathbb{C}$. Of course, we modified f on the complement of V and so this equation holds for our original f only on V, but this is what was to be shown.

If $K \subset \mathbb{C}^n$ is compact, we denote by $\mathcal{E}^{p,q}(K)$ the space $\Gamma(K; \mathcal{E}^{p,q})$ of $\mathcal{C}^\infty$ forms of bidegree (p, q) defined in a neighborhood of K. The Dolbeault cohomology $\mathcal{H}^{p\cdot q}(K)$ for K is then the cohomology of the complex $\{\mathcal{E}^{p,q}(K), \bar{\partial}\}$.

11.4 Dolbeault's Lemma. *If $\bar{\Delta}$ is a compact polydisc in $\mathbb{C}^n$ then $\mathcal{H}^{p,q}(\bar{\Delta}) = 0$ for $q > 0$ and for each p.*

Proof. Suppose $q > 0$ and let $\phi \in \mathcal{E}^{p,q}(\bar{\Delta})$ be a form such that $\bar{\partial}\phi = 0$. We must show that $\phi = \bar{\partial}\psi$ for some form $\psi \in \mathcal{E}^{p,q-1}(\bar{\Delta})$. Let k be the least integer such that the expression for ϕ involves no conjugate differential $d\bar{z}_j$ with $j > k$; that is, ϕ can be written in terms of the conjugate differentials $d\bar{z}_1, \cdots, d\bar{z}_k$ and the differentials $dz_1, \dots, dz_n$. We proceed by induction on k. If $k = 0$ then $\phi = 0$ since $q > 0$ and, hence, there is nothing to prove. Thus, we assume that $k > 0$ and the result is true for integers less than k. We write ϕ as

$$\phi = d\bar{z}_k \wedge \alpha + \beta$$

where α and β involve only the conjugate differentials $d\bar{z}_1, \cdots d\bar{z}_{k-1}$. Then

$$0 = \bar{\partial}\phi = -d\bar{z}_k \wedge \bar{\partial}\alpha + \bar{\partial}\beta$$

If $j > k$ then no cancellation can occur between terms of $\bar{\partial}\alpha \wedge d\bar{z}_k$ involving $d\bar{z}_j$ and terms of $\bar{\partial}\beta$ involving $d\bar{z}_j$. It follows that such terms individually vanish and, hence, the coeficients of α and β are holomorphic in the variables $z_{k+1}, \cdots, z_n$. Now it follows from Theorem 11.3 that if f is a coeficient of α then $f = \partial g/\partial \bar{z}_k$ for some $g \in \mathcal{C}^\infty(\bar{\Delta})$ which is also holomorphic in the variables $z_{k+1}, \cdots, z_n$. That the solution g given by Theorem 10.3 is actually $\mathcal{C}^\infty$ in all the variables and not just in ζ_k and that it is holomorphic in the variables $z_{k+1}, \cdots, z_n$ follows from the fact that these things are true of f and the solution g is given in terms of f by an explicit integral formula which commutes with the differential operators in question. By replacing each coeficient f of α by the corresponding g as above, we obtain a $(p, q-1)$-form γ with the property that

$$\bar{\partial}\gamma = \delta + d\bar{z}_k \wedge \alpha$$

where δ is a form involving only the conjugate differentials $d\bar{z}_1, \cdots, d\bar{z}_{k-1}$. Then $\phi - \bar{\partial}\gamma = \beta - \delta$ involves only the conjugate differentials $d\bar{z}_1, \cdots, d\bar{z}_{k-1}$ and is $\bar{\partial}$ closed. That is, $\bar{\partial}(\phi - \bar{\partial}\gamma) = 0$. Thus, by the induction hypothesis, we conclude that $\bar{\partial}\eta = \phi - \bar{\partial}\gamma$ for some $\eta \in \mathcal{E}^{p,q-1}(\bar{\Delta})$, from which it follows that $\phi = \bar{\partial}\psi$ with $\psi = \eta + \gamma \in \mathcal{E}^{p,q-1}(\bar{\Delta})$. This completes the proof.

Note that it was not really important in the above argument that $\bar{\Delta}$ be a polydisc. It was important that it be a Cartesian product – that is, a set of the form $K_1 \times K_2 \times \cdots \times K_n$ for some collection of compact sets $K_i \subset \mathbb{C}$. This is due to the fact that the solution was obtained by applying Theorem 11.3 in each variable separately while treating the other variables as parameters.

The following theorem concerns an open polydisc $\Delta = \Delta(\lambda, r)$. Note that we allow some or all of the radii r_i to be infinite, Thus, $\mathbb{C}^n$ itself is an open polydisc.

11.5 Theorem. *Let Δ be an open polydisc. Then $\mathcal{H}^{p,q}(\Delta) = 0$ for $q > 0$ and for all p.*

Proof. Let Δ_j be a sequence of open polydiscs with compact closure such that $\bar{\Delta}_j \subset \Delta_{j+1}$ and $\bigcup_j \Delta_j = \Delta$. If $\phi \in \mathcal{E}^{p,q}(\Delta)$ and $\bar{\partial}\phi = 0$ then we will construct $\psi \in \mathcal{E}^{p,q-1}(\Delta)$ such that $\bar{\partial}\psi = \phi$ inductively using Dolbeault's Lemma on the sets $\bar{\Delta}_j$. We first take care of the case $q > 1$ which is different and considerably easier than the case $q = 1$.

If $q > 1$ we inductively construct a sequence of forms $\{\psi_j\}$ such that $\psi_j \in \mathcal{E}^{p,q-1}(\bar{\Delta}_i)$, $\bar{\partial}\psi_j = \phi$ on Δ_j and $\psi_{j+1}|_{\Delta_j} = \psi_j$. Clearly this will give the desired result, since we can then define a solution $\psi \in \mathcal{E}^{p,q-1}(\Delta)$ by $\psi|_{\Delta_j} = \psi_j$. Suppose the sequence $\{\psi_j\}$ has been constructed with the above properties for $j < k$. Then we use Dolbeault's Lemma to find $\theta \in \mathcal{E}^{p,q-1}(\bar{\Delta}_k)$ such that $\bar{\partial}\theta = \phi$ in a neighborhood of $\bar{\Delta}_k$. We then have

$$\bar{\partial}(\theta - \psi_{k-1}) = 0$$

in a neighborhood of $\bar{\Delta}_{k-1}$ and, since $q > 1$, we may apply Dolbeault's Lemma again to find $\eta \in \mathcal{E}^{p,q-2}(\bar{\Delta}_{k-1})$ such that $\bar{\partial}\eta = \theta - \psi_{k-1}$ in a neighborhood of $\bar{\Delta}_{k-1}$. By multiplying by a $\mathcal{C}^\infty$ function which is one in a neighborhood of $\bar{\Delta}_{k-1}$ and has compact support in an appropriate slightly larger neighborhood of $\bar{\Delta}_{k-1}$, we may assume that η is actually in $\mathcal{E}^{p,q-2}(\bar{\Delta}_k)$. Then $\psi_k = \theta - \bar{\partial}\eta \in \mathcal{E}^{p,q-1}(\bar{\Delta}_k)$ gives the required next function in our sequence since

$$\bar{\partial}\psi_k = \bar{\partial}(\theta - \bar{\partial}\eta) = \phi$$

on Δ_k and $\psi_k = \psi_{k-1}$ on Δ_{k-1}. This completes the proof in the case $q > 1$.

In the case $q = 1$ we use the sequence $\{\Delta_j\}$ as before, but this time we inductively construct a sequence $\psi_j \in \mathcal{E}^{p,q-1}(\bar{\Delta}_j)$ such that $\bar{\partial}\psi_j = \phi$ on $\bar{\Delta}_j$ and $||\psi_j - \psi_{j+1}||_j < 2^{-j}$ for each j. Here, $||\theta||_j$ is the sum of the supremum norms on $\bar{\Delta}_j$ of the coeficients of the form θ. Suppose such a sequence $\{\psi_j\}$ has been constructed for all indices $j < k$. We use Dolbeault's Lemma to find $\theta \in \mathcal{E}^{p,0}(\bar{\Delta}_k)$ such that $\bar{\partial}\theta = \phi$ in a neighborhood of $\bar{\Delta}_k$. As before,

$$\bar{\partial}(\theta - \psi_{k-1}) = 0$$

in a neighborhood of $\bar{\Delta}_{k-1}$. This means that $\theta - \psi_{k-1}$ has coeficients which are holomorphic in a neighborhood of $\bar{\Delta}_{k-1}$. If we represent these coeficients as convergent power series about the point which is the center of $\bar{\Delta}_{k-1}$, then it is clear that we may choose a form η, with polynomial coeficients, such that $||\theta - \psi_{k-1} - \eta||_{k-1} < 2^{-k+1}$. Then $\psi_k = \theta - \eta$ has the properties that $\bar{\partial}\psi_k = \phi$ in a neighborhood of Δ_k and $||\psi_k - \psi_{k-1}||_{k-1} < 2^{-k+1}$. Thus, by induction, we may construct the sequence $\{\psi_j\}$ as claimed. Now on a given Δ_k consider the sequence $\{(\psi_j)|_{\Delta_k}\}_{j=k}^{\infty}$. This is a Cauchy sequence in the norm $||\cdot||_k$ since $||\psi_{j+1} - \psi_j||_k < ||\psi_{j+1} - \psi_j||_j < 2^{-j}$ for $j \geq k$. Furthermore, the terms of this sequence differ from the first term by forms with holomorphic coeficients (since the difference is killed by $\bar{\partial}$. Thus, the sequence may be regarded as a fixed form plus a uniformly convergent sequence of forms with holomorphic coeficients. It follows that this sequence actually converges in the topology of $\mathcal{E}^{p,q-1}(\Delta_k)$ – the topology in which all derivatives of all coeficients converge uniformly on compact subsets of Δ_k. Now since this is true for each k, the limit determines a form $\psi \in \mathcal{E}^{p,q-1}(\Delta)$ which clearly satifies $\bar{\partial}\psi = \phi$. This completes the proof.

Note that, as with the previous Lemma, it was not really important in the above argument that Δ be a polydisc. It was important that it be a Cartesian product – that is, a set of the form $U_1 \times U_2 \times \cdots \times U_n$ for some collection of open sets $U_i \subset \mathbb{C}$.

Clearly, the Dolbeault complex and Dolbeault cohomology can be defined on any complex manifold, since the division of $\mathcal{E}^r$ into bigraded terms $\mathcal{E}^{p,q}$ is independent of the choice of complex coordinate system (Problem 11.1).

11.6 Corollary. *If X is any complex manifold, then for each p there is a natural isomorphism $\mathcal{H}^{p,q}(X) \to H^q(X;\mathcal{H}^p)$ between Dolbeault cohomology on X and sheaf cohomology of $\mathcal{H}^p$ on X. In particular, $\mathcal{H}^{0,q}(X)$ is isomorphic to the sheaf cohomology $H^q(X;\mathcal{H})$ of the sheaf of holomorphic functions.*

Proof. Since a complex manifold has a neighborhood base consisting of sets which are biholomorphic to open polydiscs, it follows from Theorem 11.5 that the complex

$$0 \to \mathcal{H}^p \to \mathcal{E}^{p,0} \to \mathcal{E}^{p,1} \to \cdots \to \mathcal{E}^{p,n} \to 0$$

is exact. Since each $\mathcal{E}^{p,q}$ is a sheaf of $\mathcal{C}^\infty$-modules, this sequence is a fine resolution of $\mathcal{H}^P$ and, hence, may be used to compute its sheaf cohomology.

11.7 Corollary. *If $\Delta \subset \mathbb{C}^n$ is an open polydisc, then for each p the sheaf cohomology $H^q(\Delta;\mathcal{H}^p)$ vanishes for $q > 0$.*

11.8 Corollary. *If X is a complex manifold of dimension n, then for each p the sheaf cohomology $H^q(X;\mathcal{H}^p)$ vanishes for $q > n$.*

11. Problems

1. Give a coordinate free definition of the space $\mathcal{E}^{p,q}$.
2. Prove that if Δ is an open polydisc and $f \in \mathcal{H}(\Delta)$ then a function $g \in \mathcal{H}(\Delta)$ belongs to the ideal generated by f in $\mathcal{H}(\Delta)$ if and only if its germ at λ belongs to the ideal generated by the germ of f at λ in $\mathcal{H}_\lambda$ for each $\lambda \in f^{-1}(\{0\})$.
3. Prove Hartog's extension theorem: *If K is a compact subset of an open set $U \subset \mathbb{C}^n$, $U - K$ is connected and $n > 1$, then each function f which is holomorphic in $U - K$ extends to be holomorphic in U.* Hint: Let K' be chosen so that K' is a compact set containing K in its interior, $K' \subset U$ and $U - K'$ is connected. Then multiply f by a $\mathcal{C}^\infty$ function which is zero in a neighborhood of K and one in a neighborhood of $U - K'$. The resulting function then extends to a $\mathcal{C}^\infty$ function g in U which agrees with f on the connected set $U - K'$. Now use Corollary 11.7 and Problem 1.5 to show that you can find a $\mathcal{C}^\infty$ function h on U which vanishes on $U - K'$ and is such that $g - h$ is holomorphic in all of U.
4. Prove that $H^q(U, \mathcal{H}^p) = 0$ for $q > 0$ and for any open subset $U \subset \mathbb{C}$. Hint: use Theorem 11.3 and an approximation argument like that used in the second half of the proof of Theorem 11.5.
5. Prove that $H^q(P^1(\mathbb{C}), \mathcal{H}^p) = \mathbb{C} \cdot \delta_{p,q}$, where $P^1(\mathbb{C})$ is the Riemann sphere. Hint: Cover the sphere with two open disks which overlap in an annulus. By the previous problem, this is a Leray cover for the sheaves $\mathcal{H}^p$. Now compute Čech cohomology for each of the sheaves $\mathcal{H}^p$ and this cover.

12. Coherent Analytic Sheaves

A coherent algebraic sheaf $\mathcal{S}$ on an algebraic variety X is a sheaf of ${}_X\mathcal{O}$ modules which, in an affine neighborhood V of each point, can be realized as the localization $\tilde{M}$ of a finitely generated $\mathcal{O}(V)$ module M. Since M is finitely generated, there is a surjection $\mathcal{O}(V)^n \to M$ and, since $\mathcal{O}(V)$ is Noetherian, the kernel of this map is finitely generated also. Thus, there is an exact sequence

$$\mathcal{O}(V)^m \to \mathcal{O}(V)^n \to M \to 0$$

If we localize this sequence over V we obtain an exact sequence of sheaves

$${}_V\mathcal{O}^m \to {}_V\mathcal{O}^n \to \mathcal{S}|_V \to 0$$

Thus, a coherent algebraic sheaf is locally the cokernel of a morphism ${}_V\mathcal{O}^m \to {}_V\mathcal{O}^n$ between free finite rank sheaves of ${}_X\mathcal{O}$-modules. The converse is also true by Problem 10.9. Thus, this condition characterizes coherent algebraic sheaves. We will use the analogous condition as the definition of coherent analytic sheaves:

12.1 Definition. *Let X be a holomorphic variety. An analytic sheaf on X is a sheaf of ${}_X\mathcal{H}$-modules and a morphism of analytic modules is a morphism of sheaves of ${}_X\mathcal{H}$-modules. An analytic sheaf $\mathcal{S}$ will be called a coherent analytic sheaf if each point of X is contained in a neighborhood V such that $\mathcal{S}|_V$ is the cokernel of a morphism ${}_V\mathcal{H}^m \to {}_V\mathcal{H}^n$ between free finite rank analytic sheaves.*

There are two reasons why a definition in terms of localization, like that of chapter 10, would not work well: $\mathcal{H}(V)$ is never Noetherian and, tensor product relative to $\mathcal{H}(V)$ is not a well behaved operation.

The purpose of this chapter is to show that the sheaves of greatest interest in holomorphic function theory are, in fact, coherent analytic sheaves and to show that the category of coherent analytic sheaves on X is an abelian category. Later chapters will deal with vanishing theorems for the cohomology of coherent analytic sheaves and applications to holomorphic function theory.

If X is a holomorphic variety, note that each morphism of analytic sheaves $\phi : {}_X\mathcal{H} \to {}_X\mathcal{H}$ is given by multiplication by a global holomorphic function f. In fact, f is just $\phi(1)$ – the image of the identity section of ${}_X\mathcal{H}$. Similarly, each morphism of analytic sheaves $\phi : {}_X\mathcal{H}^k \to {}_X\mathcal{H}^m$ is given by a $m \times k$ matrix with entries which are holomorphic functions on X.

12.2 definition. *An analytic sheaf $\mathcal{S}$ is said to be locally finitely generated if for each point $x \in X$ there is a neighborhood U of x and finitely many sections $s_1, \dots, s_k$ of $\mathcal{S}$ such that the germs of these sections at y generate $\mathcal{S}_y$ for every $y \in U$.*

In other words $\mathcal{S}$ is locally finitely generated if for each $x \in X$ there is a neighborhood U of x and a surjective morphism ${}_U\mathcal{H}^k \to \mathcal{S}|_U$ of analytic sheaves. Thus, $\mathcal{S}$ is coherent if it is locally finitely generated in such a way that the resulting morphisms ${}_U\mathcal{H}^k \to \mathcal{S}|_U$ have kernels which are also locally finitely generated. It will take some work to show that this is a reasonable condition. The key result is Oka's Theorem, which is a substitute for the Noetherian property:

12.3 Oka's Theorem. *Let $U \subset \mathbb{C}^n$ be an open set. Then the kernel of any ${}_U\mathcal{H}$-module morphism ${}_U\mathcal{H}^m \to {}_U\mathcal{H}^k$ is locally finitely generated.*

Proof. We prove this by induction on n. It is trivial for $n = 0$, since $\mathbb{C}^0$ is a point and ${}_U\mathcal{H}^m \to {}_U\mathcal{H}^k$ is a linear map between finite dimensional vector spaces over $\mathbb{C}$ in this case. Thus, we assume that $n > 0$ and the theorem is true in dimension less than n.

To prove the theorem is true for dimension n, we first show that we may reduce the proof to the case where $k = 1$. We do this using induction on k. Thus, suppose $k > 1$ and the theorem is true in dimension n for all morphisms ${}_U\mathcal{H}^m \to {}_U\mathcal{H}^j$ with $j < k$. Consider a morphism $\alpha : {}_U\mathcal{H}^m \to {}_U\mathcal{H}^k$. We may write $\alpha = (\beta, \gamma)$ where $\beta : {}_U\mathcal{H}^m \to {}_U\mathcal{H}^{k-1}$ is α followed by the projection $(f_1, \dots, f_k) \to (f_1, \dots, f_{k-1})$ and γ is α followed by the projection $(f_1, \dots, f_k) \to f_k$. Now by the induction hypothesis, $\ker \beta$ is locally finitely generated. Hence, if $x \in U$ then there is a neighborhood V of x, contained in U, a $p > 0$ and a morphism $\phi : {}_V\mathcal{H}^p \to {}_V\mathcal{H}^m$ which maps onto $\ker \beta$. Thus, we have a diagram

$$\begin{array}{ccccc} {}_V\mathcal{H}^p & \xrightarrow{\phi} & {}_V\mathcal{H}^m & \xrightarrow{\beta} & {}_V\mathcal{H}^{k-1} \\ & & {\scriptstyle\gamma}\big\downarrow & & \\ & & {}_V\mathcal{H} & & \end{array}$$

with the top row exact. We also have, by assumption, that the kernel of $\gamma \circ \phi : {}_V\mathcal{H}^p \to {}_V\mathcal{H}$ is locally finitely generated. This means that, after shrinking V if necessary, we can find $q > 0$ and a morphism $\psi : {}_V\mathcal{H}^q \to {}_V\mathcal{H}^p$ which maps onto $\ker \gamma \circ \phi$. But $\ker \alpha = \ker \beta \cap \ker \gamma = \phi(\ker \gamma \circ \phi) = \operatorname{im} \phi \circ \psi$. Thus, the kernel of α is also locally finitely generated as was to be shown.

Thus, we have reduced the proof to showing, in dimension n, that the kernel of a morphism of the form $\alpha : {}_U\mathcal{H}^m \to {}_U\mathcal{H}$ is locally finitely generated, under the assumption that the theorem holds for all m and k in dimensions less than n. The strategy of the proof is to use the Weierstrass theorems to reduce the problem to an analogous one involving polynomials of a fixed degree in z_n and then to apply the induction assumption to the coeficients of these polynomials.

Given a point $x \in U$ we must show that $\ker \alpha$ is finitely generated in some neighborhood of x. Without loss of generality, we may assume that the point x is the origin. The map α has the form $\alpha(g_1, \dots, g_m) = \sum f_i g_i$ for an m-tuple of functions $f_i \in {}_n\mathcal{H}(U)$. By appropriate choice of coordinates, we may assume that the germ at 0 of each f_i is regular of some degree at 0 and, hence, by the Weierstrass preparation theorem, has the form $u_i p_i$ where p_i is a Weierstrass polynomial and u_i is a unit. We may replace these germs by their representatives in some neighborhood V of 0 and by choosing V small enough we may assume the u_i are non vanishing in V. Then the map $(g_1, \dots, g_m) \to (u_1 g_1, \dots, u_m g_m)$ is an automorphism of ${}_V\mathcal{H}^m$ which maps the kernel of α to the kernel of the morphism determined by the m-tuple $(p_1, \dots, p_m)$. Thus, without loss of generality, we may replace the f_i's with the p_i's and assume that α has the form $\alpha(g_1, \dots, g_m) = \sum p_i g_i$. Let d be the maximum of the degrees of the p_i's. We may assume that V has the form $V' \times V''$ for open sets $V' \subset \mathbb{C}^{n-1}$ and $V'' \subset \mathbb{C}$. Let $\mathcal{K}_d$ denote the sheaf on V' defined as follows

$$\mathcal{K}_d(W) \subset \ker \alpha : {}_V\mathcal{H}^m(W \times V'') \to {}_V\mathcal{H}(W \times V'')$$

is the subspace consisting of m-tuples whose entries are polynomials of degree less than or equal to d in the variable z_n with coeficients in ${}_{n-1}\mathcal{H}(W)$. This is a sheaf of ${}_{V'}\mathcal{H}$-modules. We shall show that it is locally finitely generated on V'. Now for each neighborhood $W \subset V'$, the space of m-tuples $(q_1, \dots, q_m)$ where each q_i is a polynomial in z_n of degree at most d, with coeficients which are in ${}_{n-1}\mathcal{H}(W)$ is a free module of rank $m(d+1)$ over ${}_{n-1}\mathcal{H}(W)$. The map, α determines a ${}_{n-1}\mathcal{H}(W)$-module morphism of this free module into the free ${}_{n-1}\mathcal{H}(W)$-module of rank $2d+1$ consisting of polynomials of degree at most $2d$ in z_n with coeficients in ${}_{n-1}\mathcal{H}(W)$. Furthermore, $\mathcal{K}_d(W)$ is the kernel of this morphism. In other words, we may regard α as determining a morphism of sheaves of ${}_{V'}\mathcal{H}$-modules, ${}_{V'}\mathcal{H}^{m(d+1)} \to {}_{V'}\mathcal{H}^{2d+1}$ and our sheaf $\mathcal{K}_d$ is its kernel. By the induction hypothesis, $\mathcal{K}_d$ is locally finitely generated as a sheaf of ${}_{V'}\mathcal{H}$-modules. To complete the proof, we will show that $\mathcal{K}_d$ generates $\ker \alpha$ as a sheaf of ${}_V\mathcal{H}$-modules.

Let λ be a point of $W \times V''$. We must show that the stalk of K_d at λ generates the stalk of $\ker \alpha$ at λ. By performing a translation, we may assume that the point λ is the origin. In the process, however, we lose the fact that the polynomials p_i are Weierstrass polynomials at 0. They are, however, still polynomials in z_n and, by the Weierstrass preparation theorem, we may factor p_1 as $p_1 = p_1' p_1''$ where the germ at 0 of p_1' is a Weierstrass polynomial and the germ at 0 of p_1'' is a unit. We set

$$d_1' = \deg p_1', \quad d_1'' = \deg p_1''$$

and note that

$$d_1' + d_1'' = \deg p_1 < d.$$

If $h = (h_1, \dots, h_m) \in {}_n\mathcal{H}_0^m$ we then use the Weierstrass division theorem on each component of h to write $h = p_1' h'' + r'$, where $r', h'' \in {}_n\mathcal{H}_0^m$ and r' has components which are polynomials in z_n of degree less than d_1'. If we set $h' = (p_1'')^{-1} h$ then

$$h = p_1 h' + r'$$

Also,

$$p_1 h' = q + \sum_{j=2}^m h_j' e_j$$

where

$$q = (\sum p_i h_i', 0, \dots, 0) \quad \text{and} \quad e_j = (-p_j, 0, \dots, p_1, 0. \dots, 0)$$

with p_1 occuring in the j th position of e_j. Note that each e_j belongs to $(\mathcal{K}_d)_0$. Thus, if we set $r = r' + q$, we have

$$h = r + \sum_{j=2}^m h_j' e_j$$

where $e_j \in (\mathcal{K}_d)_0$ and $r = (r_1, \dots, r_m)$ an element of ${}_n\mathcal{H}_0$ with $r_2, \dots, r_m$ polynomials in z_n of degree less than d_1'. Now suppose h belongs to $\ker \alpha$. Then so does each e_j and, hence, so does r. This means that

$$p_1 r_1 = -(p_2 r_2 + \cdots + p_m r_m)$$

which implies that $p_1'p_1''r_1 = p_1r_1$ is a polynomial in z_n of degree less than $d+d_1'$ since this is true of each term on the right above. However, since p_1' is a Weierstrass polynomial, the Weierstrass division theorem implies that $p_1''r_1$ must be a polynomial of degree at most d. Then the entries of $p_1''r$ all have degree at most d and we conclude that $p_1''r \in (\mathcal{K}_d)_0$. Since p_1'' is a unit we conclude that r and, thus, h belongs to the submodule of ${}_n\mathcal{H}_0$ generated by $\mathcal{K}_d$. Thus, we have shown that at every point of V the stalk of $\mathcal{K}_d$ generates the stalk of the kernel of α. Since, $\mathcal{K}_d$ is itself locally finitely generated over ${}_{V'}\mathcal{O}$, the proof is complete.

12.4 Corollary. *If U is an open set in $\mathbb{C}^n$ and $\mathcal{M}$ is a locally finitely generated sheaf of submodules of ${}_U\mathcal{H}^k$, then $\mathcal{M}$ is coherent. In particular, if $\phi : {}_U\mathcal{H}^m \to {}_U\mathcal{H}^k$ is a morphism of analytic sheaves , then $\ker\phi$ and $\operatorname{im}\phi$ are coherent analytic sheaves on U.*

Proof. Since $\mathcal{M}$ is locally finitely generated, for each point of U there is a neighborhood V on which $\mathcal{M}$ is the image of a morphism $\theta : {}_V\mathcal{H}^m \to {}_V\mathcal{H}^k$ of sheaves of ${}_V\mathcal{H}$-molules. By Oka's Theorem we know that $\ker\theta$ is locally finitely generated. Thus, by shrinking V if necessary, we may assume $\ker\theta$ is the image of a morphism $\psi : {}_V\mathcal{H}^p \to {}_V\mathcal{H}^m$. Then, on V, $\mathcal{M}$ is the cokernel of ψ. Thus, $\mathcal{M}$ is coherent.

Since the image of a morphism $\phi : {}_U\mathcal{H}^m \to {}_U\mathcal{H}^k$ of analytic sheaves is finitely generated by definition and the kernel is locally finitely generated by Oka's Theorem, they are both coherent by the result of the previous paragraph.

12.5 Corollary. *If U is an open set in $\mathbb{C}^n$ and $\mathcal{M}$ and $\mathcal{N}$ are coherent sheaves of submodules of ${}_U\mathcal{H}^m$, then so is $\mathcal{M}\cap\mathcal{N}$.*

Proof. Every point of U has a neighborhood V on which there are morphisms $\phi : {}_V\mathcal{H}^p \to {}_V\mathcal{H}^m$ with image $\mathcal{M}|_V$ and $\psi : {}_V\mathcal{H}^q \to {}_V\mathcal{H}^m$ with image $\mathcal{N}|_V$. Consider the map $\theta : {}_V\mathcal{H}^{p+q} \to {}_V\mathcal{H}^m$ defined by writing ${}_V\mathcal{H}^{p+q}$ as ${}_V\mathcal{H}^p \oplus {}_V\mathcal{H}^q$ and setting $\theta(f\oplus g) = \phi(f) - \psi(g)$. The kernel of θ is coherent by Oka's Theorem and, hence, is locally finitely generated. Furthermore, on V, $\mathcal{M}\cap\mathcal{N}$ is the image of the kernel of θ under ϕ. Thus, after shrinking V if neccessary, we may choose a finite set of generators for $\ker\theta$. Then the image of this set under ϕ will generate $\mathcal{M}\cap\mathcal{N}$ on V. Thus, $\mathcal{M}\cap\mathcal{N}$ is locally finitely generated. In view of the previous corollary, this proves that it is coherent.

If $U \subset \mathbb{C}^n$ is an open set and $\mathcal{I}$ and $\mathcal{J}$ are sheaves of ideals of ${}_U\mathcal{H}$ then $\mathcal{I} : \mathcal{J}$ will denote the sheaf which assigns to the open set $V \subset U$ the ideal $\mathcal{I}(V) : \mathcal{J}(V) = \{f \in {}_n\mathcal{H}(V) : f\mathcal{J}(V) \subset \mathcal{I}(V)\}$.

12.6 Corollary. *If U is an open set in $\mathbb{C}^n$ and $\mathcal{I}$ and $\mathcal{J}$ are coherent ideal sheaves, then so is $\mathcal{I} : \mathcal{J}$.*

Proof. First, suppose that $\mathcal{J}$ is generated by a single element $h \in {}_n\mathcal{H}(U)$ and $\mathcal{I}$ is generated by the elements $g_1, \dots, g_k \in {}_n\mathcal{H}(U)$. Then consider the map $\phi : {}_U\mathcal{H}^{k+1} \to {}_U\mathcal{H}$ defined by

$$\phi(f_0, \dots, f_k) = hf_0 - g_1f_1 - \cdots - g_kf_k$$

Then $\mathcal{I} : \mathcal{J}$ is the image of the kernel of ϕ under the projection $(f_0, \dots, f_k) \to f_0 : {}_U\mathcal{H}^{k+1} \to {}_U\mathcal{H}$. The kernel of ϕ is locally finitely generated by Oka's Theorem and, thus, so is $\mathcal{I} : \mathcal{J}$. It is then coherent by Corollary 12.4.

For the general case, for any point of U we may choose a neighborhood V in which $\mathcal{J}$ is finitely generated, say by $h_1, \dots, h_m$. Also, on V, $\mathcal{I} : \mathcal{J}$ is just the intersection of the sheaves $\mathcal{I} : \mathcal{J}_i$, where $\mathcal{J}_i$ is the sheaf of ideals on V generated by h_i. Thus, it follows from the previous corollary that $\mathcal{I} : \mathcal{J}$ is coherent.

We will prove that the ideal sheaf $\mathcal{I}_Y$ of a subvariety $Y \subset U$ of an open subset of $\mathbb{C}^n$ is coherent. The key to the proof is the following lemma:

12.7 Lemma. *Let U be an open set in $\mathbb{C}^n$ containing the origin. Let $f_1, \dots, f_p$ be a set of holomorphic functions on U with Y as its set of common zeroes and $\mathcal{I}$ as the ideal sheaf it generates. Suppose that $\mathcal{I}_0$ is a prime ideal of ${}_n\mathcal{H}_0$, and $\mathcal{I}_\lambda = \operatorname{id} Y_\lambda$ at all points $\lambda \in Y - Z$ where Z is a holomorphic subvariety of Y with $Z_0 \neq Y_0$. Then there is a polydisc $\Delta \subset U$ centered at 0 such that $\mathcal{I}_\lambda = \operatorname{id} Y_\lambda$ at all points $\lambda \in \Delta$.*

Proof. Since the functions $f_1, \dots, f_p$ determine Y, it follows from the Nullstellensatz that $\mathcal{I}_0 \subset \operatorname{id} Y_0 \subset \sqrt{\mathcal{I}_0}$. However, by assumption $\mathcal{I}_0$ is prime and, hence, we have that $\mathcal{I}_0 = \operatorname{id} Y_0 = \sqrt{\mathcal{I}_0}$. Let d be a function holomorphic in a polydisc Δ centered at 0 such that its germ d_0 at 0 belongs to $\operatorname{id} Z_0$ but not to $\operatorname{id} Y_0$. We may assume (by shrinking Δ if neccessary) that d vanishes on $Z \cap \Delta$ but does not vanish identically on $Y \cap \Delta$. By Corollary 12.6, the sheaf $\mathcal{I} : {}_n\mathcal{H}d$ on Δ is locally finitely generated and, hence, we may assume it is finitely generated by shrinking Δ. Let $g_1, \dots, g_q \in {}_n\mathcal{H}(\Delta)$ be a set of generators for this sheaf. Then, $dg_j \in \mathcal{I}(\Delta)$ and, in particular, its germ at 0 belongs to $\mathcal{I}_0$. However, this is a prime ideal and it does not contain the germ of d. Hence, we must have that the germ of g_j at 0 belongs to $\mathcal{I}_0$ for every j. That is, the germs at 0 of each g_j belong to the ideal at 0 generated by the f_i. If this is true at 0, it is true in a neighborhood of 0 and, hence, we may as well choose Δ small enough that it is true at every point of Δ. This implies that

$$\mathcal{I} : {}_n\mathcal{H}_0 d = \mathcal{I} \quad \text{on} \quad \Delta$$

For any point $\lambda \in Y \cap \Delta$ and any germ $f_\lambda \in \operatorname{id} Y_\lambda$, we choose a representative f of f_λ in a neighborhood V_λ of λ and note that, if the neighborhood is sufficiently small, then $\mathcal{I} : {}_n\mathcal{H}f$ will be finitely generated on V_λ by Corollary 12.6. The hypothesis that $\mathcal{I} = \operatorname{id} Y$ at points of $Y - Z$ implies that $\mathcal{I} : {}_n\mathcal{H}f = {}_n\mathcal{H}$ at points of $Y - Z$. Thus, if $h_1, \dots, h_m \in {}_n\mathcal{H}(V_\lambda)$ are generators of $\mathcal{I} : {}_n\mathcal{H}f$ on V_λ then the set of common zeroes of the h_i must lie in $Z \cap V_\lambda$. The function d vanishes on Z and, hence, by the Nullstellensatz, $d_\lambda^r \in (\mathcal{I} : {}_n\mathcal{H}f)_\lambda$ for some r. This means that $(d^r f)_\lambda \in \mathcal{I}_\lambda$ and, since $\mathcal{I} : {}_n\mathcal{H}_0 d = \mathcal{I}$ on Δ, it follows that $f \in \mathcal{I}_\lambda$. Therefore, $\mathcal{I} = \operatorname{id} Y$ at all points of $\Delta \cap Y$ and, hence, at all points of Δ.

Finally, we can prove:

12.8 Theorem. *If Y is a holomorphic subvariety of an open set $U \subset \mathbb{C}^n$, then its ideal sheaf $\mathcal{I}_Y$ is coherent.*

Proof. Fix a point of Y which we may assume is the origin. We assume at first that the germ Y_0 is irreducible. We may assume (after a change of variables if necessary) that the ideal $\operatorname{id} Y_0$ is strictly regular in the variables $z_{m+1}, \dots, z_n$ (see chapter 5). We choose a set of generators $f_1, \dots, f_p$ for $\operatorname{id} Y_0$ which includes the polynomials $p_j \in {}_m\mathcal{H}_0[z_j]$ for $m + 1 \leq j \leq n$ and $q_j \in {}_m\mathcal{H}_0[z_{m+1}, z_j]$ for $m + 2 \leq j \leq n$ of Lemma 5.9. Now, by

Corollary 5.18 the subvariety Y is actually a complex manifold outside a proper subvariety Z. Furthermore, it can be seen from the proof of Lemma 5.13 and Theorem 5.17 that the functions $p_{m+1}, q_{m+2}, \dots, q_n$ can be taken as the last $n-m$ coordinate functions of a coordinate system near any point of $Y - Z$ – a coordinate system in which Y is expressed as the set where these last $n-m$ coodinates vanish. It follows that the functions $f_1, \dots, f_p$ generate the germ of $\operatorname{id} Y$ at such points. It then follows from the previous lemma that these functions generate $(\mathcal{I}_Y)_\lambda = \operatorname{id} Y_\lambda$ at all points λ of some polydisc containing 0. Thus, $\mathcal{I}_Y$ is coherent.

If the germ Y_0 is not irreducible, then we write $Y = Y_1 \cup \dots \cup Y_q$ in some neighborhood of 0, where the varieties Y_j have irreducible germs at 0. Then $\mathcal{I}_{Y_j}$ is finitely generated in some neighborhood of 0 for each i by the previous paragraph. We may choose Δ to be a neighborhood in which this is true for all j. Then each ideal sheaf $\mathcal{I}_{Y_j}$ is coherent on Δ and, hence, so is the intersection $\mathcal{I}_{Y_j} \cap \dots \cap \mathcal{I}_{Y_q}$ by Corollary 12.5. But this intersection is just $\mathcal{I}_Y$ and, thus, the proof is complete.

We are now in a position to prove the strong form of Oka's Theorem:

12.9 Oka's Theorem on Varieties. *Let $\phi : {}_X\mathcal{H}^m \to {}_X\mathcal{H}^k$ be a morphism of analytic sheaves on an analytic variety X. Then $\ker \phi$ is coherent.*

Proof. This is a local result and so we may assume that X is a subvariety of an open set U contained in $\mathbb{C}^n$. The morphism ϕ is determined by a $k \times m$-matrix with entries which are holomorphic functions on X, and, hence, extend locally to be holomorphic in neighborhoods in U. Again, since we are proving a local result, we may as well assume that the entries of this matrix extend to be holomorphic in U. Then ϕ extends to a morphism $\tilde{\phi} : {}_U\mathcal{H}^m \to {}_U\mathcal{H}^k$ of analytic sheaves. Furthermore, by the previous result, we know that the ideal sheaf $\mathcal{I}_X$ is locally finitely generated and, hence, we may as well assume that it is finitely generated on U. Then we may represent it as the image of a ${}_U\mathcal{H}$-module morphism $\psi : {}_U\mathcal{H}^p \to {}_U\mathcal{H}$. Let $\psi^k : {}_U\mathcal{H}^{kp} \to {}_U\mathcal{H}^k$ denote the morphism that is just the direct sum of k copies of ψ. Then consider the morphism

$$\theta : {}_U\mathcal{H}^m \oplus {}_U\mathcal{H}^{kp} \to {}_U\mathcal{H}^k \quad \text{where} \quad \theta(f, g) = \tilde{\phi}(f) - \psi^k(g)$$

Then $f \in {}_U\mathcal{H}^m(V)$ is the first element of a pair $(f, g) \in \ker \theta$ over a neighborhood V if and only if $\tilde{\phi}(f) \in \mathcal{I}_X^k(V)$ – that is, if and only if the restriction of f to $X \cap V$ is zero. Thus, the kernel of ϕ on V can be characterized as those functions which are restrictions to $X \cap V$ of first elements of pairs (f, g) in the kernel of θ on V. Thus, it is clear that $\ker \phi$ is locally finitely generated if $\ker \theta$ is locally finitely generated. But $\ker \theta$ is locally finitely generated by Oka's Theorem. Thus, $\ker \phi$ is locally finitely generated. However, knowing this for all such morphisms ϕ allows us to prove in exactly the same manner as in Corollary 12.4 that a locally finitely generated subsheaf of a free analytic sheaf ${}_X\mathcal{H}^m$ is coherent. It follows that $\ker \phi$ is coherent.

We used in the preceding proof, that once we know that the kernel of any morphism ${}_X\mathcal{H}^m \to {}_X\mathcal{H}^k$ is locally finitely generated, then we may prove, as in Corollary 12.4, that any locally finitely generated subsheaf of a free finite rank analytic sheaf ${}_X\mathcal{H}^m$ is coherent. In fact, the analogs for varieties of Corollaries 12.4, 12.5, and 12.6 are all true and have exactly the same proofs. We state this fact as a Corollary to the above theorem.

12.10 Corollary. *If X is a holomorphic variety then*

(i) *any locally finitely generated analytic subsheaf of ${}_X\mathcal{H}^m$ is coherent;*
(ii) *the image and kernel of any morphism of analytic sheaves $\phi : {}_X\mathcal{H}^m \to {}_X\mathcal{H}^k$ are coherent;*
(iii) *the intersection of two coherent subsheaves of ${}_X\mathcal{H}^m$ is coherent;*
(iv) *if $\mathcal{I}$ and $\mathcal{J}$ are two coherent ideal sheaves in ${}_X\mathcal{H}$, then $\mathcal{I} : \mathcal{J}$ is also coherent.*

Once we have Theorem 12.8, it is a simple matter to prove the following:

12.11 Cartan's Theorem. *If X is any holomorphic variety and $Y \subset X$ is a holomorphic subvariety, then the ideal sheaf $\mathcal{I}_Y \subset {}_X\mathcal{H}$ is coherent.*

Proof. Again, this is a local result and so we may assume that X is a subvariety of an open set U in $\mathbb{C}^n$. Then Y is also a subvariety of U and, as such, its ideal sheaf is a coherent sheaf of ${}_U\mathcal{H}$-modules by Theorem 12.8. In particular, it is locally finitely generated as an ${}_U\mathcal{H}$-module and we may as well assume that it is finitely generated on U. But then the restriction to X of a set of generators of this ideal sheaf will be a set of generators over ${}_X\mathcal{H}$ of its ideal sheaf $\mathcal{I}_Y$. Thus, $\mathcal{I}_Y$ is locally finitely generated and, hence, coherent by Corollary 12.10.

The main result remaining to be proved in this chapter is that the kernel, image and cokernel of a morphism between coherent analytic sheaves are also coherent. The next three results lead up to this theorem.

12.12 Lemma. *If X is a holomorphic variety and $\phi : {}_X\mathcal{H}^m \to \mathcal{S}$ is a surjective morphism of analytic sheaves, then for each morphism of analytic sheaves $\psi : {}_X\mathcal{H}^k \to \mathcal{S}$ and each point $x \in X$ there is a neighborhood U of x in which ψ lifts to a morphism of analytic sheaves $\rho : {}_U\mathcal{H}^k \to {}_U\mathcal{H}^m$ such that $\phi \circ \rho = \psi$ on U.*

Proof. The stalk ${}_U\mathcal{H}^k_x$ is a free, hence projective, ${}_U\mathcal{H}_x$-module and so the morphism $\psi : {}_U\mathcal{H}^k_x \to \mathcal{S}_x$ lifts to a morphism $\rho : {}_U\mathcal{H}^k_x \to {}_U\mathcal{H}^m_x$ such that $\phi \circ \rho = \psi$. The morphism ρ is represented by a matrix with entries from ${}_U\mathcal{H}_x$ and we may assume that U is chosen small enough that each of these entries is represented by a holomorphic function on U. The resulting matrix defines a morphism of analytic sheaves $\rho : {}_U\mathcal{H}^k \to {}_U\mathcal{H}^m$ and the identity $\phi \circ \rho - \psi = 0$ is an identity which holds at x and, hence, in some neighborhood of x since ${}_U\mathcal{H}^k_x$ is finitely generated. Again, we may as well assume that neighborhood is U. This completes the proof.

12.13 Theorem. *If X is a holomorphic variety, then any locally finitely generated analytic subsheaf $\mathcal{M}$ of a coherent analytic sheaf $\mathcal{S}$ is also coherent.*

Proof. Since $\mathcal{S}$ is coherent, each point x of X has a neighborhood U for which there is a surjective morphism of analytic sheaves $\phi : {}_U\mathcal{H}^m \to \mathcal{S}|_U$ with a locally finitely generated kernel. Also, since $\mathcal{M}$ is locally finitely generated, we may assume U is chosen small enough that $\mathcal{M}|_U$ is finitely generated. Thus, there is a morphism of analytic sheaves $\psi : {}_U\mathcal{H}^k \to \mathcal{S}|_U$ which has $\mathcal{M}|_U$ as image. In order to prove that $\mathcal{M}$ is coherent, we need to prove that $\ker \psi$ is locally finitely generated, or, since our chosen point x is quite general, that $\ker \psi$ is finitely generated if U is chosen small enough. Also if U is chosen small enough, we may lift ψ to a morphism of analytic sheaves $\rho : {}_U\mathcal{H}^k \to {}_U\mathcal{H}^m$ such that $\phi \circ \rho = \psi$ on U by the preceding lemma.

Now the image of ρ is a finitely generated subsheaf of ${}_U\mathcal{H}^m$ as is the kernel of ϕ. Thus, both sheaves are coherent, as is their intersection, by Corollary 12.10. Since, $\ker\phi \cap \operatorname{im}\rho$ is coherent, it is, after shrinking U if necessary, the image of a morphism of analytic sheaves $\sigma : {}_U\mathcal{H}^p \to {}_U\mathcal{H}^m$. Using Lemma 12.12, we conclude that if U is chosen small enough, there is a morphism of analytic sheaves $\lambda : {}_U\mathcal{H}^p \to {}_U\mathcal{H}^k$ such that the following diagram is commutative:

$$\begin{tikzcd}
{}_U\mathcal{H}^m \arrow[r, equal] & {}_U\mathcal{H}^m \arrow[r, "\phi"] & \mathcal{S} \\
{}_U\mathcal{H}^p \arrow[u, "\sigma"] \arrow[r, "\lambda"] & {}_U\mathcal{H}^k \arrow[u, "\rho"] \arrow[r, equal] & {}_U\mathcal{H}^k \arrow[u, "\psi"]
\end{tikzcd}$$

Now for $y \in U$, $f \in {}_U\mathcal{H}^k_y$ is in the kernel of ψ if and only if $\rho(f) \in \ker\phi \cap \operatorname{im}\rho = \operatorname{im}\sigma$, that is, if and only if there exists $g \in {}_U\mathcal{H}^p_y$ such that $\rho(f) = \sigma(g) = \rho\lambda(g)$. The latter is equivalent to $f - \lambda(g) \in \ker\rho$. Since $\ker\rho$ is coherent, we may, after shrinking U again if necessary, assume that $\ker\rho$ is the image of a morphism $\tau : {}_U\mathcal{H}^q \to {}_U\mathcal{H}^k$. Then f belongs to $\ker\psi$ if and only if $f \in \operatorname{im}\lambda + \operatorname{im}\tau$. This implies that $\ker\psi$ is locally finitely generated and, hence, that $\mathcal{M}$ is coherent.

The following has a proof much like that of Corollary 12.5 (Problem 12.1):

12.14 Corollary. *If $\mathcal{M}$ and $\mathcal{N}$ are coherent subsheaves of a coherent analytic sheaf over a holomorphic variety X, then $\mathcal{M} \cap \mathcal{N}$ is also coherent.*

12.15 Theorem. *If $\mathcal{M}$ and $\mathcal{N}$ are coherent analytic sheaves over a holomorphic variety X, then the image, kernel, and cokernel of any morphism of analytic sheaves $\phi : \mathcal{M} \to \mathcal{N}$ are also analytic.*

Proof. Since $\mathcal{M}$ is locally finitely generated, so is $\operatorname{im}\phi$. Hence, $\operatorname{im}\phi$ is coherent by Theorem 12.13.

Fix a point $x \in X$. Since $\mathcal{N}$ and $\operatorname{im}\phi$ are locally finitely generated, there is a neighborhood U of x and surjective morphisms of analytic sheaves $\psi : {}_U\mathcal{H}^n \to \mathcal{N}|_U$ and $\rho : {}_U\mathcal{H}^p \to \operatorname{im}\phi|_U$. In fact, by choosing U small enough and using Lemma 12.12 to construct λ, we may construct the following commutative diagram of morphisms of analytic sheaves:

$$\begin{tikzcd}
0 \arrow[r] & \operatorname{im}\phi|_U \arrow[r] & \mathcal{N}|_U \arrow[r] & \operatorname{coker}\phi|_U \arrow[r] & 0 \\
& {}_U\mathcal{H}^p \arrow[u, "\rho"] \arrow[r, "\lambda"] & {}_U\mathcal{H}^n \arrow[u, "\psi"] \arrow[r, equal] & {}_U\mathcal{H}^n \arrow[u, "\sigma"] &
\end{tikzcd}$$

The top row of this diagram is exact and the vertical maps are all surjective. It evident from the diagram that the kernel of σ is $\psi^{-1}(\operatorname{im}\phi) = \operatorname{im}\lambda$. Thus, by definition, $\operatorname{coker}\phi$ is coherent.

It remains to prove that $\ker\phi$ is coherent. To this end, we fix $x \in X$ and choose a neighborhood U of x small enough that we may find surjective morphisms of analytic sheaves $\sigma : {}_U\mathcal{H}^m \to \mathcal{M}$ and $\rho : {}_U\mathcal{H}^n \to \mathcal{N}$ such that $\ker\rho$ and $\ker\sigma$ are finitely generated and, thus, coherent. We then construct the following commutative diagram after shrinking U appropriately:

$$\begin{tikzcd}
0 \arrow[r] & \ker\phi|_U \arrow[r] & \mathcal{M}|_U \arrow[r, "\phi"] & \mathcal{N}|_U \\
& {}_U\mathcal{H}^p \arrow[r, "\psi"] & {}_U\mathcal{H}^m \arrow[u, "\sigma"] \arrow[r, "\theta"] & {}_U\mathcal{H}^n \arrow[u, "\rho"] \\
& & {}_U\mathcal{H}^q \arrow[u, "\eta"] \arrow[r, equal] & {}_U\mathcal{H}^q \arrow[u, "\lambda"]
\end{tikzcd}$$

Here we use Lemma 12.12 to construct the map θ. Its image is finitely generated and, hence, coherent. Then $\ker\rho \cap \operatorname{im}\theta$ is coherent and, hence, is the image of a morphism λ as above. Then η is obtained as another application of Lemma 12.12. The map ψ is a morphism of analytic sheaves which maps onto the kernel of θ and it exists, for small enough U, by Corollary 12.10(ii). Now a simple diagram chase shows that $\sigma\circ(\psi+\eta) : {}_U\mathcal{H}^p \oplus {}_U\mathcal{H}^q \to \mathcal{M}|_U$ has $\ker\phi|_U$ as its image. This completes the proof that $\ker\phi$ is locally finitely generated and, hence, coherent.

The above theorem implies that the category of coherent analytic sheaves is an abelian category.

Finally, we have the following result:

12.16 Theorem. *If*

$$0 \longrightarrow \mathcal{K} \xrightarrow{\alpha} \mathcal{M} \xrightarrow{\beta} \mathcal{N} \longrightarrow 0$$

is an exact sequence of analytic sheaves and if any two of the three are coherent, then the third is also coherent.

Proof. Two of the three cases have already been proved in the preceding theorem. It remains to prove that if $\mathcal{K}$ and $\mathcal{N}$ are coherent then $\mathcal{M}$ is coherent. To this end, let x be a point of X and U a neighborhood of x for which we may find surjective morphisms of analytic sheaves $\rho : {}_U\mathcal{H}^k \to \mathcal{K}|_U$ and $\tau : {}_U\mathcal{H}^n \to \mathcal{N}|_U$ with finitely generated kernels. We may then use Lemma 12.12 to construct the following commutative diagram with exact rows:

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{K}|_U \arrow[r, "\alpha"] & \mathcal{M}|_U \arrow[r, "\beta"] & \mathcal{N}|_U \arrow[r] & 0 \\
0 \arrow[r] & {}_U\mathcal{H}^k \arrow[u, "\rho"] \arrow[r] & {}_U\mathcal{H}^{k+n} \arrow[u, "\sigma"] \arrow[r] & {}_U\mathcal{H}^n \arrow[u, "\tau"] \arrow[r] & 0
\end{tikzcd}$$

Here the morphisms on the bottom row are just the canonical injection and projection associated with writing ${}_U\mathcal{H}^{k+n}$ as ${}_U\mathcal{H}^k \oplus {}_U\mathcal{H}^n$. The morphism σ is constructed, for sufficiently small U, by lifting τ to a morphism $\tau' : {}_U\mathcal{H}^n \to \mathcal{M}|_U$ with $\beta\circ\tau' = \tau$ using Lemma 12.12, then writing ${}_U\mathcal{H}^{k+n}$ as ${}_U\mathcal{H}^k \oplus {}_U\mathcal{H}^n$ and defining σ by

$$\sigma(f \oplus g) = \alpha\circ\rho(f) + \tau'(g).$$

Now the fact that ρ and τ are surjective implies that σ is also surjective. Also, the kernels of the vertical maps form a short exact sequence

$$0 \longrightarrow \ker \rho|_U \longrightarrow \ker \sigma|_U \longrightarrow \ker \tau|_U \longrightarrow 0$$

in which the first and third terms are locally finitely generated and, hence, coherent. Thus, we may repeat the above argument for this sequence and conclude that, for sufficiently small U, $\ker \sigma|_U$ is also finitely generated. It follows that $\mathcal{M}$ is coherent and the proof is complete.

12. Problems

1. Prove Corollary 12.14.
2. Prove that if $\phi : {}_X\mathcal{H}^k \to {}_X\mathcal{H}^m$ is a morphism of analytic sheaves then $\phi_x : {}_X\mathcal{H}^k_x \to {}_X\mathcal{H}^m_x$ is surjective if and only if the matrix of holomorphic functions on X defining ϕ has rank m at x.
3. Use the result of the preceding problem to prove that if $\mathcal{M}$ is a coherent analytic sheave on a holomorphic variety X, then $\text{Support}(\mathcal{M}) = \{x \in X : \mathcal{M}_x \neq 0\}$ is a holomorphic subvariety of X.
4. Use the result of the preceding problem to prove that if $\phi : \mathcal{M} \to \mathcal{N}$ is a morphism of analytic sheaves between two coherent analytic sheaves, then $\{x \in X : \phi_x \text{ is not injective}\}$ and $\{x \in X : \phi_x \text{ is not surjective}\}$ are holomorphic subvarieties of X.

13. Projective Varieties

Complex projective space of dimension n is the algebraic variety defined in the following fashion: We consider the point set P^n which is $\mathbb{C}^{n+1} - \{0\}$ modulo the equivalence relation defined by

$$(\lambda z_0, \dots, \lambda z_n) \sim (z_0, \dots, z_n) \quad \forall \lambda \in \mathbb{C} - \{0\}$$

We shall define a topology and a ringed space structure on P^n and show that the resulting ringed space is an algebraic variety.

Let $p \in \mathbb{C}[z_0, \dots, z_n]$ be a homogeneous polynomial of some degree – say k. Then p does not define a function on P^n but the relation

$$p(\lambda z_0, \dots, \lambda z_n) = \lambda^k p(z_0, \dots, z_n)$$

means that the zero set of p is a union of equivalence classes and, hence, defines a subset of P^n. Similarly, the set of common zeroes of any set of homogeneous polynomials is a union of equivalence classes and defines a subset of P^n.

13.1 Definition. *An algebraic subset of P^n is a subset which is the set of common zeroes of some family of homogeneous polynomials.*

Of course, because subvarieties of $\mathbb{C}^{n+1}$ satisfy the descending chain condition, any algebraic subset of P^n is actually the zero set of a finite set of homogeneous polynomials.

It is easy to see that the collection of algebraic subsets of P^n is closed under finite union and arbitrary intersection. Thus, this collection may be taken as the closed sets in a topology for P^n.

13.2 Definition. *The Zariski topology on P^n is the topology in which the open sets are the complements of algebraic sets.*

We now associate to each integer k a sheaf $\mathcal{O}(k)$ on P^n in the following way: Let $\pi : \mathbb{C}^{n+1} - \{0\} \to P^n$ be the projection. If U is open in P^n then $\pi^{-1}(U)$ is open in $\mathbb{C}^{n+1} - \{0\}$ and, in fact, is the complement of the set of common zeroes of a finite set of homogeneous polynomials.

Definition 13.3. *If U is an open subset of P^n and k is an integer, we define $\mathcal{O}(k)(U)$ to be the space of regular functions on $\pi^{-1}(U)$ which are homogeneous of degree k.*

Clearly $\mathcal{O}(k)$ forms a sheaf for each k. Furthermore, if $f \in \mathcal{O}(j)(U)$ and $g \in \mathcal{O}(k)(U)$ then $fg \in \mathcal{O}(j+k)(U)$. Thus, $\oplus_{k=-\infty}^{\infty} \mathcal{O}(k)$ is a sheaf of graded rings. In particular, $\mathcal{O}(0)$ is a sheaf of rings and each $\mathcal{O}(k)$ is a sheaf of modules over $\mathcal{O}(0)$.

Definition 13.4. *We make P^n into a ringed space by defining the structure sheaf to be the sheaf of rings $\mathcal{O} = \mathcal{O}(0)$.*

Consider the open set U_i in P^n which is the complement of the algebraic set defined by the zero set of the i th coordinate function z_i. We define a map $\phi_i : U_i \to \mathbb{C}^n$ by $\phi_i \circ \pi = \psi_i$ where

$$\psi_i(z_0, \dots, z_n) = \left(\frac{z_0}{z_i}, \dots, \frac{z_{i-1}}{z_i}, \frac{z_{i+1}}{z_i}, \dots \frac{z_n}{z_i} \right)$$

The map ϕ_i is clearly well defined and, in fact:

13.5 Theorem. *The map $\phi_i : U_i \to \mathbb{C}^n$ is an isomorphism of ringed spaces.*

Proof. We may as well assume that $i = 0$. We first prove that ϕ_0 is a homeomorphism. The map ϕ_0 is clearly bijective and so, to prove that it is a homeomorphism, we must show that a set is closed in U_i if and only its image is closed in $\mathbb{C}^n$.

Every closed subset of U_i is a finite intersection of sets of the form $U_i \cap Z(p)$ where p is a homogeneous polynomial and $Z(p) \subset P^n$ is its zero set. But if $z_0 \neq 0$ then

$$p(z_0, z_1, \ldots, z_n) = 0 \quad \Leftrightarrow \quad p(1, z_1/z_0, \ldots, z_n/z_0) = 0$$

Thus, the image of $U_i \cap Z(p)$ under ϕ_0 is a subvariety of $\mathbb{C}^n$ – that is, a closed set. It follows that the image of every closed subset of U_i is closed in $\mathbb{C}^n$.

The closed subsets of $\mathbb{C}^n$ are finite intersections of zero sets of polynomials. Let q be a polynomial in $z_1, \ldots, z_n$ of degree k. Then we may define a homogeneous polynomial of degree k in $z_0, \ldots, z_n$ by

$$p(z_0, z_1, \ldots, z_n) = z_0^k q(z_1/z_0, \ldots, z_n/z_0)$$

and p is zero exactly at points which ϕ_0 maps to zeros of q. That is, $Z(p) \cap U_0$ is the inverse image under ϕ_0 of the zero set of q. It follows that the inverse image under ϕ_0 of any closed set in $\mathbb{C}^n$ is a closed set in U_0.

To finish the proof, we must show that ϕ_0 induces an isomorphism between the structure sheaves of U_0 and $\mathbb{C}^n$. This amounts to showing that, for each open set $V \subset U_0$, a complex valued function on V has the form $f \circ \pi$ for a homogeneous regular function f on $\pi^{-1}(V)$ if and only if it has the form $g \circ \phi_0$ for a regular function g on $\phi_0(V)$. In other words, for each open set $W = \phi_0(V)$, we must show that $g \to g \circ \psi_0$ is an isomorphism from the regular functions on W to the homogeneous regular functions on $\pi^{-1}(W)$. Since ψ_0 is algebraic and homogeneous it is clear that $g \to g \circ \psi_0$ is a ring homomorphism from $\mathcal{O}(W)$ to regular homogeneous functions on $\pi^{-1}(W)$. To see that it is an isomorphism, we simply note that its inverse is given by $f \to \tilde{f}$ where $\tilde{f}(z_1, \ldots, z_n) = f(1, z_1, \ldots, z_n)$. This completes the proof.

13.6 Theorem. *The ringed space P^n is an algebraic variety.*

Proof. We have that, $\{U_i\}$ is a cover of P^n by open subsets which are isomorphic as ringed spaces to $\mathbb{C}^n$. Thus, P^n is an algebraic prevariety It is also clear that, given any two points p and q of P^n, we may choose our coordinate system in $\mathbb{C}^{n+1}$ in such a way that one of the U_i contains both p and q. In other words, given any two points, there is an affine open subset containing both. By Theorem 10.6 this implies that P^n is actually an algebraic variety.

13.7 Definition. *A projective variety is an algebraic variety which is isomorphic to a closed subvariety of P^n for some n.*

To any algebraic variety we may associate in a canonical way a holomorphic variety. An algebraic variety is locally isomorphic, as a ringed space, to an algebraic subvariety of $\mathbb{C}^n$. There is a canonical way to associate to an algebraic subvariety V of $\mathbb{C}^n$ an holomorphic subvariety V^h – we simply give V the Euclidean topology instead of the Zariski topology and let its structure sheaf be the sheaf of holomorphic functions rather than the sheaf of regular functions. This means that for any algebraic variety we have a canonical way of changing the topology and ringed space structure on any affine open subset in such a way as to make it a holomorphic variety. This is canonical because any two ways of representing an affine open subset as subvarieties of complex Euclidean space are related by a biregular map between the two subvarieties – which will necessarily be a biholomorphic map between the associated holomorphic subvarieties. In particular, this implies that holomorphic space structures defined in this way on affine subsets of an algebraic variety will agree on intersections and, thus, define a global structure of a holomorphic variety. One thing that does need to be checked is that the resulting topological space is Hausdorff (Problem 13.2).

From the above discussion, we conclude that projective space P^n may also be considered as a holomophic variety – in fact, as a complex manifold, since the maps $\phi_i : U_i \to \mathbb{C}^n$ give local biholomorphic maps onto $\mathbb{C}^n$. We initially defined regular functions on $U \subset P^n$ as just regular functions on $\pi^{-1}(U)$ which are homogeneous of degree zero. It turns out that the analogous statement is true of holomorphic functions:

13.8 Definition. *For each integer k we define a sheaf $\mathcal{H}(k)$ on P^n with the Euclidean topology as follows: If $U \subset P^n$ is open in the Euclidean topology, then $\mathcal{H}(k)(U)$ is the space of functions in $\mathcal{H}(\pi^{-1}(U))$ which are homogeneous of degree k.*

13.9 Theorem. *The sheaf $\mathcal{H}(0)$ is a sheaf of rings canonically isomorphic to the structure sheaf $\mathcal{H}$ of P^n and the sheaves $\mathcal{H}(k)$ are sheaves of $\mathcal{H}$-modules.*

Proof. Clearly $\mathcal{H}(0)$ is a sheaf of rings and each $\mathcal{H}(k)$ is a sheaf of modules over $\mathcal{H}(0)$. Thus, we need only show that $\mathcal{H}(0)$ is canonically isomorphic to $\mathcal{H}$. The isomorphism is obviously the one which sends $f \in \mathcal{H}(U)$ to $f \circ \pi$. It is easy to see, using the definition of the holomorphic structure on P^n, that $\pi : Cn - 0 \to P^n$ is holomorphic, so that $f \to f \circ \pi$ is an homomorphism (obviously injective) of $\mathcal{H}$ to $\mathcal{H}(0)$. It only remains to show that this is surjective. If g is a homogeneous holomorphic function on an open set $\pi^{-1}(U)$ then $f(\pi(z)) = g(z)$ certainly defines a function f on U. The only question is whether or not it is holomorphic. It suffices to prove this in the case where $U \subset U_i$ for some i and, without loss of generality, we may assume that $i = 0$. Then to prove that f is holomorphic, we must show that $f \circ \phi_0^{-1}$ is holomorphic on $\phi_0(U) \subset \mathbb{C}^n$. But $f \circ \phi_0^{-1}(z_1, \dots, z_n) = g(1, z_1, \dots, z_n)$ and this is certainly holomorphic.

We now turn to the study of the sheaves $\mathcal{O}(k)$ and $\mathcal{H}(k)$ introduced above. The sheaves $\mathcal{O}(k)$ are sheaves of $\mathcal{O} = \mathcal{O}(0)$ modules. If V is an open subset of U_i then it is easy to see that $f \in \mathcal{O}(V)$ if and only if $z_i^k f \in \mathcal{O}(k)(V)$. Thus, $f \to z_i^k f$ defines an $\mathcal{O}|_{U_i}$-module isomorphism from $\mathcal{O}|_{U_i}$ to $\mathcal{O}(k)|_{U_i}$. In other words, $\mathcal{O}(k)$ is locally free of rank one as a sheaf of $\mathcal{O}$-modules. Exactly the same thing is obviously true of $\mathcal{H}(k)$ as a sheaf of $\mathcal{H}$ modules. A sheaf of modules with this property is called an *invertible sheaf* since such a sheaf always has an inverse under tensor product relative to the structure sheaf. In this case, the inverse of $\mathcal{O}(k)(\mathcal{H}(k))$ is obviously $\mathcal{O}(-k)(\mathcal{H}(-k))$ in view of the following theorem:

13.10 Theorem. *If j and k are integers then multiplication defines an isomorphism*

$$\mathcal{O}(j) \otimes_{\mathcal{O}} \mathcal{O}(k) \to \mathcal{O}(j + k).$$

The analogous statement is true for the sheaves $\mathcal{H}(k)$.

Proof. We multiplication map $f \otimes g \to fg$ defines a morphism of algebraic sheaves from $\mathcal{O}(j)\otimes_{\mathcal{O}}\mathcal{O}(k)$ to $\mathcal{O}(j+k)$. The only question is whether or not it is bijective. However, this map is bijective if and only if it is locally bijective. Thus, it suffices to show that the map is bijective on U_i for each i. However, on U_i, each $\mathcal{O}(k)$ is free of rank one with generator z_i^k while $\mathcal{O}(j) \otimes_{\mathcal{O}} \mathcal{O}(k)$ is free of rank one with generator $z_i^j \otimes z_i^k$. The multiplication map sends $z_i^j \otimes z_i^k$ to z_i^{j+k}, which is the generator of $\mathcal{O}(j+k)$ on U_i. This proves the theorem in the algebraic case. The proof in the analytic case is the same.

The sheaf of sections of a (finite dimensional) holomorphic vector bundle is a locally free finite rank sheaf of $\mathcal{H}$- modules and vice-verse. This is due to the fact that a holomorphic vector bundle is locally trivial and, thus, its sections may locally be identified with holomorphic vector valued functions. If the vector space has dimension m and a basis is chosen, then the space of holomorphic vector valued functions over an open set U is isomorphic to $\mathcal{H}(U)^m$. Thus, the module of holomorphic sections is locally free and of finite rank. The converse is just as easy (problem 13.3). Sheaves of modules which are locally free of rank one (invertible sheaves), such as our sheaves $\mathcal{H}(k)$, can be realized as sheaves of sections of holomorphic line bundles – vector bundles with one dimensional fiber. Note that locally free sheaves of finite rank are, of course, coherent, since coherence is a local property and locally free finite rank sheaves locally have the form $\mathcal{H}^m$ (or $\mathcal{O}^m$ in the algebraic case).

Our next main objective is to compute the sections and cohomology of the sheaves $\mathcal{O}(k)$ and $\mathcal{H}(k)$. It will be convenient to introduce three additional sheaves on P^n:

13.11 Definition. *Let $\mathcal{S}$ denote the sheaf which assigns to an open set $U \subset P^n$ the ring $\mathcal{O}(\pi^{-1}(U))$ and let $\mathcal{T}$ denote the sheaf which assigns to U the ring $\mathcal{H}(\pi^{-1}(U))$. We also let $\mathcal{T}_0$ be the subsheaf of $\mathcal{T}$ spanned by the subsheaves $\mathcal{H}(k)$.*

Thus, $\mathcal{O}(k)$ is the subspace of $\mathcal{S}$ consisting of elements of degree k and $\mathcal{S} = \oplus_{k=-\infty}^{\infty}\mathcal{O}(k)$, while $\mathcal{H}(k)$ is the subspace of $\mathcal{T}$ consisting of elements of degree k. Note that $\oplus_{k=-\infty}^{\infty}\mathcal{H}(k)$ is not $\mathcal{T}$ but the subsheaf $\mathcal{T}_0$.

We will compute cohomology using the Čech complex for the open cover $\{U_i\}$. Note that if $\alpha = (i_0, \dots, i_p)$ is a multi-index, then $U_\alpha = U_{i_0} \cap \cdots \cap U_{i_p} = U_{z_\alpha}$, where $z_\alpha = z_{i_0} \dots z_{i_p}$ and U_{z_α} is the complement of the algebraic set in P^n determined by the vanishing of z_α. Thus, U_α is the subset of $U_{z_{i_0}} \simeq \mathbb{C}^n$ on which the regular function $z_{i_0}^{-p} z_{i_1} \dots z_{i_p}$ does not vanish. This implies that U_α is an affine variety for each multi-index α from which it follows that $\mathcal{O}(k)$ is acyclic on U_α for each α and each k and, hence, that $\{U_i\}$ is a Leray cover for $\mathcal{O}(k)$ for each k. Thus, Leray's Theorem applies and we may, in fact, compute the sheaf cohomology of $\mathcal{O}(k)$ or $\mathcal{S}$ using the Čech complex for $\{U_i\}$. We also have that each U_i is biholomorphically equivalent to $\mathbb{C}^n$, each $\mathcal{H}(k)$ is a free $\mathcal{H}$ module on each U_i and each U_α is a Cartesian product of planes and punctured planes. It follows from our work on Dolbeault cohomology (see the remark following Theorem 11.5) that $\mathcal{H}(k)$ is also acyclic on each U_α and, hence, that $\{U_i\}$ is also a Leray cover for the sheaf $\mathcal{H}(k)$. Again, Leray's theorem implies that we may compute the cohomology of $\mathcal{H}(k)$ or $\mathcal{T}_0$ using Čech cohomology for the cover $\{U_i\}$.

For what follows, we need to use the definition of Čech cohomology in which only alternating cochains are used – that is, cochains f for which $f(\sigma(\alpha)) = \mathrm{sgn}(\sigma)f(\alpha)$ for each permutation σ and each multi-index α. This is equivalent to the definition we used earlier – the proof of this may be found in nearly any standard text in which Čech Theory is developed (or do Problem 13.1).

We will compute Čech cohomology for $\mathcal{O}(k)$ and $\mathcal{H}(k)$ by computing it for $\mathcal{S}$ and $\mathcal{T}_0$ and then projecting out the part which is homogeneous of degree k. We begin by expressing each element of $\mathcal{S}(U_\alpha)$ and $\mathcal{T}(U_\alpha)$ in terms of its *Fourier coeficients.* That is, if $f \in \mathcal{T}(U_\alpha)$ for $\alpha = (i_0, \dots, i_p)$ then f is a holomorphic function on the set of all $(z_0, \dots, z_n) \in \mathbb{C}^{n+1}$ such that $z_{i_j} \neq 0$ for $j = 0, \dots, p$. Such a function has a Laurent series expansion

$$\sum a_{m_0 \dots m_n} z_0^{m_0} \dots z_n^{m_n}$$

in which $a_{m_0 \dots m_n}$ is allowed to be non-zero only for terms such that either $m_i \geq 0$ or $m_i < 0$ and i is one of the i_j appearing in α. For such an expansion to converge on U_α, the coeficients $a_{m_0 \dots m_n}$ must decay faster than $r^{|m_0| + \dots |m_n|}$ for every positive number r. In other words, we may identify $\mathcal{T}(U_\alpha)$ with the space of functions $m \to a_m : Z^{n+1} \to \mathbb{C}$ which decay at infinity faster than any geometric series and which are supported on

$$\{m = (m_0, \dots, m_n) : m_i \geq 0 \quad \text{if} \quad i \notin \{i_0, \dots, i_p\}\}$$

Of course, $\mathcal{S}(U\alpha)$ is the subspace of $\mathcal{T}(U\alpha)$ consisting of sums, as above, with coeficients which are non-vanishing for only finitely many indices.

The spaces of global sections $\mathcal{S}(P^n)$ and $\mathcal{T}(P^n)$ may be described as above if we simply use for α the empty index. Thus, $\mathcal{T}(P^n)$ is the set of sums as obove, which satisfy the decay condition and have Fourier coeficients which vanish except when all indices m_i are non-negative. Similarly, $\mathcal{S}(P^n)$ is the set of sums, as above, with coeficients which are finitely non-zero and are zero except for all non-negative indices m_i.

In $\mathcal{S}(U_\alpha)$ and $\mathcal{T}(U_\alpha)$, the elements which are homogeneous of degree k are those whose Fourier coeficient functions are supported on the set where $m_0 + m_1 + \dots + m_n = k$. The coeficient functions of global sections are also supported on the set where all $m_i \geq 0$. Then, since there are only finitely many tuples $(m_0, \dots, m_n)$ of non-negative integers with $m_0 + m_1 + \dots + m_n = k$ we conclude that $\mathcal{O}(k)(P^n)$ and $\mathcal{H}(k)(P^n)$ are finite dimensional, are equal, and, in fact:

13.12 Theorem. *For each integer k, $H^0(P^n, \mathcal{O}(k))$ and $H^0(P^n, \mathcal{H}(k))$ are both equal to the space of homogeneous polynomials of degree k in $\mathbb{C}[z_0, \dots, z_n]$. In particular, $H^0(P^n, \mathcal{O}) = H^0(P^n, \mathcal{H}) = \mathbb{C}$ and $H^0(P^n, \mathcal{O}(k)) = H^0(P^n, \mathcal{H}(k)) = 0$ for $k < 0$.*

13.13 Theorem. *For each integer k we have*

(a) $H^n(P^n, \mathcal{O}(-n-1)) \simeq \mathbb{C}$;
(b) $H^n(P^n, \mathcal{O}(-n-k-1))$ *is the vector space dual of* $H^0(P^n, \mathcal{O}(k))$ *under the pairing*

$$H^0(P^n, \mathcal{O}(k)) \times H^n(P^n, \mathcal{O}(-n-k-1)) \to H^n(P^n, \mathcal{O}(-n-1)) \simeq \mathbb{C}$$

induced by multiplication;

(c) $H^p(P^n, \mathcal{O}(k)) = 0$ *if p is not 0 or n.*

Also, the same statements are true with $\mathcal{O}(k)$ replaced by $\mathcal{H}(k)$.

Proof. An alternating Čech cochain vanishes on any multi-index with a repeated entry. Thus, n is the largest value of p for which there are non-zero alternating p-cochains for $\{U_i\}$, from which it follows that $H^p(P^n, S) = 0$ for $p > n$.

Via the Fourier representation, we regard a p-cochain for $\mathcal{S}$ and $\{U_i\}$ as an alternating function a which assigns to each $\alpha = (i_0, \ldots, i_p)$ a finitely non-zero function $(m_0, \ldots, m_n) \to a^{\alpha}_{m_0 \ldots m_n} : Z^{n+1} \to \mathbb{C}$ with support in the set

$$W_\alpha = \{(m_0, \ldots, m_n : m_i \geq 0 \quad \text{for} \quad i \notin \{i_0, \ldots, i_n\}\}.$$

Now multiplication by the characteristic function of a given subset L of Z^{n+1} is a projection operator P_L on cochains which commutes with the coboundary operator and has as range the cochains with support in L. Thus, P_L also defines a projection operator on Čech cohomology as well. In particular, the cohomology of $\mathcal{O}(q)$ can be obtained from that of $\mathcal{S}$ by applying the projection P_{L_k} determined by the set $L_k = \{(m_0, \ldots, m_n) : \sum m_i = k\}$. Also, it makes sense to talk about an element of cohomology of $\mathcal{S}$ being supported in a set L – that is, $\xi \in H^p(P^n, \mathcal{S})$ is supported in L if $P_L \xi = \xi$. With this in mind, the strategy of the proof will be to prove that, for $p > 0$, every element of $H^p(P^n, \mathcal{S})$ is supported in the set $K = \{(m_0, \ldots, m_n) : m_i < 0 \quad \forall i\}$.

We know that $H^p(U_j, \mathcal{S}) = 0$ for $p > 0$, since U_j is affine and $\mathcal{S}$ is quasi-coherent (it is an infinite direct sum of coherent sheaves). In terms of our Fourier coeficient picture, this means that, for a fixed j and $p > 0$, any p-cocycle $a = \{a^{\alpha}_{m_0, \ldots, m_n}\}$ for $\mathcal{S}$ and $\{U_i\}$ is, when restricted to U_j, the coboundary of some $p-1$-cochain $b = \{b^{\beta}_{m_0, \ldots, m_n}\}$ for $\mathcal{S}$ and $\{U_j \cap U_i\}_i$. Now generally b will have some non-zero coeficients with negative index m_j and a β which does not include j among its entries. This, of course, just means that b is not generally a cochain for $\{U_i\}$. However, if $W_j = \{m \in Z^{n+1} : m_j \geq 0\}$ then $P_{W_j} b$ is a cochain for $\{U_i\}$ which is mapped by the coboundary onto $P_{W_j} a$. This implies that P_{W_j} kills elements of $H^p(P^n, \mathcal{S})$ for $p > 0$ and, hence, that $H^p(P^n, \mathcal{S})$ is supported on the complement of W_j. Since this is true for every j, we conclude that $H^p(P^n, \mathcal{S})$ is supported on K as claimed. This implies that $H^p(P^n, \mathcal{S}) = 0$ for $0 < p < n$ since every p-cochain for $p < n$ is supported on the complement of K. We already know that $H^p(P^n, \mathcal{S}) = 0$ for $p < 0$ and for $p > n$. This completes the proof of (c) in the algebraic case, since we know that we can obtain $H^p(\mathcal{O}(k))$ as a direct summand of $H^p(P^n, \mathcal{S})$.

For an n-cochain, the only multi-index we need consider is $\alpha = (0, 1, \ldots, n)$, since all others without repeated entries are permutations of this one. For this index α the set W_α is all of Z^{n+1}. Thus, an n-cochain for $\mathcal{S}$ and $\{U_i\}$ is just a single coeficient function $a = a_{m_0, \ldots, m_n}$ with no restriction other than that only finitely many coeficients are non-zero. Now we know from the previous paragraph that $H^n(P^n, S)$ is supported on K. However, every n- cochain is a cocycle and the space of $(n-1)$ cochains is supported in the complement of K. Thus, it follows that $H^n(P^n, S)$ is actually isomorphic to the space of coeficient functions supported in K. This proves (a) since an index $(m_0, \ldots, m_n)$ which belongs to K, that is $m_i < 0 \quad \forall i$, can satisfy $\sum m_i = -n-1$ if and only if each m_i is -1. Thus, $H^{-n-1}(P^n, \mathcal{O}(k))$ consists of functions with a single non-vanishing coeficient corresponding to the index $(-1, \ldots, -1)$. This proves part (a).

We have from the previous theorem that $H^0(P^n, \mathcal{S})$ may be regarded as the space of finitely non-zero coeficient functions which are supported in $K_+ = \{(m_0, \dots, m_n) : m_i \geq 0 \quad \forall i\}$ and from the previous paragraph, that $H^n(P^n, \mathcal{S})$ may be regarded as the space of finitely non-zero coeficient functions which are supported on K. Clearly reflection through the hyperplane $\sum m_i = (-n-1)/2$ is a bijection between K_+ and K. This pairs basis elements for $H^0(P^n, \mathcal{S})$ and $H^n(P^n, \mathcal{S})$. Clearly it pairs functions supported on L_k with functions supported on L_{-n-k-1} and, thus, it pairs basis elements for $H^0(P^n, \mathcal{O}(k))$ and $H^n(P^n, \mathcal{O}(-n-k-1k))$ as require in the theorem. This completes the proof of part (a) in the algebraic case.

The argument is almost the same in the analytic case, where $\mathcal{O}(k)$ is replaced by $\mathcal{H}(k)$. The difference is that we must work with $\mathcal{T}_0$ rather than $\mathcal{S}$. However, $\mathcal{T}_0$ is just the direct sum of the $\mathcal{H}(k)$ and, hence, also has vanishing cohomology on each U_i. The rest of the argument is the same if one observes that the projection operators which play such a prominent role preserve the faster than exponential decay at infinity that defines coeficient functions in $\mathcal{T}(U_\alpha)$.

Since we use the same open cover by Zariski open sets to compute both, there is a map $H^p(P^n, \mathcal{O}(k)) \to H^p(P^n, \mathcal{H}(k))$ defined by inclusion. Clearly, the above theorem has as a consequence:

13.14 Corollary. *The map $H^p(P^n, \mathcal{O}(k)) \to H^p(P^n, \mathcal{H}(k))$ is an isomorphism for every p and every k.*

We complete this section with some results which show how to make strong use of the sheaves $\mathcal{O}(k)$.

13.15 Definition. *If $\mathcal{F}$ is a coherent algebraic sheaf on P^n then we set $\mathcal{F}(k) = \mathcal{F} \otimes_{\mathcal{O}} \mathcal{O}(k)$. The sheaf $\mathcal{F}(k)$ is said to be obtained from $\mathcal{F}$ through twisting it by $\mathcal{O}(k)$.*

Note that $\oplus_k \mathcal{F}(k)$ is a sheaf of modules over the sheaf of rings $\mathcal{S} = \oplus_k \mathcal{O}(k)$ where the product hf of $h \in \mathcal{O}(k)$ and $f \in \mathcal{F}(m)$ is $h \otimes f \in \mathcal{F}(k+m)$.

Note also, that, since $\mathcal{O}(k)$ is locally free of rank one, $\mathcal{F}(k)$ is locally isomorphic to $\mathcal{F}$, but it is not generally globally isomorphic to $\mathcal{F}$ because of the twist introduced by tensoring with $\mathcal{O}(k)$.

13.16 Theorem. *If $\mathcal{F}$ is a coherent algebraic sheaf on P^n then for some $k \geq 0$, the sheaf $\mathcal{F}(k)$ is generated by finitely many of its global sections.*

Proof. For each i the open set U_i is affine and, hence, the coherent sheaf $\mathcal{F}|_{U_i}$ is the image under localization of $\Gamma(U_i, \mathcal{F})$ which is finitely generated. Let $\{f_{ij}\}_j$ be a finite generating set of sections for $\Gamma(U_i, \mathcal{F})$. Now z_i is a global section of $\mathcal{O}(1)$ which vanishes exactly on the complement of U_i. For some integer m, which may be chosen large enough to work for all i, j, the product $z_i^m f_{ij}$ extends to be a global section of $\mathcal{F}(m)$ (problem 13.4). The resulting collection of global sections $\{z_i^m f_{ij}\}_{ij}$ clearly generate $\mathcal{F}(m)$ on each U_i and, hence, on all of P^n.

13.17 Theorem. *If $\mathcal{F}$ is a coherent algebraic sheaf on P^n, then $\mathcal{F}$ is the quotient of a sheaf which is a finite direct sum of sheaves of the form $\mathcal{O}(k)$*

Proof. By the previous theorem, there exists a non-negative integer m such that $\mathcal{F}(m)$ is generated by finitely many global sections. If there are r of these sections then they determine a surjection $\mathcal{O}^r \to \mathcal{F}(m)$. On tensoring this with $\mathcal{O}(-m)$, we get a surjection $\mathcal{O}(-m)^r \to F$ as required.

13.18. *If $\mathcal{F}$ is a coherent algebraic sheaf on P^n then*

(a) *$H^p(P^n, \mathcal{F})$ is finite dimensional for each p and vanishes for $p > n$;*
(b) *there exists an integer m_0 such that $H^p(P^n, \mathcal{F}(m)) = 0$ for all $p > 0$ and all $m > m_0$.*

Proof. That $H^p(P^n, \mathcal{F})$ vanishes for $p > n$ is just the fact that we can use alternating Čech cochains for the cover $\{U_i\}$ to compute it. Note that this implies that the m_0 in part (b) can be chosen independent of p if it can be chosen depending on p since there are only finitely many ps to worry about. With these things in mind, the proof is by reduction on p. We assume both statements are true for $p+1 \leq n+1$ and then prove they are also true for p. We express $\mathcal{F}$ as a quotient of a finite direct sum $\oplus_i \mathcal{O}(k_i)$ so that we have a short exact sequence

$$0 \to \mathcal{K} \to \oplus_i \mathcal{O}(k_i) \to \mathcal{F} \to 0$$

Then part(a) for p follows from the long exact sequence for coholomogy associated to this short exact sequence, the assumption that the theorem is true for $p+1$ and Theorem 13.13. Part(b) follows in a similar fashion from Theorem 13.13 and the long exact sequence associated to the short exact sequence

$$0 \to \mathcal{K}(m) \to \oplus_i \mathcal{O}(m + k_i) \to \mathcal{F}(m) \to 0.$$

obtained by twisting by the sheaf $\mathcal{O}(m)$. This completes the proof.

13. Problems

1. Verify that Čech cohomology does not change if the Čech cochains of chapter 9 are replaced by alternating Čech cochains.
2. Prove that if V is an algebraic variety then the topological space V^h constructed by giving each affine open subset of V its Euclidean topology is Hausdorph. Hint: Use the condition of Definition 10.5 to prove that the diagonal in $V^h \times V^h$ is closed.
3 Prove that if V is a holomorphic variety and $\mathcal{M}$ is a locally free finite rank sheaf of ${}_V\mathcal{H}$-modules, then there is a holomorphic vector bundle such that $\mathcal{M}$ is isomorphic to its sheaf of holomorphic sections.
4. Model the proof of Theorem 10.12(ii) to prove that if $\mathcal{F}$ is a coherent sheaf on P^n and $f \in \Gamma(U_i, \mathcal{F})$, then for some positive integer m the section $z_i^m f \in \Gamma(U_i, \mathcal{F}(m))$ extends to a global section of $\mathcal{F}(m)$.

14. Algebraic vs. Analytic Sheaves – Serre's Theorems

In this section we prove that, for all practical purposes, analytic projective varieties and algebraic projective varieties and their coherent sheaves are the same. These are the results of Serre's famous GAGA paper *Geométrie Algébrique et Géométrie Analytique.*

Note that an algebraic subvariety V of $\mathbb{C}^n$ also has the structure of a holomorphic subvariety V^h of $\mathbb{C}^n$. This structure does not depend on the embedding of V in $\mathbb{C}^n$, since, if $V_1 \subset \mathbb{C}^n$ and $V_2 \subset \mathbb{C}^m$ are algebraic subvarieties which are isomorphic as algebraic varieties, then there is a biregular map ϕ of V_1 onto V_2. Such a map is also biholomoprhic and, hence, is an isomorphism between V_1^h and V_2^h. In other words, an affine variety V has a unique structure of a holomorphic variety V^h with the property that any algebraic embedding of V in the Zariski space $\mathbb{C}^n$ is also a holomorphic embedding of V^h in the Euclidean space $\mathbb{C}^n$.

14.1 Theorem. *If X is an algebraic variety then there is a unique holomorphic variety X^h which is X as a point set, for which every open subset of X is open in X^h and for which every affine open subset has its natural holomorphic structure.*

Proof. The uniqueness is clear, since the condition on affine subvarieties fixes the holomorphic structure locally and that fixes it period. To show existence, we must show that the natural holomorphic structures on two affine open sets agree on their intersection. This just amounts to the fact that a biregular map between a Zariski open subset U of an algebraic subvariety of $\mathbb{C}^n$ and a Zariski open subset of an algebraic subvariety of $\mathbb{C}^m$ is also a biholomorphic map. Thus, the topology and ringed space stucture of X^h are well defined and X^h is locally isomorphic to a subvariety of Euclidean space. By problem 13.2, X^h is Hausdorff. The topology is second countable because X is covered by finitely many affine open sets and each of these clearly has a second countable Euclidean topology. This completes the proof.

Our next task will be to show that a coherent algebraic sheaf $\mathcal{M}$ on an algebraic variety X gives rise to a coherent analytic sheaf $\mathcal{M}^h$ on X^h. First, if $\mathcal{M}$ is any sheaf on X let $\mathcal{M}'$ be the sheaf on X^h which is the inverse image of $\mathcal{M}$ under the continuous map $X^h \to X$. Thus,

$$\mathcal{M}'(U) = \cap\{\mathcal{M}(W) : U \subset W \quad \text{open in} \quad X\}.$$

Note that at each point $x \in X$, the stalks $\mathcal{M}_x$ and $\mathcal{M}'_x$ are the same.

14.2 Definition. *If $\mathcal{M}$ is any sheaf of $\mathcal{O}$-modules on an algebraic variety X, then we define a corresponding analytic sheaf $\mathcal{M}^h$ on X^h by $\mathcal{M}^h = \mathcal{H} \otimes_{\mathcal{O}'} \mathcal{M}'$.*

The sheaf $\mathcal{M}^h = \mathcal{H} \otimes_{\mathcal{O}'} \mathcal{M}'$ is the sheaf of germs of the presheaf $U \to \mathcal{H}(U) \otimes_{\mathcal{O}'(U)} \mathcal{M}'(U)$ and so its stalk at $x \in X$ is

$$\mathcal{M}^h_x = \varinjlim\{\mathcal{H}(U) \otimes_{\mathcal{O}'(U)} \mathcal{M}'(U) : x \in U\} = \mathcal{H}_x \otimes_{\mathcal{O}'_x} \mathcal{M}'_x = \mathcal{H}_x \otimes_{\mathcal{O}_x} \mathcal{M}_x.$$

The second equality above follows easily from the definitions of direct limit and tensor product (Problem 14.1).

The next theorem is a direct application of our work on faithful flatness in Chapter 7.

14.3 Theorem. *If X is an algebraic variety, then $\mathcal{M} \to \mathcal{M}^h$ is an exact functor from the category of sheaves of $\mathcal{O}$-modules on X to the category of sheaves of $\mathcal{H}$-modules on X^h.*

Proof. It is clear that $\mathcal{M} \to \mathcal{M}^h$ is a functor from sheaves of $\mathcal{O}$-modules to sheaves of $\mathcal{H}$-modules. We have that $\mathcal{M}^h_x = \mathcal{H}_x \otimes_{\mathcal{O}_x} \mathcal{M}_x$ and, by Corollary 7.15, $\mathcal{H}_x$ is faithfully flat over $\mathcal{O}_x$. Since exactness is a stalkwise property, it follows that $\mathcal{M} \to \mathcal{M}^h$ is an exact functor.

14.4 Theorem. *For an algebraic variety X, the functor $\mathcal{M} \to \mathcal{M}^h$ takes*
- (a) *$\mathcal{O}^m$ to $\mathcal{H}^m$ for each m;*
- (b) *coherent algebraic sheaves to coherent analytic sheaves;*
- (c) *the ideal sheaf in $\mathcal{O}$ of an algebraic subvariety V to the ideal sheaf in $\mathcal{H}$ of the corresponding holomorphic subvariety V^h.*

Proof. Part (a) is trivial since $\mathcal{H}(U) \otimes_{\mathcal{O}'(U)} \mathcal{O}'(U) = \mathcal{H}(U)$ for any Euclidean open set U.

If $\mathcal{M}$ is a coherent algebraic sheaf on X, then for each affine open subset $U \subset X$, there is a finitely generated $\mathcal{O}(U)$-module M such that $\mathcal{M}|_U \simeq \mathcal{O} \otimes_{\mathcal{O}(U)} M$. Since M is finitely generated, there is a surjective morphism $\mathcal{O}(U)^m \to M$ and, since $\mathcal{O}(U)$ is Noetherian the kernel of this morphism is also finitely generated. Thus, there is a morphism $\mathcal{O}(U)^k \to \mathcal{O}(U)^m$ so that the sequence

$$\mathcal{O}(U)^k \longrightarrow \mathcal{O}(U)^m \longrightarrow M \longrightarrow 0$$

is exact. If we apply the localization functor $\mathcal{O} \otimes_{\mathcal{O}(U)} (\quad)$ to this sequence, we obtain an exact sequence of sheaves

$$\mathcal{O}^k|_U \longrightarrow \mathcal{O}^m|_U \longrightarrow \mathcal{M}|_U \longrightarrow 0$$

On applying the functor $\mathcal{M} \to \mathcal{M}^h$ to this sequence and then using the previous theorem and part (a) above, we conclude there is an exact sequence

$$\mathcal{H}^k|_U \longrightarrow \mathcal{H}^m|_U \longrightarrow \mathcal{M}^h|_U \longrightarrow 0$$

Thus, by definition, $\mathcal{M}^h$ is coherent. This completes the proof of part (b).

If V is an algebraic subvariety of X, then the inclusion $0 \to \mathcal{I}_V \to \mathcal{O}$ yields an inclusion $0 \to \mathcal{I}_V^h \to \mathcal{H}$ by Theorem 14.3. The image consists of the sheaf of ideals generated in $\mathcal{H}$ by the image of $\mathcal{I}'_V$. By Theorem 7.13, this is the ideal sheaf of V^h in $\mathcal{H}$. This completes the proof of (c).

Now suppose that $i : Y \to X$ is the inclusion of a holomorphic subvariety in an holomorphic variety X and $\mathcal{M}$ is a sheaf of ${}_Y\mathcal{H}$-modules on Y. Then $i_*\mathcal{M}$ – the extension of $\mathcal{M}$ by zero – is a sheaf on X which is supported on Y. It is certainly a sheaf of $i_{*Y}\mathcal{H}$-modules. However, $i_{*Y}\mathcal{H}$ is the quotient sheaf ${}_X\mathcal{H}/\mathcal{I}_Y$ and, hence, $i_*\mathcal{M}$ is also a sheaf of ${}_X\mathcal{H}$-modules. Exactly the same thing is true for an embedding $i : Y \to X$ of algebraic varieties and a sheaf $\mathcal{M}$ of ${}_Y\mathcal{O}$-modules – the sheaf $i_*\mathcal{M}$ is then a sheaf of ${}_X\mathcal{O}$-modules.

14.5 Theorem. *If $i : Y \to X$ is an embedding of a holomorphic variety Y in an holomorphic variety X and $\mathcal{M}$ is a coherent analytic sheaf on Y, then $i_*\mathcal{M}$ is a coherent analytic sheaf on X. The analogous statement is true for an embedding of algebraic varieties and a coherent algebraic sheaf.*

Proof. . We prove this in the holomorphic case. The proof in the algebraic case is even easier.

If $\mathcal{M}$ is a coherent analytic sheaf on Y then for each $y \in Y$ there is a neighborhood $V \subset Y$ of y and an exact sequence

$$ {}_V\mathcal{H}^k \xrightarrow{\beta} {}_V\mathcal{H}^m \xrightarrow{\alpha} \mathcal{M}|_V \longrightarrow 0 $$

If $V = U \cap Y$ for an open set $U \subset X$, then we may rewrite this in terms of the image of this sequence under the exact functor i_* restricted to U:

$$ (i_{*Y}\mathcal{H}^k)|_U \xrightarrow{\beta} (i_{*Y}\mathcal{H}^m)|_U \xrightarrow{\alpha} (i_*\mathcal{M})|_U \longrightarrow 0 $$

Now $i_{*Y}\mathcal{H}$ is the cokernel of the inclusion $j : \mathcal{I}_Y \to {}_X\mathcal{H}$ and the ideal sheaf $\mathcal{I}_Y$ is coherent and, hence, locally finitely generated, by Theorem 12.11. Thus, after possibly shrinking U if necessary, we may find a morphism $\gamma : {}_X\mathcal{H}^p \to {}_X\mathcal{H}$ with cokernel equal to $i_{*Y}\mathcal{H}$. This and the lifting lemma (Lemma 12.12) allow us to construct the following commutative diagram

$$
\begin{tikzcd}
{}_X\mathcal{H}^{kp}|_U \arrow[r, "\mu"] \arrow[d, "\gamma^k"'] & {}_X\mathcal{H}^{mp}|_U \arrow[d, "\gamma^m"'] & & \\
{}_X\mathcal{H}^{k}|_U \arrow[r, "\nu"] \arrow[d, "j^k"'] & {}_X\mathcal{H}^{m}|_U \arrow[d, "j^m"'] & & \\
i_*({}_Y\mathcal{H}^k)|_U \arrow[r, "\beta"] \arrow[d] & i_*({}_Y\mathcal{H}^m)|_U \arrow[r, "\alpha"] \arrow[d] & i_*\mathcal{M}|_U \arrow[r] & 0 \\
0 & 0 & &
\end{tikzcd}
$$

with exact rows and columns. From the diagram it is evident that $\alpha \circ j^m : {}_X\mathcal{H}^m|_U \to i_*\mathcal{M}|_U$ is surjective with kernel equal to the sum of the images of $\gamma^m : {}_X\mathcal{H}^{mp}|_U \to {}_X\mathcal{H}^m|_U$ and $\nu \circ \gamma^k : {}_X\mathcal{H}^{kp}|_U \to {}_X\mathcal{H}^m|_U$. Thus, $i_*\mathcal{M}$ is coherent.

This is a very important fact. It will allow us to prove a great many theorems about coherent sheaves on projective varieties once we know they are true for coherent sheaves on projective space. To make use of it in our study of the functor $\mathcal{M} \to \mathcal{M}^h$ we need the following:

14.6 Theorem. *If $i : Y \to X$ is an embedding of algebraic varieties and $\mathcal{M}$ is a coherent algebraic sheaf on Y, then $(i_*\mathcal{M})^h = i_*(\mathcal{M}^h)$.*

Proof. There is an inclusion $\mathcal{M}' \to \mathcal{M}^h$ of sheaves of ${}_Y\mathcal{O}'$-modules which yields a morphism $i_*(\mathcal{M}') \to i_*(\mathcal{M}^h)$ of sheaves of ${}_X\mathcal{O}'$-moldules. Clearly, $i_*(\mathcal{M}') = (i_*\mathcal{M})'$ and so we have a morphism $(i_*\mathcal{M})' \to i_*(\mathcal{M}^h)$ of sheaves of ${}_X\mathcal{O}'$-modules. On tensoring this morphism with ${}_X\mathcal{H}$ we obtain a morphism $(i_*\mathcal{M})^h \to (i_*(\mathcal{M}^h))^h = {}_X\mathcal{H} \otimes_{{}_X\mathcal{O}'} i_*(\mathcal{M}^h)$. Since $i_*(\mathcal{M}^h)$ is already an ${}_X\mathcal{H}$-module, there is a morphism ${}_X\mathcal{H} \otimes_{{}_X\mathcal{O}'} i_*(\mathcal{M}^h) \to i_*(\mathcal{M}^h)$ given by the module action. Thus, we have a morphism

$$(i_*\mathcal{M})^h \to i_*(\mathcal{M}^h).$$

To complete the argument, we need to show that this is an isomorphism on each stalk. For this, we only need to consider stalks at points $x \in Y$, since the stalks of both sheaves are zero off Y. At such a point x, this amounts to showing that the natural map

$${}_X\mathcal{H}_x \otimes_{{}_X\mathcal{O}'_x} \mathcal{M}_x \to {}_Y\mathcal{H}_x \otimes_{{}_Y\mathcal{O}'_x} \mathcal{M}_x$$

is an isomorphism. The latter sheaf can be written as

$$({}_X\mathcal{H}_x \otimes_{{}_X\mathcal{O}'_x} {}_Y\mathcal{O}'_x) \otimes_{{}_X\mathcal{O}'_x} \mathcal{M}_x$$

by using the fact that ${}_X\mathcal{H}_x \otimes_{{}_X\mathcal{O}'_x} {}_Y\mathcal{O}'_x \simeq {}_Y\mathcal{H}_x$, which is deduced from tensoring ${}_X\mathcal{H}_x$ with the exact sequence $0 \to \mathcal{I}_Y \to {}_X\mathcal{O} \to {}_Y\mathcal{O} \to 0$ and using flatness and Theorem 14.4(c). The theorem then follows from the associativity of tensor product and the fact that ${}_Y\mathcal{O}'_x \otimes_{{}_X\mathcal{O}'_x} \mathcal{M}_x \simeq \mathcal{M}_x$.

We are now in a position to prove the three main theorems of Serre's GAGA paper.

Note that if X is an algebraic variety and $\mathcal{M}$ is a sheaf on X then $\mathcal{M}$ and $\mathcal{M}'$ have the same global sections and, in fact, the same cohomology. The latter is due to the fact that a flabby resolution of $\mathcal{M}'$ will also be a flabby resolution of $\mathcal{M}$ when restricted to the Zariski open sets. Thus, $H^p(X, \mathcal{M}) = H^p(X^h, \mathcal{M}')$. Furthermore, the map $m \to 1 \otimes m : \mathcal{M}' \to \mathcal{M}^h$ induces a morphism $H^p(X^h, \mathcal{M}') \to H^p(X^h, \mathcal{M}^h)$. Combining these facts gives us a morphism $H^p(X, \mathcal{M}) \to H^p(X^h, \mathcal{M}^h)$. The following is the first of Serre's three GAGA theorems:

14.7 Theorem. *If X is a projective algebraic variety and $\mathcal{M}$ is a coherent algebraic sheaf on X, then the natural map $H^p(X, \mathcal{M}) \to H^p(X^h, \mathcal{M}^h)$ is an isomorphism for every p.*

Proof. Since X is projective, we may assume that it is embedded as a subvariety of P^n. If $i : X \to P^n$ is the embedding, then $i_*\mathcal{M}$ is a coherent algebraic sheaf on P^n by Theorem 14.5. Furthermore, the cohomology groups of $\mathcal{M}$ and $i_*\mathcal{M}$ are the same (Problem 14.2). Likewise, $i_*\mathcal{M}^h$ is a coherent analytic sheaf on P^n by Theorem 14.5. and the cohomology groups of $\mathcal{M}^h$ and $i_*\mathcal{M}^h$ are the same. By Theorem 14.6, $(i_*\mathcal{M})^h = i_*(\mathcal{M}^h)$. It follows that the theorem is true in general if it is true for coherent algebraic sheaves on P^n. Thus, we may assume that $X = P^n$.

By Theorem 13.17, $\mathcal{M}$ is the quotient of a sheaf $\mathcal{F}$ which is a finite direct sum of sheaves of the form $\mathcal{O}(k)$. If $\mathcal{K}$ is the kernel of $\mathcal{F} \to \mathcal{M}$, then we have a short exact sequence of coherent algebraic sheaves

$$0 \longrightarrow \mathcal{K} \longrightarrow \mathcal{F} \longrightarrow \mathcal{M} \longrightarrow 0$$

If we apply $(\ )^h$ to this sequence and use Theorems 14.3 and 14.4 then we have an exact sequence

$$0 \longrightarrow \mathcal{K}^h \longrightarrow \mathcal{F}^h \longrightarrow \mathcal{M}^h \longrightarrow 0$$

of coherent analytic sheaves with $\mathcal{F}^h$ a finite direct sum of sheaves of the form $\mathcal{H}(k)$. If we apply the morphism of cohomology induced by $(\ )^h$ to the long exact sequences of cohomology corresponding to these short exact sequences, we obtain a commutiative diagram

$$\begin{array}{ccccccccc}
H^p(\mathcal{K}) & \longrightarrow & H^p(\mathcal{F}) & \longrightarrow & H^p(\mathcal{M}) & \longrightarrow & H^{p+1}(\mathcal{K}) & \longrightarrow & H^{p+1}(\mathcal{F}) \\
{\scriptstyle \epsilon_1}\big\downarrow & & {\scriptstyle \epsilon_2}\big\downarrow & & {\scriptstyle \epsilon_3}\big\downarrow & & {\scriptstyle \epsilon_4}\big\downarrow & & {\scriptstyle \epsilon_5}\big\downarrow \\
H^p(\mathcal{K}^h) & \longrightarrow & H^p(\mathcal{F}^h) & \longrightarrow & H^p(\mathcal{M}^h) & \longrightarrow & H^{p+1}(\mathcal{K}^h) & \longrightarrow & H^{p+1}(\mathcal{F}^h)
\end{array}$$

where, to save space, we have suppressed the P^n in each cohomology group and written, for example, $H^p(\mathcal{M})$ rather than $H^p(P^n, \mathcal{M})$. In this diagram, ϵ_2 and ϵ_5 are always isomorphisms by Corollary 13.14 and the fact that $\mathcal{F}$ is a direct sum of sheaves $\mathcal{O}(k)$. Suppose that $H^{p+1}(\mathcal{M}) \to H^{p+1}(\mathcal{M}^h)$ is a isomorphism for every coherent algebraic sheaf $\mathcal{M}$. Then ϵ_4 is an isomrphism. This implies that ϵ_3 is surjective for every $\mathcal{M}$ and, hence, that ϵ_1 is surjective. This, along with the fact that ϵ_2 and ϵ_4 are isomorphisms, implies that ϵ_3 is also injective and, thus, is an isomorphim. Thus, the proof will be complete if we can show that there is a p_0 so that $H^p(\mathcal{M}) \to H^p(\mathcal{M}^h)$ is an isomorphism for all $\mathcal{M}$ for $p > p_0$. However, $H^p(\mathcal{M}) = 0$ for $p > n$ because it can be computed using Čech cohomology for the Leray cover $\{U_i\}$ for $\mathcal{M}$. Using Cartan's Theorem B from the next Chapter one can show that $\{U_i\}$ is also a Leray cover for $\mathcal{M}^h$, but there is a more elementary way to see that $H^p(\mathcal{M}^h) = 0$ for large p. We know that P^n is a $2n$ dimensional topological manifold, from which it follows that every open cover of P^n has a refinement in which no $2n+2$ sets have non-empty intersection. Thus, $H^p(\mathcal{M}^h) = 0$ for $p > 2n$ and $H^p(\mathcal{M}) \to H^p(\mathcal{M}^h)$ is trivially an isomorphism for $p > 2n$. This completes the proof.

The next theorem is the second of Serre's GAGA theorems. To prove it, we first need a lemma. If A is a subring of a ring B and E and F are A-modules, then the functor $(\ ) \otimes_A B$ defines a natural morphism of A-modules

$$\hom_A(E, F) \to \hom_B(E \otimes_A B, F \otimes_A B)$$

which, since the right side is a B-module, induces a B-module morphism

$$\iota : \hom_A(E, F) \otimes_A B \to \hom_B(E \otimes_A B, F \otimes_A B).$$

14.8 Lemma. *The morphism ι, defined above, is an isomorphism if A is Noetherian, E is finitely generated over A and B is faithfully flat over A.*

Proof. For a fixed module F, consider the two functors T and T' from A-modules to B-modules defined by

$$T(E) = \hom_A(E, F) \otimes_A B$$

and

$$T'(E) = \quad \hom_B(E \otimes_A B, F \otimes_A B) = \hom_A(E, F \otimes_A B).$$

Then $\iota : T \to T'$ is a transformation of functors and we are to show that it is an isomorphism on finitely generated modules. Clearly ι is an isomorphism if $E = A$, since $T(A)$ and $T'(A)$ are both equal to $F \otimes_A B$ in this case. Similarly ι is an isomorphism if $E = A^n$ for some n. Note that T and T' are left exact since hom is left exact in its first variable and B is faithfully flat over A. Since A is Noetherian, for each finitely generated A-module E we can construct an exact sequence of the form

$$A^n \longrightarrow A^m \longrightarrow E \longrightarrow 0$$

On applying T and T', we obtain a diagram

$$\begin{array}{ccccccc}
0 & \longrightarrow & T(E) & \longrightarrow & T(A^m) & \longrightarrow & T(A^n) \\
 & & \iota\big\downarrow & & \iota\big\downarrow & & \iota\big\downarrow \\
0 & \longrightarrow & T'(E) & \longrightarrow & T'(A^m) & \longrightarrow & T'(A^n)
\end{array}$$

with exact rows and with the last two vertical maps isomorphisms. It follows that the first vertical map is also an isomorphism and the proof is complete.

14.9 Theorem. *If $\mathcal{M}$ and $\mathcal{N}$ are two coherent algebraic sheaves on a projective algebraic variety X, then every morphism of analytic sheaves $\mathcal{M}^h \to \mathcal{N}^h$ is induced by a morphism $\mathcal{M} \to \mathcal{N}$ of algebraic sheaves.*

Proof. Let $\mathcal{A}$ denote the sheaf $hom(\mathcal{M}, \mathcal{N})$. This is the sheaf which assigns to an open $U \subset X$ the $\mathcal{O}(U)$-module consisting of morphisms $\mathcal{M}|_U \to \mathcal{N}|_U$ in the category of sheaves of ${}_U\mathcal{O}$-modules. Similarly, let $\mathcal{B} = hom(\mathcal{M}^h, \mathcal{N}^h)$ be the analogous sheaf for the sheaves of $\mathcal{H}$ modules $\mathcal{M}^h$ and $\mathcal{N}^h$. The functor $(\ )^h$ clearly defines a sheaf morphism $\mathcal{A}' \to \mathcal{B}$ and, since $\mathcal{B}$ is an $\mathcal{H}$-module, this induces a morphism $\mathcal{A}^h = \mathcal{H} \otimes \mathcal{A}' \to \mathcal{B}$, where, in this argument, $\otimes$ will mean tensor product relative to $\mathcal{O}$. We claim that this morphism $\mathcal{A}^h \to \mathcal{B}$ is an isomorphism. As usual, it suffices to check this for the stalks at each point of X. The fact that $\mathcal{M}$ is coherent and, hence, locally finitely generated implies that each $\mathcal{O}_x$-module homomorphism from $\mathcal{M}_x$ to $\mathcal{N}_x$ extends to a morphism from $\mathcal{M}|_U$ to $\mathcal{N}|_U$ for some neighborhood U of x and that a morphism from $\mathcal{M}|_U$ to $\mathcal{N}|_U$ which vanishes at x also vanishes in a neighborhood of x. These two statements, taken together, and their analogues for coherent analytic sheaves mean that

$$\mathcal{A}_x = \hom(\mathcal{M}_x, \mathcal{N}_x) \quad \text{and} \quad \mathcal{B}_x = \hom(\mathcal{M}^h_x, N^h_x)$$

We also have that

$$\mathcal{A}^h_x = \hom(\mathcal{M}_x, \mathcal{N}_x) \otimes \mathcal{H}_x, \quad M^h_x = M_x \otimes \mathcal{H}_x \quad \text{and} \qquad N^h_x = N_x \otimes \mathcal{H}_x$$

Thus, our claim will be established if we can show that the natural homomorphism

$$\hom(\mathcal{M}_x, \mathcal{N}_x) \otimes \mathcal{H}_x \to \hom(M_x \otimes \mathcal{H}_x, N_x \otimes \mathcal{H}_x)$$

is bijective. But this follows from the previous lemma since $\mathcal{O}_x$ is Noetherian, M_x is finitely generated and $\mathcal{H}_x$ is faithfully flat over $\mathcal{O}_x$ by Corollary 7.15.

To finish the proof, we consider the morphisms

$$H^0(X,\mathcal{A}) \to H^0(X^h,\mathcal{A}^h) \to H^0(X,\mathcal{B}).$$

The first of these morphisms is an isomorphism by Theorem 14.7 provided we can show that $\mathcal{A}$ is coherent. This is done in Problem 14.3. The second morphism is an isomorphism by the claim proved in the previous paragraph. Thus, the composition $H^0(X,\mathcal{A}) \to H^0(X,\mathcal{B})$ is an isomorphism. This completes the proof of the theorem since a global section of $\mathcal{B}$ is a morphism $\mathcal{M}^h \to \mathcal{N}^h$ while a global section of $\mathcal{A}$ is a morphism $\mathcal{M} \to \mathcal{N}$.

The *geometric fiber* of a coherent analytic sheaf $\mathcal{S}$ at a point x is the $\mathcal{H}_x$ module $\mathcal{S}_x/M_x\mathcal{S}_x$, where M_x is the maximal ideal of H_x. The geometric fiber of a coherent algebraic sheaf is defined analogously.

14.10 Lemma. *If $\mathcal{S}$ is a coherent analytic sheaf on an holomorphic variety X, $x \in X$ and $F \subset H^0(X,\mathcal{S})$ is a set of sections which generates the geometric fiber of $\mathcal{S}$ at x, then F generates $\mathcal{S}|_U$ for some neighborhood U of x. The analogous statement is true for coherent algebraic sheaves.*

Proof. It follows from Nakayama's Lemma that if F generates $\mathcal{S}_x/M_x\mathcal{S}_x$ then it generates $\mathcal{S}_x$ (Problem 14.5). However, by Problem 12.4 the set of y at which F fails to generate $\mathcal{S}_y$ is a closed subvariety. Hence, there is a neighborhood U of x such that F generates $\mathcal{S}|_U$.

The third and most difficult of Serre's GAGA theorems is the following:

14.11 Theorem. *If X is a projective algebraic variety and $\mathcal{M}$ is a coherent analytic sheaf on X^h, then there is a coherent algebraic sheaf $\mathcal{N}$ on X such that $\mathcal{N}^h \simeq \mathcal{M}$. Furthermore, $\mathcal{N}$ is unique up to isomorphism.*

Proof. The uniqueness is an immediate consequence of the preceding theorem.

Claim 1. *The theorem is true if it is true for $X = P^n$ for all n.*

If $i : X \to P^n$ is an embedding of X as a subvariety of P^n, then $i_*\mathcal{M}$ is a coherent analytic sheaf on P^n. Suppose there is a coherent algebraic sheaf $\mathcal{S}$ on P^n with $\mathcal{S}^h \simeq i_*\mathcal{M}$. Then we claim that $\mathcal{S}$ is $i_*\mathcal{N}$ for a coherent algebraic sheaf $\mathcal{N}$ on X. In fact, if $\mathcal{I}$ is the ideal sheaf of X and $f \in \mathcal{I}_x$ for some $x \in P^n$, then multiplication by f determines an endomorphism $\phi : \mathcal{S}|_U \to \mathcal{S}|_U$ for some nieghborhood U of x. The corresponding morphism $\phi^h : \mathcal{S}^h|_U \to \mathcal{S}^h|_U$ is still multiplication by f and is zero in a neighborhood of x because $\mathcal{S}^h \simeq i_*\mathcal{M}$ and $\mathcal{M}$ is a sheaf of ${}_X\mathcal{H}$-modules and the $\mathcal{H}$-module action on $i_*\mathcal{M}$ factors through the quotient map $\mathcal{H} \to {}_X\mathcal{H}$. But if ϕ^h vanishes in a neighborhood then so does ϕ by Theorem 14.3. Thus, we have proved that $\mathcal{I}\mathcal{S} = 0$. This implies that $\mathcal{S}$ is supported on X and its restriction to X is a sheaf $\mathcal{N}$ of ${}_X\mathcal{O}$-modules. It is easy to see that $\mathcal{N}$ is a coherent algebraic sheaf on X (Problem 14.4) and, obviously, $i_*\mathcal{N} = \mathcal{S}$. Now Theorem 14.6 implies that $i_*\mathcal{N}^h \simeq (i_*\mathcal{N})^h = \mathcal{S}^h \simeq i_*\mathcal{M}$. On restricting to X, this implies that $\mathcal{N}^h \simeq \mathcal{M}$. Thus, the theorem is true for any projective variety if it is true for P^n.

We have reduced the proof to the case where $X = P^n$. We will now prove it in this case by induction on n. It is trivial when $n = 0$ since P^0 is a point and coherent algebraic and analytic sheaves are just finite dimensional vector spaces. Thus, we will assume that $n > 0$ and that the theorem is true in dimensions less than n.

Note that for a coherent analytic sheaf $\mathcal{M}$ on P^n we may *twist by* $\mathcal{H}(k)$ to construct a coherent analytic sheaf $\mathcal{M}(k) = \mathcal{M} \otimes_{\mathcal{H}} \mathcal{H}(k)$ for each k just as we did for coherent algebraic sheaves. Note also that $\mathcal{N}(k)^h \simeq (\mathcal{N}^h)(k)$ due to the fact that $\mathcal{O}(k)^h = \mathcal{H}(k)$.

Claim 2. *Let E be a hyperplane in P^n and $\mathcal{A}$ a coherent analytic sheaf on E. Then, under our induction assumption, $H^p(E, \mathcal{A}(k)) = 0$ for large enough k.*

We have that E is the subvariety of P^n defined by the zero set of a linear functional on $\mathbb{C}^{n+1}$ (a homogeneous polynomial of degree one). Then E is a copy of P^{n-1} and so, by assumption, the theorem is true for $X = E$. That $H^p(E, \mathcal{A}(k)) = 0$ for a coherent analytic sheaf $\mathcal{A}$ on E and large enough k then follows from Theorems 14.7 and 13.18 and our induction assumption which implies that $\mathcal{A} = \mathcal{B}^h$ for some coherent algebraic sheaf B.

The key to the proof of the Theorem is the next claim:

Claim 3. *Under our induction assumption, for every coherent analytic sheaf $\mathcal{M}$ on X there is an integer $k_{\mathcal{M}}$ such that for every $k > k_{\mathcal{M}}$ the sheaf $\mathcal{M}(k)$ is generated over $\mathcal{H}$ by its space of global sections $H^0(P^n, \mathcal{M}(k))$.*

Note that if $\mathcal{M}(k)_x$ is generated by $H^0(P^n, \mathcal{M}(k))$ then $\mathcal{M}(k+p)_x$ is generated by $H^0(P^n, \mathcal{M}(k+p))$ for all $p > 0$. This is due to the fact that $\mathcal{M}(k+p)_x = \mathcal{M}(k)_x \otimes_{\mathcal{H}_x} \mathcal{H}(p)_x$ and $\mathcal{H}(p)_x$ is generated by its global sections if $p \geq 0$ (since $\mathcal{H}(p)$ is the sheaf of holomorphic sections of a line bundle, one only needs to have one global section which is non-vanishing at x in order to have the global sections generate $\mathcal{H}(p)_x$ and the existence of such a section follows from Theorem 13.12). In view of these remarks, the compactness of P^n, and Lemma 14.10, to prove Claim 3 it suffices to prove that for each $x \in P^n$ there is a k for which the geometric fiber of the module $\mathcal{M}(k)_x$ is generated by $H^0(P^n, \mathcal{M}(k))$.

Now let E be a hyperplane in P^n and let $\mathcal{I}_E$ be its ideal sheaf in $\mathcal{H}$. We may as well assume that E is the hyperplane on which the coordinate function z_0 vanishes. Consider the exact sequence

$$0 \longrightarrow \mathcal{I}_E \longrightarrow \mathcal{H} \longrightarrow {}_E\mathcal{H} \longrightarrow 0$$

Now z_0 may be regarded as a section of $\mathcal{H}(1)$ and so multiplication by z_0 defines a morphism from $\mathcal{H}(-1)$ to $\mathcal{H}$. The image of this morphism is the ideal sheaf $\mathcal{I}_E$. Thus, we have an isomorphism $\mathcal{H}(-1) \to \mathcal{I}_E$. Thus, the above exact sequence becomes

$$0 \longrightarrow \mathcal{H}(-1) \longrightarrow \mathcal{H} \longrightarrow {}_E\mathcal{H} \longrightarrow 0$$

Which, on tensoring with $\mathcal{M}$ relative to $\mathcal{H}$, yields an exact sequence

$$0 \longrightarrow \mathcal{C} \longrightarrow \mathcal{M}(-1) \longrightarrow \mathcal{M} \longrightarrow \mathcal{B} \longrightarrow 0$$

Where $\mathcal{B} = {}_E\mathcal{H} \otimes_{\mathcal{H}} \mathcal{M}$ and $\mathcal{C} = \operatorname{tor}_1^{\mathcal{H}}({}_E\mathcal{H}, \mathcal{M})$. If we tensor this with $\mathcal{H}(k)$, we get the sequence of coherent analytic sheaves

$$0 \longrightarrow \mathcal{C}(k) \longrightarrow \mathcal{M}(k-1) \longrightarrow \mathcal{M}(k) \longrightarrow \mathcal{B}(k) \longrightarrow 0.$$

If we set $\mathcal{L}_k = \ker(\mathcal{M}(k) \to \mathcal{B}(k))$, then this sequence breaks up into two short exact sequences of coherent analytic sheaves

$$0 \longrightarrow \mathcal{C}(k) \longrightarrow \mathcal{M}(k-1) \longrightarrow \mathcal{L}_k \longrightarrow 0$$

and

$$0 \longrightarrow \mathcal{L}_k \longrightarrow \mathcal{M}(k) \longrightarrow \mathcal{B}(k) \longrightarrow 0.$$

Now we can apply the long exact sequences of cohomology to these two short exact sequences. The relevant parts for us are

$$H^1(P^n, \mathcal{M}(k-1)) \longrightarrow H^1(P^n, \mathcal{L}_k) \longrightarrow H^2(P^n, \mathcal{C}(k))$$

and

$$H^1(P^n, \mathcal{L}_k) \longrightarrow H^1(P^n, \mathcal{M}(k)) \longrightarrow H^1(P^n, \mathcal{B}(k))$$

Now $\mathcal{B} = {}_E\mathcal{H} \otimes_{\mathcal{H}} \mathcal{M}$ and $\mathcal{C} = \operatorname{tor}_1^{\mathcal{H}}({}_E\mathcal{H}, \mathcal{M})$ are coherent sheaves of $\mathcal{H}$-modules but the action of $\mathcal{H}$ factors through ${}_E\mathcal{H}$ and, hence, they are actually coherent analytic sheaves on E (Problem 14.4). It follows that $\mathcal{B}(k)$ and $\mathcal{C}(k)$ are also coherent analytic sheaves on E. By our induction assumption this means that they are images under the functor $(\ )^h$ of coherent algebraic sheaves on E. It then follows from Theorem 14.7 and Therem 13.18 that $\mathcal{B}(k)$ and $\mathcal{C}(k)$ have vanishing pth cohomology for $p > 0$ and k sufficiently large. Thus, the above sequences imply that for large k we have surjective maps

$$H^1(P^n, \mathcal{M}(k-1)) \to H^1(P^n, \mathcal{L}_k) \quad \text{and} \quad H^1(P^n, \mathcal{L}_k) \to H^1(P^n, \mathcal{M}(k))$$

which implies that

$$\dim H^1(P^n, \mathcal{M}(k-1)) \geq \dim H^1(P^n, \mathcal{L}_k) \geq \dim H^1(P^n, \mathcal{M}(k)).$$

At this point we must appeal to the results of Chapter 16 – specifically, to the Cartan-Serre Theorem which states that all cohomology spaces of a coherent analytic sheaf on a compact holomorphic variety are finite dimensional vector spaces. Thus, the spaces $H^1(P^n, \mathcal{M}(k))$ are finite dimensional and, by the above, the dimension is a non-increasing function of k for sufficiently large k. This implies that eventually the dimension must become constant as k increases. Thus, for sufficiently large k we have that $H^1(P^n, \mathcal{L}_k) \to H^1(P^n, \mathcal{M}(k))$ is a surjective map between finite dimensional vector spaces of the same dimension. This implies that it is injective as well. If we use this fact on the long exact sequence of cohomology

$$H^0(P^n, \mathcal{M}(k)) \longrightarrow H^0(P^n, \mathcal{B}(k)) \longrightarrow H^1(P^n, \mathcal{L}_k) \longrightarrow H^1(P^n, \mathcal{M}(k))$$

we conclude that $H^0(P^n, \mathcal{M}(k)) \to H^0(P^n, \mathcal{B}(k))$ is surjective for k sufficiently large. We also know, by the induction hypothesis, that the coherent sheaf $\mathcal{B}$ on E is the image under $(\ )^h$ of a coherent algebraic sheaf on E and, hence, that $\mathcal{B}(k)_x$ is generated by $H^0(P^n, \mathcal{B}(k))$ if k is sufficiently large (since the analogous thing is true of coherent algebraic sheaves by Theorem 13.16). We conclude that $H^0(P^n, \mathcal{M}(k))$ generates $\mathcal{B}(k)_x = \mathcal{M}(k)_x \otimes_{\mathcal{H}_x} {}_E\mathcal{H}_x$.

The latter module is just the quotient module $\mathcal{M}(k)_x/(I_E)_x\mathcal{M}(k)_x$. Since $H^0(P^n, \mathcal{M}(k))$ generates this quotient module of $\mathcal{M}(k)_x$, it generates the geometric fiber of $\mathcal{M}(k)_x$. This completes the proof of Claim 3 and puts us in a position to complete the proof of the Theorem.

By Claim 3 we know that if $\mathcal{M}$ is a coherent analytic sheaf then there is an integer k such that $\mathcal{M}(k)$ is generated by $H^0(P^n, \mathcal{M}(k))$. Since $H^0(P^n, \mathcal{M}(k))$ is finite dimensional by Theorem 16.18, there is a surjection $\mathcal{H}^p \to \mathcal{M}(k)$ for some p. If we twist this morphism by the sheaf $\mathcal{H}(-k)$, we obtain a surjection $\mathcal{H}^p(-k) \to \mathcal{M}$. Now by applying the same analysis to the kernel of this map, which is also a coherent analytic sheaf, we obtain an exact sequence of the form

$$\mathcal{H}^q(-j) \xrightarrow{\alpha} \mathcal{H}^p(-k) \longrightarrow \mathcal{M} \longrightarrow 0$$

Now $\mathcal{H}^q(-j) = \mathcal{O}^q(-j)^h$ and $\mathcal{H}^p(-k) = \mathcal{O}^p(-k)^h$ and so, by Theorem 14.9, the morphism α is induced by a morphism of coherent algebraic sheaves $\beta : \mathcal{O}^q(-j) \to \mathcal{O}^p(-k)$. If $\mathcal{N}$ is the cokernel of β, then the exact functor $(\ )^h$ applied to the exact sequence

$$\mathcal{O}^q(-j) \xrightarrow{\beta} \mathcal{O}^p(-k) \longrightarrow \mathcal{N} \longrightarrow 0$$

yields an exact sequence

$$\mathcal{H}^q(-j) \xrightarrow{\alpha} \mathcal{H}^p(-k) \longrightarrow \mathcal{N}^h \longrightarrow 0.$$

But this implies that $\mathcal{M} \simeq \mathcal{N}^h$ which completes the induction and the proof of the theorem.

The results of Theorems 14.7, 14.9 and 14.11 (Serre's Theorems 1, 2, and 3) can be summarized as follows:

14.12 Theorem. *If X is a projective algebraic variety, then the functor $\mathcal{M} \to \mathcal{M}^h$ is a cohomology preserving equivalence of categories from the category of coherent algebraic sheaves on X to the category of coherent analytic sheaves on X^h.*

The above results have a wide variety of applications. We state some of these below, but prove only a couple of them. The proofs of the others require knowledge of results from algebraic geometry which we have not developed here. For a more complete discussion of applications we refer the reader to Serre's paper.

At this point, we will drop the use of the X^h notation except in situations in which it is needed to avoid confusion. If X is an algebraic variety then will generally also use X to denote the corresponding holomorphic variety – that is, we will think of an algebraic variety as having two ringed space structures – one algebraic and one holomorphic. We will call a holomorphic variety X *algebraic* if it is the holomorphic variety associated to some algebraic variety.

Our first application is the following theorem of Chow:

14.13 Corollary. *If X is a projective variety, then every holomorphic subvariety of X is algebraic.*

Proof. Let Y be an analytic subvariety of X and consider the ideal sheaf $\mathcal{I}_Y$. Then the quotient $\mathcal{H}/\mathcal{I}_Y$ is a coherent analytic sheaf on X which is isomorphic to $i_{*Y}\mathcal{H}$, where

$i : Y \to X$ is the inclusion. The support of the sheaf $\mathcal{H}/\mathcal{I}_Y$ is clearly the subvariety Y. Now by Theorem 14.12, there is a coherent algebraic sheaf $\mathcal{N}$ with the property that $\mathcal{N}^h \simeq \mathcal{H}/\mathcal{I}_Y$. The support of $\mathcal{N}$ is the same pointset as the support of $\mathcal{N}^h$ due to the fact that $\mathcal{H}_x$ is faithfully flat over $\mathcal{O}_x$ for each x. Thus, the support of $\mathcal{N}$ is Y. However, the support of a coherent algebraic sheaf is an algebraic subvariety (see Problem 12.4 and note that the proof works equally well in the algebraic case). This completes the proof.

One can combine this with another result of Chow (on representing a general algebraic variety as the image under a regular map of a dense open subset of a projective variety) to obtain (cf. Serre):

14.13 Corollary. *If X is an algebraic variety, then every compact holomorphic subvariety of X^h is algebraic.*

With this and a little additional work, one can prove (cf. Serre):

14.14 Corollary. *Every holomorphic map from a compact algebraic variety to an algebraic variety is regular.*

This has the obvious consequence that:

14.15 Corollary. *A compact holomorphic variety has at most one structure of an algebraic variety (up to isomorphism).*

The category of algebraic vector bundles on an algebraic variety may be identified with the category of locally free finite rank sheaves of $\mathcal{O}$ modules – a vector bundle is identified with its sheaf of sections. In the same way, the category of holomorphic vector bundles may be identified with the category of locally free finite rank sheaves of $\mathcal{H}$-modules. Clearly, the equivalence of categories $\mathcal{M} \to \mathcal{M}^h$ of Theorem 14.12 has the property that $\mathcal{M}$ is free of finite rank if and only if $\mathcal{M}^h$ is free of finite rank. The corresponding functor on vector bundles is just the functor which assigns to an algebraic vector bundle $\pi : E \to X$ over an algebraic variety the holomorphic bundle $\pi^h : E^h \to X^h$ obtained by putting the canonical analytic structure on both total space and base. Thus, we have proved:

14.16 Corollary. *If X is a projective algebraic variety then the category of algebraic vector bundles on X is equivalent to the category of holomorphic vector bundles on X^h under the natural correspondence.*

Serre's paper contains a more general result of the above type which concerns bundles with structure groups other than $Gl_n(\mathbb{C})$. It also contains, as an entirely different kind of application of the GAGA theorems, a proof of the following conjecture of A. Weil:

14.17 Corollary. *If V is a projective non-singular variety defined over an algebraic number field K. Then the complex projective variety determined by an embedding of K in $\mathbb{C}$ has Betti numbers which are independent of the embedding that is chosen.*

14. Problems

1. Prove that if $\mathcal{M}$ and $\mathcal{N}$ are sheaves of modules over a sheaf of rings $\mathcal{R}$, then

$$(\mathcal{M} \otimes_{\mathcal{R}} \mathcal{N})_x = \varinjlim \{\mathcal{M}(U) \otimes_{\mathcal{R}(U)} \mathcal{N}(U) : x \in U\} = \mathcal{M}_x \otimes_{\mathcal{R}_x} \mathcal{N}_x$$

   where the first equality is the definition of $(\mathcal{M} \otimes_{\mathcal{R}} \mathcal{N})_x$.
2. Prove that if $i : Y \to X$ is an embedding of Y as a closed subspace of a topological space X and if $\mathcal{S}$ is any sheaf on Y, then $H^p(X, i_*\mathcal{S}) \simeq H^p(Y, \mathcal{S})$. Hint: Use a flabby resolution of $\mathcal{S}$ and take i^* of it.
3. Prove that if $\mathcal{M}$ and $\mathcal{N}$ are coherent algebraic (analytic) sheaves on an algebraic (holomorphic) variety X, then the sheaf $hom(\mathcal{M}, \mathcal{N})$ is also a coherent algebraic (analytic) sheaf. Hint: For coherent algebraic sheaves prove this first when M is $\mathcal{O}^m$ and then represent a general M as the cokernel of a morphism $\mathcal{O}^k \to \mathcal{O}^m$. The proof is analogous for coherent analytic sheaves.
4. Suppose Y is a subvariety of X, $\mathcal{I}$ is the ideal sheaf of Y and $\mathcal{S}$ is a coherent sheaf on X which satisfies $\mathcal{I}\mathcal{S} = 0$. Prove that $\mathcal{S}$ is supported on Y and its restriction to Y is a coherent sheaf on Y. It doesn't matter whether the sheaves and spaces are algebraic or analytic.
5. Use Nakayama's Lemma to prove that if S is a finitely generated module over a Noetherian local ring A, M is the maximal ideal of A and F is a subset of S whose image in S/MS generates S/MS, then F generates S.

15. Stein Spaces

Stein spaces play the role in the study of holomorphic varieties and functions that is played by the affine varieties in the algebraic theory. The most important theorems in the subject of several complex variables are Cartan's Theorems A and B. Cartan's Theorem A, says that a coherent analytic sheaf on a Stein space is generated over $\mathcal{H}$ by its global sections and Cartan's Theorem B which says that a coherent analytic sheaf on a Stein space has vanishing cohomology in degrees greater than zero. In this section we define the category of Stein spaces and lay the groundwork for proving Cartan's theorems. The proofs of the theorems themselves will be carried out in the next chapter, where we develop and employ some approximation results for coherent analytic sheaves.

The key result in this section is a vanishing theorem which states that a coherent analytic sheaf defined in a neighborhood of a compact polydisc has vanishing higher cohomology on the polydisc. The strategy of the proof is to construct a free resolution for each such sheaf like the one in the following theorem – that is, a terminating chain of syzygys.

15.1 Lemma. *Let Δ be an open polydisc in $\mathbb{C}^n$ and suppose there is an exact sequence of sheaves of $\mathcal{H}$-modules on Δ of the form*

$$0 \longrightarrow \mathcal{H}^{p_m} \longrightarrow \cdots \longrightarrow \mathcal{H}^{p_1} \longrightarrow \mathcal{H}^{p_0} \longrightarrow \mathcal{S} \longrightarrow 0$$

Then $H^p(\Delta, \mathcal{S}) = 0$ for $p > 0$ and the sequence of global sections

$$0 \longrightarrow \Gamma(\Delta, \mathcal{H}^{p_m}) \longrightarrow \cdots \longrightarrow \Gamma(\Delta, \mathcal{H}^{p_0}) \longrightarrow \Gamma(\Delta, \mathcal{S}) \longrightarrow 0$$

is also exact.

Proof. The exact sequence of sheaves in the hypothesis can be decomposed into a collection of short exact sequences of coherent analytic sheaves

$$0 \longrightarrow \mathcal{L}_k \longrightarrow \mathcal{H}^{p_k} \longrightarrow \mathcal{L}_{k-1} \longrightarrow 0$$

where $\mathcal{L}_k = \ker\{\mathcal{H}^{p_k} \to \mathcal{H}^{p_{k-1}}\}$ for $k > 0$, $\mathcal{L}_m = 0$ and $\mathcal{L}_{-1} = \mathcal{S}$. For such a sequence, the long exact sequence of cohomology and the fact that $\mathcal{H}^{p_k}$ has vanishing qth cohomology for $q > 0$ imply that if $\mathcal{L}_k$ is also a sheaf with vanishing qth cohomology for all $q > 0$ then $\mathcal{L}_{k-1}$ is as well and the sequence

$$0 \longrightarrow \Gamma(\Delta, \mathcal{L}_k) \longrightarrow \Gamma(\Delta, \mathcal{H}^{p^k}) \longrightarrow \Gamma(\Delta, \mathcal{L}_{k-1}) \longrightarrow 0$$

is exact. The theorem follows from descending induction using this fact, beginning on the left at $k = m$.

In a corollary to Hilbert's Syzygy Theorem (Corollary 3.14) we proved that every finitely generated module over the local ring ${}_n\mathcal{H}$ has a terminating free finite rank resolution – a terminating syzygy. Thus, given a coherent sheaf $\mathcal{S}$ on an open set in $\mathbb{C}^n$, a sequence like the one in the above lemma can always be constructed at each point. Using the results on coherence of Chapter 12, we are able to construct such a sequence in a neighborhood of any point:

15.2 Lemma. *If $\mathcal{S}$ is a coherent analytic sheaf defined on an open set U in $\mathbb{C}^n$, then for each point $x \in U$ there is a neighborhood W of x in U and an exact sequence:*

$$0 \longrightarrow \mathcal{H}^{p_n}|_W \xrightarrow{\phi_n} \dots \longrightarrow \mathcal{H}^{p_1}|_W \xrightarrow{\phi_1} \mathcal{H}^{p_0}|_W \xrightarrow{\phi_0} \mathcal{S}|_W \longrightarrow 0$$

In other words, $\mathcal{S}$ locally has a free finite rank resolution of length n.

Proof. Let x be any point of U. By the definition of analytic coherence and the results of Chapter 12, there is a neighborhood W of x and a surjection $\phi_0 : \mathcal{H}^{p_0}|_W \to \mathcal{S}|_W \to 0$ such that $\ker \phi$ is also coherent. Thus, by shrinking W if necessary, we may express $\ker \phi$ as the image of a morpism $\phi_1 : \mathcal{H}^{p_1}|_W \to \mathcal{H}^{p_0}|_W$. Continuing in this manner, we construct the sequence as above up to stage $n-1$, where we have a morphism $\phi_{n-1} : \mathcal{H}^{p_{n-1}}|_W \to \mathcal{H}^{p_{n-2}}|_W$. By Hilbert's Syzygy Theorem (Theorem 3.13), the stalk $(\ker \phi_{n-1})_x$ is a free finite rank $\mathcal{H}_x$-module. We choose germs which form a basis for $(\ker \phi_{n-1})_x$ and then choose representatives in a neighborhood (which we may as well assume is W) for these germs. The resulting finite set of sections defines a morphism $\phi_n : \mathcal{H}^{p_n}|_W \to \ker \phi_{n-1}$ which is an isomorphism at x. By Problem 12.4, ϕ is an isomorphism in a neighborhood of x. Therefore, after again shrinking W if necessary, the morphism ϕ_n completes the construction of our resolution.

The final step in proving a vanishing theorem for the cohomology of coherent sheaves on a polydisc will be to piece together the local resolutions given by the above theorem to obtain such a resolution on the entire polydisc. This is not so easy to do. The construction is based on a factorization lemma (Cartan's Lemma) for holomorphic matrix valued functions. The best proof of this lemma uses the infinite dimensional implicit function theorem. In the next theorem we prove the version of this that we will use. Before we can prove this theorem, we need to establish some preliminary facts about Banach space operators.

Let $A : X \to Y$ be a surjective bounded linear map between two Banach spaces. By the open mapping theorem, A is an open map. This implies that there is a constant K with the property that for each $y \in Y$ there is an $x \in X$ such that $Ax = y$ and $||x|| \leq K||y||$. The least such K is the *inversion constant* for A.

15.3 Lemma. *If $A : X \to Y$ is a surjective continuous linear map between Banach spaces, then there is a $\delta > 0$ and a $K > 0$ such that B is surjective with inversion constant less than or equal to K whenever $B : X \to Y$ is a bounded linear map with $||A - B|| < \delta$.*

Proof. Let K_0 be the inversion constant for A and choose $\delta = (2K_0)^{-1}$. For $B : X \to Y$ with $||A - B|| < \delta$ and $y \in Y$ we seek an x such that $Bx = y$. We choose u_0 such that $Au_0 = y$ and $||u_0|| \leq K_0||y||$. We now choose inductively a sequence $\{u_n\}$ of elements of X such that

$$Au_n = (A - B)u_{n-1}, \quad ||u_n|| \leq 2^{-n}K_0||y||$$

The choice of δ clearly makes this possible. If $x_n = u_0 + u_1 + \cdots + u_n$, then the sequence x_n converges in X and

$$\begin{aligned} y - Bx_n &= Au_0 - Bx_n = Au_1 - B(x_n - u_0) \\ &= Au_2 - B(x_n - u_0 - u_1) = \cdots = Au_{n+1} \end{aligned}$$

and so it is clear that $Bx = y$ if $x = \lim x_n$. It also follows from our estimate on the norms of the u_n that $||x|| \leq 2K_0$. This establishes the Lemma with $K = 2K_0$.

If f is a function from an open subset U of a Banach space X to a Banach space Y, then the *derivative* (if it exists) of f at $x \in U$ is a bounded linear map $f'(x) : X \to Y$ with the property that

$$\lim_{u \to 0} ||u||^{-1} ||f(x+u) - f(x) - f'(x)u|| = 0.$$

15.4 Theorem. *Let X and Y be Banach spaces and let $f : U \to Y$ be a (non-linear) function from a neighborhood of 0 in X into Y. If $f'(x)$ exists at each $x \in U$, is a continuous function of x and is surjective at $x = 0$, then the image of f contains a neighborhood of $f(0)$.*

Proof. By the previous Lemma and the continuity of f', we may assume that U is small enough that $f'(x)$ is surjective with inversion constant bounded by some constant K for all $x \in U$. We may also assume that U is convex so that if x and $x+u$ lie in U then so does the line segment joining them. In this case,

$$f(x+u) - f(x) - f'(x)u = \int_0^1 [f'(x+tu) - f'(x)]u \, dt$$

Thus, by shrinking U if necessary and using the continuity of f' again, we may assume that

$$||f(x+u) - f(x) - f'(x)u|| < (2K)^{-1} ||u||$$

for $x, x+u \in U$.

The remainder of the proof is just an application of Newton's method. We choose $\delta > 0$ so that $||u|| < 2\delta$ implies that $u \in U$. We will show that for $||y - f(0)|| < K^{-1}\delta$ we can solve the equation $f(x) = y$. We proceed as in Newton's method, using $x_0 = 0$ as our initial guess. We then choose $x_1 \in X$ so that $f'(0)x_1 = y - f(0)$ and $||x_1|| \leq K||y - f(0)||$. Note that $||x_1|| < \delta$. We then inductively choose x_n so that $x_n = x_{n-1} + u_n$ where

$$f'(x_{n-1})u_n = y - f(x_{n-1}).$$

and

$$||u_n|| \leq K||y - f(x_{n-1})||.$$

Then we have

$$\begin{aligned} ||y - f(x_n)|| &= ||y - f(x_{n-1}) - (f(x_n) - f(x_{n-1}))|| \\ &= ||f'(x_{n-1})(x_n - x_{n-1}) - (f(x_n) - f(x_{n-1}))|| < (2K)^{-1} ||u_n|| \end{aligned}$$

So that

$$||u_n|| < 2^{-1} ||u_{n-1}||$$

Now $||u_1|| < \delta$ implies that $||u_2|| < 2^{-1}\delta$ and, in general, that $||u_n|| < 2^{-n+1}\delta$. This implies that the sequence $\{x_n\}$ is contained in U and converges to an element $x \in U$. Our estimate above on $||y - f(x_n)||$ shows that $f(x) = y$. This completes the proof.

Roughly speaking, the above theorem says that a certain non-linear problem has a solution if the linearized version is solvable. In our application of this result, the solvability of the linearized problem is given by the next lemma. In this lemma and in what follows we will use the following geometric situation. By an open (compact) *box* in $\mathbb{C}^n$ we will mean an open (compact) set U which is the cartesian product of intervals – one from each of the $2n$ real and imaginary coordinate axes. An aligned pair of open (compact) boxes will be a pair (U_1, U_2) which is the cartesian product of two ordered sets of intervals which are identical except in one (real or imaginary) coordinate and in that coordinate the two intervals are overlapping. The coordinate in which the defining intervals are allowed to be different will be called the exceptional coordinate. It is clear from the definition that if (U_1, U_2) is an aligned pair of boxes, then $U_1 \cap U_2$ and $U_1 \cup U_2$ are also boxes and they are obtained from U_1 and U_2 by taking intersection or union of the two defining intervals in the exceptional coordinate and leaving the defining intervals in all other coordinates the same.

15.5 Lemma. *Let (U_1, U_2) be an aligned pair of open boxes in $\mathbb{C}^n$. Then each bounded holomorphic function f on $U_1 \cap U_2$ is the difference $f_1 - f_2$ of a bounded holomorphic function f_1 on U_1 and a bounded holomorphic function f_2 on U_2.*

Proof. Without loss of generality we may assume that the exceptional coordinate for (U_1, U_2) is x_1, so that the pair (U_1, U_2) has the form

$$U_1 = I_1 \times iK \times W$$

$$U_2 = I_2 \times iK \times W$$

where I_1 and I_2 are overlapping open intervals on the line, K is an open interval on the line and W is an open box in $\mathbb{C}^{n-1}$.

We choose a bounded C^∞ function ϕ on $\mathbb{R}$ that is one in a neighborhood of $I_1 - I_2$ and is zero in a neighborhood of $I_2 - I_1$. We then consider ϕ to be a function defined on $\mathbb{C}^n$ which is constant in all variables except x_1. Then $(1-\phi)f$ extends by zero to be a bounded C^∞ function g_1 in U_1 while $-\phi f$ extends by zero to be a bounded C^∞ function g_2 in U_2. Furthermore, on $U_1 \cap U_2$, $g_1 - g_2 = f$. In other words, we have solved our problem in the class of bounded functions which are C^∞ in the variable z_1 and holomorphic in the remaining variables. Now we need to modify this solution to arrive at one which is holomorphic in z_1 as well.

The fact that $g_1 - g_2 = f$ is holomorphic on $U_1 \cap U_2$ implies that

$$\frac{\partial g_1}{\partial \bar{z}_1} = \frac{\partial g_2}{\partial \bar{z}_1}$$

on $U_1 \cap U_2$ and that implies that they define a bounded $\mathbb{C}^\infty$ function g on $U_1 \cup U_2$ which is $\dfrac{\partial g_i}{\partial \bar{z}_1}$ on U_i. Let V be an open set with compact closure in $(I_1 \cup I_2) \times K$ and let λ_1 be a C^∞ function of z_1 which is 1 on V and which has compact support in $(I_1 \cup I_2) \times K$. Set $\lambda_2 = 1 - \lambda_1$. We then proceed as in the proof of Dolbeault's Lemma. We set $D = U_1 \cup U_2$ and

$$h_i(z) = \frac{1}{2\pi i} \int_D \frac{\lambda_i(\zeta_1) g(\zeta_1, z_2, \ldots, z_n)\, d\zeta_1 \wedge d\bar{\zeta_1}}{\zeta_1 - z_1}$$

and let $h = h_1 + h_2$. As we noted in the proof of Dolbeault's Lemma, $(\zeta_1 - z_1)^{-1}$ is integrable on any bounded subset of $\mathbb{C}$ so that the functions h_i are bounded C^∞ functions on $U_1 \cup U_2$. They are also holomorphic in $(z_2, \ldots, z_n)$. Since λ_1 is compactly supported in $D\ z_1$, so that the line integral term in the generalized Cauchy theorem can be made to vanish, we may conclude as in the proof of Theorem 11.3 that $\frac{\partial h_1}{\partial \bar{z}_1} = \lambda_1 g$ in $U_1 \cup U_2$. This means that $\frac{\partial h_1}{\partial \bar{z}_1} = g$ on $V \times W$. On the other hand, because λ_2 vanishes on V the function h_2 is holomorphic on $V \times W$. Thus, $\frac{\partial h}{\partial \bar{z}_1} = g$ in $V \times W$. However, h is independent of the choice of V and λ and, hence, $\frac{\partial h}{\partial \bar{z}_1} = g$ in all of $U_1 \cup U_2$.

We now have a bounded C^∞ function h on $U_1 \cup U_2$ which is holomorphic in $(z_2, \ldots, z_n)$ and which satisfies

$$\frac{\partial h}{\partial \bar{z}_1} = \frac{\partial g_i}{\partial \bar{z}_1}, \quad \text{in} \quad U_i$$

then $f_i = g_i - h$ is bounded and holomorphic in U_i and $f = f_1 - f_2$ on the intersection $U_1 \cap U_2$. This completes the proof.

A Banach algebra is a complex algebra A with a norm which makes A a Banach space and which is submultiplicative – that is, satisfies $||ab|| \le ||a||||b||$ for all $a, b \in A$. We will only deal with Banach algebras with an identity of norm one. The set of invertible elements in a Banach algebra A form a group which we will sometimes denote A^{-1}. The submultiplicative property, and the fact that Banach algebras are Banach spaces and, hence, are complete, allows one to use power series arguments. These immediately yield the following elementary results from Banach algebra theory:

15.6 Theorem. *If A is any banach algebra, then*

(a) *if $a \in A$ and $||a|| < 1$ then $1 - a$ has an inverse $(1-a)^{-1} = \sum_{n=0}^{\infty} a^n$;*
(b) *A^{-1} is open in A, inversion is continuous in A^{-1} and A^{-1} is a topological group;*
(c) *there is a map $a \to \exp(a) = \sum_{n=0}^{\infty} \frac{a^n}{n!}$ from A to A^{-1} which is a homeomorphism from some neighborhood of 0 in A to the neighborhood $\{b : ||1 - b|| < 1\}$ in A^{-1};*
(d) *on $\{b \in A : ||1 - b|| < 1$ the map $b \to \log b = -\sum_{n=1}^{\infty} \frac{(1-b)^n}{n}$ is an inverse for $\exp$;*
(e) *the subgroup of A^{-1} generated by the image of $\exp$ is open and is equal to the connected component of the identity in A^{-1}.*

In what follows, $Gl_n(\mathbb{C})$ will denote the group of invertible $n \times n$ complex matrices. This is the group of invertible elements of the algebra $M_n(\mathbb{C})$ of all $n \times n$ complex matrices. The latter is a Banach algebra under the standard matrix norm

$$||a|| = \sup\{||ax|| : x \in \mathbb{C}^n,\ ||x|| \le 1\}$$

and, hence, with the topology determined by this norm, $Gl_n(\mathbb{C})$ is a topological group. If U is a domain in $\mathbb{C}^n$, we will also be concerned with the Banach algebra $\mathcal{H}_b(U, M_n)$ of bounded M_n-valued functions on U. Here the norm is given by $||f|| = \sup\{||f(x)|| : x \in U\}$. The invertible group of this Banach algebra is the group of bounded holomorphic functions with values in $Gl_n(\mathbb{C})$.

15.7 Theorem. *If $U \subset \mathbb{C}^n$ is the Cartesian product of simply connected open subsets of $\mathbb{C}$ and K is a compact subset of U, then each holomorphic mapping $f : U \to Gl_n(\mathbb{C})$ may be uniformly approximated on K by holomorphic mappings from $\mathbb{C}^n$ to $Gl_n(\mathbb{C})$.*

Proof. By the Riemann mapping theorem we may, without loss of generality, assume that U is a polydisc centered at the origin. We let V be an open polydisc centered at the origin, containing K and with compact closure in U. Then, on V, f and f^{-1} are bounded holomorphic functions with values in $Gl_n(\mathbb{C})$ – that is, f is an element of the invertible group $\mathcal{H}_b(U, Gl_n(\mathbb{C}))$ of the Banach algebra $\mathcal{H}_b(V, M_n(\mathbb{C}))$. Then a curve $t \to f_t$, $t \in [0,1]$ in $\mathcal{H}_b(V, Gl_n(\mathbb{C}))$ joining f to the constant matrix $f(0)$ may be constructed by setting $f_t(z) = f(tz)$ for each $t \in [0,1]$. Since $Gl_n(\mathbb{C})$ itself is connected (Problem 15.2), this proves that $\mathcal{H}_b(V, Gl_n(\mathbb{C}))$ is connected. By Theorem 15.6(e), this implies that f is a product of elements in the range of the exponential function. Hence, on V,

$$f = \exp(g_1)\exp(g_2)\dots\exp(g_k) \quad \text{with} \quad g_1, g_2, \dots g_n \in \mathcal{H}_b(V, M_n(\mathbb{C})).$$

Now each g_i may be regarded as a matrix with entries which are bounded holomorphic functions on V. By truncating the power series of each entry of each g_i we may approximate each g_i by matrices h_i with polynomial entries as closely as we like in the uniform topology on K. Then

$$\tilde{f} = \exp(h_1)\exp(h_2)\dots\exp(h_k)$$

will be a holomorphic $Gl_n(\mathbb{C})$-valued function on all of $\mathbb{C}^n$. Clearly, the h_i can be chosen so that $\tilde{f}$ approximates f arbitrarily closely on K.

The next lemma is the key to the vanishing theorem we are seeking:

15.8 Cartan's Lemma. *Let (K_1, K_2) be an aligned pair of compact boxes in $\mathbb{C}^n$. Then each holomorphic $Gl_n(\mathbb{C})$-valued function f defined in a neighborhood of $K_1 \cap K_2$ may be factored as $f = f_2^{-1} f_1$ where f_i is a holomorphic $Gl_n(\mathbb{C})$-valued function in a neighborhood of K_i for $i = 1, 2$.*

Proof. We may construct an aligned pair (U_1, U_2) of open boxes such that $K_i \subset U_i$ and f and f^{-1} are holomorphic and bounded in $U_1 \cap U_2$. Let A_i be the Banach algebra $\mathcal{H}_b(U_i, M_n(\mathbb{C}))$ for $i = 1, 2$ and let B be the Banach algebra $\mathcal{H}_b(U_1 \cap U_2, M_n(\mathbb{C}))$ and consider the non-linear map $\phi : A_1 \times A_2 \to B$ defined by

$$\phi(g_1, g_2) = \log((\exp g_2)^{-1} \exp g_1) \quad \text{on} \quad U_1 \cap U_2$$

Now it is easy to see that a function from a Banach algebra to itself which is defined by a convergent power series is infinitely differentiable. It is also easy to see that the analogues of the chain rule and the product rule hold for functions between Banach algebras. It follows that ϕ is infinitely differentiable in a neighborhood of $(0,0)$. We have that $\phi(0,0) = 0$ and a simple calculation (Problem 15.3) shows that

$$\phi'(0,0)(h_1, h_2) = h_1 - h_2$$

By Lemma 15.5, $\phi'(0,0)$ is surjective. By Theorem 15.4, the image of ϕ contains a neighborhood of 0 in B. After composing ϕ with exp, we conclude that the map

$$(g_1, g_2) \to (\exp g_2)^{-1} \exp g_1$$

has image which contains a neighborhood of the identity in B. This means that the theorem is true for f sufficiently close to the identity in B. However, by the previous theorem, we may approximate f arbitrarily closely on $U_1 \cap U_2$ by a $Gl_n(\mathbb{C})$-valued function h which is holomorphic on all of $\mathbb{C}^n$. Then, for an appropriate choice of h we will have fh^{-1} sufficiently close to the identity in B that we may write $fh^{-1} = g_2^{-1} g_1$ with g_i a holomorphic $Gl_n(\mathbb{C})$-valued function on U_i. Then the desired solution is $f = f_2^{-1} f_1$ with $f_1 = g_1 h$ and $f_2 = g_2$.

The procedure used in the next theorem is sometimes known as amalgamation of syzygies.

15.9 Theorem. *Let (K, L) be an aligned pair of compact boxes in $\mathbb{C}^n$. Let $\mathcal{S}$ be a coherent analytic sheaf defined in a neighborhood of $K \cup L$ and suppose that there are exact sequences of analytic sheaves*

$$0 \longrightarrow \mathcal{H}^{p_m} \xrightarrow{\alpha_m} \dots \xrightarrow{\alpha_2} \mathcal{H}^{p_1} \xrightarrow{\alpha_1} \mathcal{H}^{p_0} \xrightarrow{\alpha_0} \mathcal{S} \longrightarrow 0$$

over a neighborhood of K and

$$0 \longrightarrow \mathcal{H}^{q_m} \xrightarrow{\beta_m} \dots \xrightarrow{\beta_2} \mathcal{H}^{q_1} \xrightarrow{\beta_1} \mathcal{H}^{q_0} \xrightarrow{\beta_0} \mathcal{S} \longrightarrow 0$$

over a neighborhood of L. Then there exists an exact sequence of analytic sheaves

$$0 \longrightarrow \mathcal{H}^{r_m} \xrightarrow{\gamma_m} \dots \xrightarrow{\gamma_2} \mathcal{H}^{r_1} \xrightarrow{\gamma_1} \mathcal{H}^{r_0} \xrightarrow{\gamma_0} \mathcal{S} \longrightarrow 0$$

defined in a neighborhood of $K \cup L$.

Proof. The proof is by induction on the length m of the syzygies. If $m = 0$ then we have a pair of isomorphisms:

$$0 \longrightarrow \mathcal{H}^{p_0} \xrightarrow{\alpha_0} \mathcal{S} \longrightarrow 0$$

$$0 \longrightarrow \mathcal{H}^{q_0} \xrightarrow{\beta_0} \mathcal{S} \longrightarrow 0$$

With the first defined over a neighborhood of K and the second over a neighborhood of L. On $K \cap L$, the composition $\phi = \beta_0^{-1} \circ \alpha_0 : \mathcal{H}^{p_0} \to \mathcal{H}^{q_0}$ is an isomorphism over a neighborhood of $K \cap L$. This implies that p_0 and q_0 are the same integer k and that ϕ is determined by a holomorphic $Gl_k(\mathbb{C})$-valued function in a neighborhood of $K \cap L$. By Cartan's Lemma, this function may be factored as $\phi = \mu^{-1} \circ \lambda$ where λ (μ) is a holomorphic $Gl_k(\mathbb{C})$-valued function in a neighborhood of K (L). Then $\beta_0 \circ \mu^{-1} = \alpha_0 \circ \lambda^{-1}$ in a neighborhood of $K \cap L$ and so these two morphisms fit together to define an isomorphism

$$0 \longrightarrow \mathcal{H}^k \xrightarrow{\gamma_0} \mathcal{S} \longrightarrow 0$$

over a neighborhood of $K \cup L$, as required.

For the induction step, we assume that the theorem is true of all pairs of sequences, as above, of length less than m and we suppose we are given a pair of length m. By applying the Riemann mapping theorem in each variable we see that $K \cap L$ has arbitrarily small neighborhoods U which are biholomorphically equivalent to open polydiscs. On a sufficiently small such neighborhood, Lemma 15.1 implies that the sequences

$$\dots \xrightarrow{\alpha_2} \Gamma(U, \mathcal{H}^{p_1}) \xrightarrow{\alpha_1} \Gamma(U, \mathcal{H}^{p_0}) \xrightarrow{\alpha_0} \Gamma(U, \mathcal{S}) \longrightarrow 0$$

$$\dots \xrightarrow{\beta_2} \Gamma(U, \mathcal{H}^{q_1}) \xrightarrow{\beta_1} \Gamma(U, \mathcal{H}^{q_0}) \xrightarrow{\beta_0} \Gamma(U, \mathcal{S}) \longrightarrow 0$$

are exact. We can use the fact that $\Gamma(U, \mathcal{H}^{p_0}) = \mathcal{H}^{p_0}(U)$ is a free $\mathcal{H}(U)$-module to construct the morphism λ which makes commutative the diagram

$$\begin{tikzcd}
\mathcal{H}^{p_0}(U) \arrow[r, "\alpha_0"] \arrow[d, "\lambda"'] & \Gamma(U, \mathcal{S}) \arrow[r] \arrow[d, equal] & 0 \\
\mathcal{H}^{q_0}(U) \arrow[r, "\beta_0"] & \Gamma(U, \mathcal{S}) \arrow[r] & 0
\end{tikzcd}$$

Now λ is a matrix with entries which are holomorphic functions on U and, as such, it defines a morphism of analytic sheaves $\lambda : \mathcal{H}^{p_0} \to \mathcal{H}^{q_0}$ over U such that $\beta_0 \circ \lambda = \alpha_0$. A similar argument shows that we may construct a morphism of analytic sheaves $\mu : \mathcal{H}^{q_0} \to \mathcal{H}^{p_0}$ over U with $\alpha_0 \circ \mu = \beta_0$. We now modify each of the sequences so that the free modules that appear in degree 0 will be identical. Thus, the first sequence is modufied by taking its direct sum with the exact sequence

$$0 \longrightarrow \mathcal{H}^{q_0} \xrightarrow{\text{id}} \mathcal{H}^{q_0} \longrightarrow 0$$

to obtain

$$\cdots \longrightarrow \mathcal{H}^{p_1} \oplus \mathcal{H}^{q_0} \xrightarrow{\tilde{\alpha}_1} \mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0} \xrightarrow{\tilde{\alpha}_0} \mathcal{S} \longrightarrow 0$$

where $\tilde{\alpha}_1 = \alpha_1 \oplus \text{id}$ and $\tilde{\alpha}_0(f \oplus g) = \alpha_0(f)$. Similarly, we modify the second sequence by taking its direct sum (on the other side) with the exact sequence

$$0 \longrightarrow \mathcal{H}^{p_0} \xrightarrow{\text{id}} \mathcal{H}^{p_0} \longrightarrow 0$$

to obtain

$$\cdots \longrightarrow \mathcal{H}^{p_0} \oplus \mathcal{H}^{q_1} \xrightarrow{\tilde{\beta}_1} \mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0} \xrightarrow{\tilde{\beta}_0} \mathcal{S} \longrightarrow 0$$

where $\tilde{\beta}_1 = \mathrm{id} \oplus \beta_1$ and $\tilde{\beta}_0(f \oplus g) = \beta_0(g)$.

Now we define two endomorphisms ϕ and ψ of $\mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0}$ on U to be the maps which have the following matrix representations relative to the direct sum decomposition:

$$\phi = \begin{bmatrix} 1 & -\mu \\ \lambda & 1 - \lambda\mu \end{bmatrix}, \quad \psi = \begin{bmatrix} 1 - \mu\lambda & \mu \\ -\lambda & 1 \end{bmatrix}$$

A calculation shows that $\tilde{\beta}_0 \circ \phi = \tilde{\alpha}_0$, $\tilde{\alpha}_0 \circ \psi = \tilde{\beta}_0$ and $\psi = \phi^{-1}$. Thus, ϕ is an isomorphism and we have the following commutative diagram over U

$$\begin{tikzcd}
\cdots \arrow[r] & \mathcal{H}^{p_1} \oplus \mathcal{H}^{q_0} \arrow[r, "\tilde{\alpha}_1"] & \mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0} \arrow[r, "\tilde{\alpha}_0"] \arrow[d, "\phi"'] & \mathcal{S} \arrow[r] \arrow[d, equal] & 0 \\
\cdots \arrow[r] & \mathcal{H}^{p_0} \oplus \mathcal{H}^{q_1} \arrow[r, "\tilde{\beta}_1"] & \mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0} \arrow[r, "\tilde{\beta}_0"] & \mathcal{S} \arrow[r] & 0
\end{tikzcd}$$

Now by Cartan's Lemma, ϕ can be factored over some neighborhood of $K \cap L$ as $\delta \circ \theta^{-1}$, where θ is an automorphism of $\mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0}$ over some neighborhood of K and δ is an automorphism of $\mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0}$ over some neighborhood of L. Then, in a neighborhood of $K \cap L$ we have

$$\tilde{\beta}_0 \circ \delta = \tilde{\alpha}_0 \circ \theta$$

which means that there is a single morphism

$$\gamma_0 : \mathcal{H}^{p_0} \oplus \mathcal{H}^{q_0} \to \mathcal{S}$$

over a neighborhood of $K \cup L$ such that $\gamma_0 = \tilde{\alpha}_0 \circ \theta$ on a neighborhood of K and $\gamma_0 = \tilde{\beta}_0 \circ \delta$ on a neighborhood of L. If $\mathcal{K}$ is the kernel of γ_0, then we have exact sequences

$$0 \longrightarrow \mathcal{H}^{p_m} \xrightarrow{\alpha_m} \cdots \xrightarrow{\alpha_2} \mathcal{H}^{p'_1} \xrightarrow{\alpha'_1} \mathcal{K} \longrightarrow 0$$

over a neighbohood of K and

$$0 \longrightarrow \mathcal{H}^{q_m} \xrightarrow{\beta_m} \cdots \xrightarrow{\beta_2} \mathcal{H}^{q'_1} \xrightarrow{\beta'_1} \mathcal{K} \longrightarrow 0$$

over a neighborhood of L. Where $p'_1 = p_1 + q_0$, $q'_1 = q_1 + p_0$, $\alpha'_1 = \theta^{-1} \circ \alpha_1$ and $\beta'_1 = \delta^{-1} \circ \alpha$. Since these sequences have length $m-1$, the induction hypothesis implies that there exists an exact sequence

$$0 \longrightarrow \mathcal{H}^{r_m} \xrightarrow{\gamma_m} \cdots \xrightarrow{\gamma_2} \mathcal{H}^{r_1} \xrightarrow{\gamma_1} \mathcal{K} \longrightarrow 0$$

in a neighborhood of $K \cup L$. Combining this with the morphism γ_0 gives us the required sequence for $\mathcal{S}$ over a neighborhood of $K \cup L$. This completes the proof of the theorem.

15.10 Theorem. *If U is an open polydisc in $\mathbb{C}^n$ and $\mathcal{S}$ is a coherent analytic sheaf on U, then for any compact subset $K \subset U$ there is an exact sequence of the form*

$$0 \longrightarrow \mathcal{H}^{p_m} \longrightarrow \dots \longrightarrow \mathcal{H}^{p_1} \longrightarrow \mathcal{H}^{p_0} \longrightarrow \mathcal{S} \longrightarrow 0$$

defined in some neighborhood of K. In other words, $\mathcal{S}$ has a free, finite rank resolution of length n in a neighborhood of any compact subset of U.

Proof. By applying the Riemann mapping theorem in each coordinate, we may reduce the problem to one in which U is an open box in $\mathbb{C}^n$. Also, without loss of generality, we may assume that K is a compact box contained in U. By Lemma 15.2, we know that $\mathcal{S}$ has a resolution like the one above in a neighborhood of each point of U. If K is a cartesian product $I_1 \times \cdots \times I_{2n}$ of closed intervals, we partition each of these intervals into m subintervals of equal length. This results in a partition of K into nm compact boxes. By choosing m large enough, we may assume that $\mathcal{S}$ has a free finite rank resolution of length n in a neighborhood of each of these boxes. Clearly, adjacent pairs of boxes are aligned pairs and, hence, we may apply the previous theorem to conclude that there is a free finite rank resolution in a neighborhood of the union of any two adjacent pairs. By working from left to right along rows in which only the variable x_1 changes we can prove that there is a free finite rank resolution of length n for $\mathcal{S}$ in a neighborhood of the union of the boxes in any such row. These unions provide a new partition of K into compact boxes in which adjacent boxes are aligned. Repeating this argument one variable at a time, we eventually end up with a free finite rank resolution of $\mathcal{S}$ in a neighborhood of K as required.

15.11 Corollary. *Let U be an open polydisc in $\mathbb{C}^n$ and let $\mathcal{S}$ be a coherent analytic sheaf on U. Then $H^p(\Delta, \mathcal{S}) = 0$ for $p > 0$ and for any open polydisc Δ with compact closure in U.*

Proof. This follows immediately from Lemma 15.1 and Theorem 15.10.

We are now in a position to define Stein spaces and to prove a vanishing theorem which is a weak form of Cartan's Theorem B.

15.12 Definition. *Let X be a holomorphic variety. Then*

(i) *if K is a compact subset of X, then the holomorphically convex hull of K in X is the set*
$$\widehat{K} = \{x \in X : |f(x)| \le \sup_{y \in K} |f(y)| \quad \forall f \in \mathcal{H}(X)\};$$

(ii) *a compact subset K of X is said to be holomorphically convex in X if $\widehat{K} = K$;*

(iii) *X is said to be holomorphically convex if $\widehat{K}$ is compact for every compact subset $K \subset X$.*

Note that $\widehat{K}$ is always a closed subset of X and so it will be compact if it is contained in a compact subset of X. Note also that if X is a holomorphic variety, U an open subset of X and K a compact subset of U, then it makes sense to talk about the holomorphically convex hull of K in X and in U. These are not necessarily the same. The open set

U is said to be holomorphically convex if the holomorphically convex hull of K in U is compact for every compact subset K of U – that is, if it is holomorphically convex as a holomorphic variety in its own right. It is easy to see that the intersection of any finite set of holomorphically convex open subsets of a holomorphic variety is also holomorphically convex (Problem 15.4).

15.13 Definition. *A holomorphic variety X is said to be a Stein space if*

(i) *X is holomorphically convex;*
(ii) *for each $x \in X$ the maximal ideal in $\mathcal{H}_x$ is generated by a set of global sections of $\mathcal{H}$;*
(iii) *the global sections of $\mathcal{H}$ separate points in X.*

Note that conditions (ii) and (iii) are automatically inherited by open subsets and so an open subset of a Stein space is also a Stein space provided it is holomorphically convex. In particular, $\mathbb{C}^n$ is clearly a Stein space (Problem 15.5) and so each of its holomorphically convex open subsets is also a Stein space. It is also easy to see that a closed subvariety of a Stein space is a Stein space (Problem 15.6).

15.14 Definition. *An open subset W of a holomorphic variety X is said to be an Oka-Weil subdomain of X if W has compact closure in X and if there is a holomorphic map $\phi : X \to \mathbb{C}^n$ which, when restricted to W, is a biholomorphic map of W onto a closed subvariety of $\Delta(0,1)$.*

15.15 Theorem. *If X is a Stein Space, K is a compact holomorphically convex subset of X and U is an open subset of X containing K, then there exist an Oka-Weil subdomain W of X such that $K \subset W \subset \overline{W} \subset U$ and every coherent analytic sheaf defined on a neighborhood of $\overline{W}$ is acyclic on W.*

Proof. We may as well assume that U has compact closure $\overline{U}$. If $x \in \partial U$ we may choose a function $f \in \mathcal{H}(X)$ such that $|f(x)| > 1 = ||f||_K$, where $||f||_K$ denotes the supremum norm of f on K. Since this inequality will continue to hold in some neighborhood of x and since ∂U is compact, we may choose a finite set of functions $\{f_i\}_{i=1}^k$ so that

$$K \subset W_1 = \{x \in \overline{U} : |f_i(x)| \leq 1, \quad i = 1, \dots k\} \subset U$$

Thus, the functions $\{f_i\}_{i=1}^k$ are the coordinate functions of a holomorphic map $\phi_0 : X \to \mathbb{C}^k$ which maps W_1 into $\Delta(0,1)$. In fact, ϕ_0 is a proper holomorphic map of W_1 into $\Delta(0,1)$ since the inverse image in W_1 of any compact subset of $\Delta(0,1)$ will be closed not only in W_1, but also in the compact set $\overline{U}$.

Now, by (ii) of Definition 15.13, for each point x of $\overline{W}_1$ there is a finite set of global sections of $\mathcal{H}$ which vanish at x and generate the maximal ideal of $\mathcal{H}_x$. Without loss of generality, we may assume that these functions all have modulus less than 1 at each point of W_1. By Theorems 6.15 and 6.16, the map $X \to \mathbb{C}^m$ with these functions as coordinate functions is a biholomorphic map of some neighborhood of x onto a closed subvariety of some neighborhood of $f(x)$. Since $\overline{W}_1$ is compact, finitely many such neighborhoods will cover $\overline{W}_1$. By adjoining all the corresponding functions to the set $\{f_i\}_{i=1}^k$ we obtain a set $\{f_i\}_{i=1}^q$ of functions which are the coordinate functions of a holomorphic map $\phi_1 : X \to \mathbb{C}^q$

which is proper on W_1 and is also has the property that in a neighborhood of each point $x \in \overline{W}_1$ it is a biholomorphic map onto a closed subvariety of a neighborhood of $f(x)$.

In particular, ϕ_1 is locally one to one. This means that the diagonal in $\overline{W}_1 \times \overline{W}_1$ has a neighborhood V in which $\phi(x) - \phi(y)$ is non-vanishing except on the diagonal itself. Another compactness argument and (iii) of Definition 15.13 show that we can find another finite set of global sections of $\mathcal{H}$ (again with modulus less than one at points of W_1) such that whenever $(x, y) \in \overline{W}_1 \times \overline{W}_1 - V$ there is some function f in this set such that $f(x) \neq f(y)$. By adjoining this finite set to the set $\{f_i\}_{i=1}^q$, we obtain a set $\{f_i\}_{i=1}^m$of functions which are the coordinate functions of a holomorphic map $\phi : X \to \mathbb{C}^m$ which is proper, injective and locally biholomorphic on W_1.

A proper, injective, continuous map from one locally compact space into another has closed image and is a homeomorphism onto its image (Problem 15.7). Thus, the image Y of ϕ on W_1 is a closed subset of $\Delta(0, 1)$ and ϕ is a holomorphic homeomorphism onto Y. Furthermore, for each $x \in W_1$, ϕ maps some neighborhood of x biholomorphically onto an open set in Y which is a closed holomorphic subvariety of an open set in $\Delta(0, 1)$. It follows that Y is a closed subvariety of $\Delta(0, 1)$ and ϕ maps W_1 biholomorphically onto Y. Thus, W_1 is an Oka-Weil subdomain of X.

For each polyradius r which has all coordinates less than or equal to 1 we set $W_r = \phi^{-1}(\Delta(0, r))$. Then each W_r is obviously also an Oka-Weil subdomain of X. If $r < s$ means that each coordinate of r is less than the corresponding coordinate of s, then $\overline{W}_r \subset W_s$ whenever $r < s$. Fix an $r < 1$ such that $K \subset W_r$. If a coherent sheaf $\mathcal{S}$ is defined in a neighborhood of $\overline{W}_r$, then we may choose an $s > r$ such that W_s is contained in that neighborhood. Then $\mathcal{S}$ may be considered a coherent analytic sheaf defined on $Y \cap \Delta(0, s)$ and may be extended by zero to a coherent analytic sheaf on all of $\Delta(0, s)$. It follows from Corollary 15.11 that such a sheaf will be acyclic on $\Delta(0, r)$ and, hence, on $Y \cap \Delta(0, r)$. Thus, if we let $W = W_r$, then W is an Oka-Weil subdomain of X which contains K, has compact closure contained in U and has the property that every coherent analytic sheaf defined in a neighborhood of $\overline{W}$ is acyclic on W. This completes the proof.

15.16 Corollary. *If X is a Stein Space then X is the union of a sequence $\{W_n\}$ of Oka-Weil subdomains such that $\overline{W}_n \subset W_{n+1}$ for each n and each coherent analytic sheaf on $\overline{W}_n$ is acyclic on W_n.*

Proof. Since X is countable at infinity, it is the union of an increasing sequence $\{K_n\}$ of compact sets. Suppose we have managed to find Oka-Weil subdomains W_j with the required acyclic property for coherent sheaves and such that $K_j \subset W_j$ and $\overline{W}_{j-1} \subset W_j$ for $j \leq n$. Then $C_n = \overline{W}_{n-1} \cup K_n$ is compact. Since X is a Stein space, $\widehat{C}_n$ is also compact. By Theorem 15.15, C_n is contained in an Oka-Weil subdomain W_n with the required acyclic property. Thus, the Corollary is true by induction.

15.17 Corollary. *Every holomorphic variety X has a neighborhood basis $\mathcal{W}$ with the property that given any $x \in X$ and any open set U containing x, there is a $W \in \mathcal{W}$, with $x \in W$ such that W has compact closure $\overline{W} \subset U$ and every coherent analytic sheaf defined on a neighborhood of $\overline{W}$ is acyclic on W.*

Proof. Every holomorphic variety has a neighborhood base consisting of closed subvarieties of open polydiscs. Since a closed subvariety of an open polydisc is a Stein space, an application of Theorem 15.15 proves the corollary.

Finally, we have the following important application of Theorem 15.11:

15.18. *Suppose Y is a closed subvariety of an open polydisc U and Δ is an open polydisc with compact closure in U. Then every holomorphic function on $\Delta \cap Y$ is the restriction of a holomorphic function on Δ.*

Proof. On U we have an exact sequence of coherent analytic sheaves:

$$0 \to \mathcal{I}_Y \to {}_U\mathcal{H} \to {}_Y\mathcal{H} \to 0$$

By Theorem 15.11, the sheaf $\mathcal{I}_Y$ has vanishing pth cohomology on Δ for $p > 0$. Thus, the long exact sequence of cohomology implies that $\mathcal{H}(\Delta) \to \mathcal{H}(\Delta \cap Y)$ is surjective.

This is as far as we can go in the direction of proving that all coherent sheaves on a Stein space are acyclic (Cartan's Theorem B), without doing an approximation argument like that used in the proof of the vanishing theorem for Dolbeault cohomology (Theorem 11.5). However, to carry out such an argument we need a topology on the space of sections of a coherent sheaf. Specifically, we need to establish that the space of sections of a coherent analytic sheaf has a canonical Frechet space topology. This project is left to the next chapter.

15. Problems

1. Prove Theorem 15.6.
2. Prove that $Gl_n(\mathbb{C})$ is connected.
3. Calculate the derivative at the point $(0, 0)$ of the function ϕ in the proof of Lemma 15.8.
4. Prove that the intersection of any finite collection of holomorphically convex open subsets of a holomorphic variety is also holomorphically convex. Then prove that a holomorphic variety has a locally finite open cover $\mathcal{W}$ with the property that any finite intersection of sets from $\mathcal{W}$ is a Stein space.
5. Prove that $\mathbb{C}^n$ is a Stein space.
6. Prove that a closed subvariety of a Stein space is a Stein space.
7. Prove that a proper, injective, continuous map from one locally compact space into another has closed image and is a homeomorphism onto its image.

16. Fréchet Sheaves – Cartan's Theorems

A topological vector space is a vector space with a topology under which the vector space operations are continuous. A topological vector space is locally convex if it has a neighborhood base at zero consisting of convex balanced sets, where a set is balanced if it is closed under multiplication by scalars of modulus one. A topological vector space is locally convex if and only if its topology is given by a family $\{\rho_\alpha\}$ of seminorms, where a semi-norm on a vector space X is a function ρ from X to the non-negative reals which satisfies

$$\rho(x+y) \leq \rho(x) + \rho(y) \quad \text{and} \quad \rho(\lambda x) = |\lambda|\rho(x)$$

for $x, y \in X$ and $\lambda \in \mathbb{C}$. The topology determined by a family $\{\rho_\alpha\}$ of seminorms is the topology in which a basis of neighborhoods at zero is given by the collection of all sets of the form

$$\{x \in X : \rho_\alpha(x) < \epsilon\}.$$

A continuous linear functional on a topological vector space is a linear complex valued function which is continuous. The space of all continuous linear functionals is called the dual of X and is denoted X^*. There are a number of topologies that can be put on X^* which make it a locally convex topological vector space. They have different properties and are used in different circumstances.

A Fréchet space is a locally convex topological vector space F which is complete and which has its topology defined by a sequence $\{\rho_n\}$ of semi-norms. Without loss of generality, the sequence $\{\rho_n\}$ may be chosen to be increasing. Equivalently, a Fréchet space is a topological vector space which is the inverse limit of a sequence of Banach spaces and bounded linear maps. Equivalently, a Fréchet space is a complete locally convex topological vector space with a topology defined by a translation invariant metric. A bounded set in a Fréchet space is a set B with the property that each of the defining semi-norms ρ_n is bounded on B – equivalently, B is bounded if for every neighborhood U of 0, there is a positive number k such that $B \subset kU$. A Fréchet space F is called a Montel space if every closed bounded subset of F is compact. A separated quotient of a topological vector space is a quotient by a closed subspace.

We will assume knowledge of the following elementary facts concerning locally convex topological vector spaces. The proofs can be found in any text on functional analysis or topological vector space theory.

TVS 1. *Closed subspaces, separated quotients and countable direct products of Fréchet spaces are Fréchet spaces.*

TVS 2. *Closed subspaces, separated quotients, and countable direct products of Montel spaces are Montel spaces.*

TVS 3 (Open Mapping Theorem). *A surjective continuous linear map from a Fréchet space to a Fréchet space is an open map.*

TVS 4 (Closed Graph Theorem). *A linear map from a Fréchet space to a Fréchet space is continuous if and only if its graph is closed.*

TVS 5 (Hahn-Banach Theorem). *The strong form says that a real linear functional on a subspace of a real vector space, dominated by a convex functional, extends to the whole space with preservation of the dominance. We will need two corollaries of this:*

(a) *every continuous linear functional on a subspace of a locally convex topological vector space extends to a continuous linear functional on the whole space;*
(b) *if B is a closed convex balanced set in a locally convex topological vector space X and $x_0 \in X$ is a point not in B, then there exists a continuous linear functional f on X such that $|f(x)| \le 1$ for all $x \in B$ but $|f(x_0)| > 1$.*

TVS 6. *Every locally compact topological vector space is finite dimensional.*

If U is a domain in $\mathbb{C}^n$ then $\mathcal{H}(U)$ is a Fréchet space in the topology of uniform convergence on compact sets. In fact, by Theorem 1.8, $\mathcal{H}(U)$ is a Montel space. The same thing is true for holomorphic functions on a variety:

16.1 Theorem. *If U is an open subset of a holomorphic variety X, then $\mathcal{H}(U)$ is a Montel space in the topology of uniform convergence on compact subsets of U.*

Proof. We may express U as the union of an increasing sequence $\{K_n\}$ of compact subsets, then the sequence $\{\rho_n\}$ of seminorms defined by

$$\rho_n(f) = \sup\{|f(x)| : x \in K_n\}$$

determine the topology of $\mathcal{H}(U)$. It is not obvious that $\mathcal{H}(U)$ is even complete in this topology. However, by Corollary 15.18, X has a neighborhood base consisting of sets W for which W may be imbedded as a closed subvariety of an open polydisc Δ and the restriction map $\mathcal{H}(\Delta) \to \mathcal{H}(W)$ is surjective. Since the kernel of this map consists of the functions which vanish on W, it is closed in $\mathcal{H}(\Delta)$. Hence, $\mathcal{H}(W)$ is a separated quotient of a Montel space and is, therefore, Montel. Now for a general open set U, let $\{W_i\}$ be a countable open cover of U by sets with the above property. Then the map $f \to \{f|_{W_i}\}$ embedds $\mathcal{H}(U)$ as a subspace of the Montel space $\prod_i \mathcal{H}(W_i)$. In fact, it is embedded as the closed subspace $\{\{f_i\} : (f_i)|_{W_i \cap W_j} = (f_j)|_{W_i \cap W_j}\ \forall i, j\}$. Thus, $\mathcal{H}(U)$ is a Montel space.

Given a coherent analytic sheaf $\mathcal{S}$ on a holomorphic variety X, we will define a canonical Fréchet space topology on its space $\mathcal{S}(U)$ of sections over an open set U. The strategy for doing this is to do it first for small neighborhoods of the kind given by Corollary 15.17 and then to argue that this is enough to determine a Fréchet space structure on spaces of sections over arbitrary open sets.

If $x \in X$ then there is a neighborhood W of x on which $\mathcal{S}$ is the cokernel of a morphism of analytic sheaves $\phi : \mathcal{H}^m \to \mathcal{H}^k$. By Corollary 15.17, by shrinking W if necessary, we may assume that W has the property that $\ker \phi$ is acyclic on W. From the long exact sequence of cohomology, we conclude that $\phi : \mathcal{H}(W)^m \to \mathcal{H}(W)^k$ has $\mathcal{S}(W)$ as cokernel, so that $\mathcal{S}(W)$ is a quotient of the Fréchet space $\mathcal{H}^k(W)$. If $\phi : \mathcal{H}(W)^m \to \mathcal{H}(W)^k$ has closed image then $\mathcal{S}(W)$ is a quotient of $\mathcal{H}^k(W)$ by a closed subspace and, hence, is also a Fréchet space. Unfortunately, it is not trivial to prove that ϕ has closed image. This requires an argument using the Weierstrass theorems which is reminiscent of the proof of Oka's Theorem. First a lemma:

16.2 Lemma. *Let U be an open set in $\mathbb{C}^n$,*

(a) *let z be in U and let M be a submodule of $\mathcal{H}^k_z$. Then the set of all $f \in \mathcal{H}^k(U)$ such that $f_z \in M$ is closed in $\mathcal{H}^k(U)$;*

(b) *if $\phi : \mathcal{H}^m \to \mathcal{H}^k$ is a morphism of analytic sheaves on U and Δ is an open polydisc with compact closure in U, then the induced map $\phi : \mathcal{H}^m(\Delta) \to \mathcal{H}^k(\Delta)$ has closed image.*

Proof. We first show that if part (a) holds for given integers n and k then part (b) holds for the same n and k and arbitrary m. In fact, by part (a), every element in the closure of the image of $\phi : \mathcal{H}^m(\Delta) \to \mathcal{H}^k(\Delta)$ has germ at z in the image of ϕ_z at every point $z \in \Delta$. That is, such an element is a section of the sheaf $\mathcal{L} = \operatorname{im}\{\phi : \mathcal{H}^m \to \mathcal{H}^k\}$. If $\mathcal{K}$ is the kermel of this sheaf morphism, then we have a short exact sequence of sheaves on U

$$0 \to \mathcal{K} \to \mathcal{H}^m \xrightarrow{\phi} \mathcal{L} \to 0.$$

From the long exact sequence of cohomology and the vanishing theorem of the previous chapter (Corollary 15.11), we conclude that the corresponding sequence of sections over Δ is also exact. Thus, $\phi : \mathcal{H}^m(\Delta) \to \mathcal{L}(\Delta)$ is surjective. Since the closure of the image of $\phi : \mathcal{H}^m(\Delta) \to \mathcal{H}^k(\Delta)$ is contained in $\mathcal{L}(\Delta)$, we conclude that the image and its closure coincide.

The proof of part (a) will be by induction on n. Part (a) is trivially true in the case $n = 0$ and so we assume that $n > 0$ and that the theorem is true whevever U is an open subset of $\mathbb{C}^{n-1}$. Under this assumption, we next prove part (a) when U is an open subset U of $\mathbb{C}^n$ and $k = 1$.

Thus, let M be an ideal in $\mathcal{H}_z$. We may as well assume that $z = 0$. After a linear change of coordinates, if necessary, we may assume that there is a polydisc Δ centered at 0 and a function $h \in \mathcal{H}(\Delta)$ which is regular in z_n and has germ h_0 belonging to M. Then the Weierstrass preparation theorem implies that h_0 is a unit times a Weierstrass polynomial. Hence, by shrinking Δ if necessary, we may assume that h itself is a Weierstrass polynomial – say of degree k. Then, by the Weierstrass division theorem, every germ $f \in \mathcal{H}_0$ has a unique representation as $f = gh + q$ where g is the germ of a function holomorphic in a neighborhood of 0 and q is a polynomial in z_n of degree less than k with coeficients which are germs at 0 of holomorphic functions in the variables $z_1, \dots, z_{n-1}$. Recall from the proof of the Weierstrass division theorem, that g is given by an integral formula

$$g(z) = \frac{1}{2\pi i} \int_{|\zeta| = r_n} \frac{f(z', \zeta)\, d\zeta}{h(z', \zeta)(\zeta - z)}$$

where $z = (z', z_n) \in \Delta(0, r)$ and $r = (r', r_n)$ is chosen so that $h(z)$ has no zeroes on the part of $\overline{\Delta}(0, r)$ where $|z_n| = r_n$. This clearly means that if Δ is chosen small enough, then g may be chosen so as to have a representative in $\mathcal{H}(\Delta)$ for all $f \in \mathcal{H}(\Delta)$ and, furthermore, that this representative will depend linearly and continuously on f. In other words, for small enough Δ, there is a continuous linear map $\gamma : \mathcal{H}(\Delta) \to \mathcal{H}(\Delta)$ such that $f - \gamma(f)h$ is a polynomial of degree less than k in ζ_n for all $f \in \mathcal{H}(\Delta)$.

Now let N be the set of elements of M which are polynomials in z_n of degree less than k. Clearly, N is a module over the ring ${}_{n-1}\mathcal{H}_0$ of germs of holomorphic functions in the variables $z_1, \dots z_{n-1}$ and, as such, it may be regarded as a submodule of ${}_{n-1}\mathcal{H}_0^k$. Now if f is in the closure of the set of functions in $\mathcal{H}(U)$ with germs in M at 0, then it follows from the continuity of γ that $f - \gamma(f)h$ is in the closure of the set of functions in $\mathcal{H}(\Delta)$ which are polynomials of degree less than or equal to k and have germs beonging to N at 0. By the induction hypothesis, this implies that $f - \gamma(f)h$ has germ at 0 belonging to N. Since $N \subset M$ and $h \in M$, this shows that $f \in M$ and completes the induction step in the case where $k = 1$.

To complete the induction on n for general k, we use induction on k. We assume that part (a) is true with k replaced by $k-1$ and prove that it is also true for k. Thus, let M be a submodule of $\mathcal{H}_0^k$ and define submodules $M_0 \subset \mathcal{H}_0^{k-1}$ and $M_n \subset \mathcal{H}_0$ by

$$M_0 = \{f = (f_1, \dots, f_k) \in M : f_k = 0\}, \quad M_k = \{f_k : (f_1, \dots, f_k) \in M\}$$

Then M_k is isomorphic to M/M_0 under the projection map which sends a k-tuple of functions to its kth element. We choose a polydisc Δ centered at 0 and a finite set of elements of $\mathcal{H}^k(\Delta)$ whose germs at 0 form a set of generators for M. This set of elements then determines a morphism of analytic sheaves over Δ

$$\psi : \mathcal{H}^m \to \mathcal{H}^k$$

with the image of ψ_0 equal to M. Over Δ we define $\phi : \mathcal{H}^m \to \mathcal{H}$ to be ψ followed by the projection of $\mathcal{H}^k$ on its last component. Then the image of ϕ_0 is M_k. Now, by shrinking Δ if necessary, we may assume that ψ and ϕ are defined in an open polydisc containing the closure of Δ and that Δ has compact closure. Under our induction assumption, we know part (a) holds for $k = 1$. Hence, Part (b) holds for $k = 1$. This means that $\phi : \mathcal{H}^m(\Delta) \to \mathcal{H}(\Delta)$ has closed image. By the open mapping theorem for Fréchet spaces, this implies that ϕ is an open mapping onto its image. Now suppose that $\{f_j\}$ is a sequence in $\mathcal{H}^k(U)$ with $(f_j)|_0 \in M$ for every j and suppose that this sequence converges to $f \in \mathcal{H}^k(U)$. If g_j and g are the last components of f_j and f, then $g_j \to g$ and, since ψ is an open map, there exists a convergent sequence $h_j \to h$ in $\mathcal{H}^m(\Delta)$ such that $g_j = \phi(h_j)$ (Problem 16.1). Then $f_j - \psi(h_j)$ is a convergent sequence in $\mathcal{H}^{k-1}(\Delta)$ consisting of elements whose germs at 0 belong to M_0. Thus, by our induction assumption on k, its limit $f - \psi(h)$ also has germ at 0 belonging to M_0. However, $\psi(h)_0 \in M$ and $M_0 \subset M$. Hence, $f_0 \in M$. This completes the induction on k and also the induction on n. Hence, part(a) and part(b) are both proved in general.

16.3 Theorem. *Let X be a holomorphic variety and $\mathcal{S}$ a coherent subsheaf of $\mathcal{H}^k$ for some k. Then $\mathcal{S}(X)$ is a closed subspace of the Fréchet space $\mathcal{H}^k(X)$.*

Proof. We first note that part(a) of Theorem 16.2 extends to varieties. That is, if $f_j \to f$ in ${}_X\mathcal{H}^k(X)$ and if $x \in X$, then Corollary 15.18 and the open mapping theorem imply that there is a neighborhood W of x which may be identified with a closed subvariety of an open polydisc $\Delta \subset \mathbb{C}^n$ and there is a convergent sequence $g_j \in \mathcal{H}^k(\Delta)$ such that $g_j|_W = f_j|_W$ (Problem 16.1). Thus, if the functions f_j all belong to some submodule $M \subset {}_X\mathcal{H}_x$, then the g_j will belong to the inverse image $N \subset {}_n\mathcal{H}_x^k$ of this submodule under the restriction map ${}_n\mathcal{H}_x^k \to {}_X\mathcal{H}_x^k$. It follows from Theorem 16.2 that if g is the limit of the sequence

$\{g_j\}$, then g has germ at x belonging to N and, hence, its restriction f to W has germ at x belonging to M. Thus, part(a) of Theorem 16.2 holds with U replaced by an arbitrary holomorphic variety X.

Now Theorem 16.3 is an immediate consequence, since a global section of $\mathcal{S}$ is just a global section of $\mathcal{H}^k$ which has germ at x belonging to $\mathcal{S}_x$ for each $x \in X$.

A Fréchet sheaf on a space X is just a sheaf of Fréchet spaces on X. Of course, the restriction maps are required to be morphisms in the category of Fréchet spaces – that is, continuous linear maps.

16.4 Theorem. *If $\mathcal{S}$ is a coherent analytic sheaf on an analytic space X, then there is a unique structure of a Fréchet sheaf on $\mathcal{S}$ with the property that if U is any open set and $\mathcal{H}^k \to \mathcal{S}$ is a surjective morphism of analytic sheaves from a free finite rank $\mathcal{H}$ module to $\mathcal{S}$ defined over U, then $\mathcal{H}^k(U) \to \mathcal{S}(U)$ is continuous. Furthermore, $\mathcal{S}$ is a Montel sheaf with this structure.*

Proof. By Corollary 15.17, we may choose a neighbohood base $\mathcal{U}$ for the topology of X consisting of sets U which have compact closure and have the property that coherent sheaves defined in a neighborhood of $\overline{U}$ are acyclic on U. Furthermore, we may choose $\mathcal{U}$ so that for each $U \in \mathcal{U}$, the sheaf $\mathcal{S}$ is the cokernel of a morphism between free finite rank sheaves on a neighborhood of $\overline{U}$.

Fix $U \in \mathcal{U}$. Then there is an exact sequence of analytic sheaves

$$\mathcal{H}^m \xrightarrow{\phi} \mathcal{H}^k \longrightarrow \mathcal{S} \longrightarrow 0$$

defined in a neighborhood of $\overline{U}$. The fact that every coherent sheaf defined in a neighborhood of $\overline{U}$ is acyclic on U implies that the induced sequence on sections over U:

$$\mathcal{H}^m(U) \xrightarrow{\phi} \mathcal{H}^k(U) \longrightarrow \mathcal{S}(U) \longrightarrow 0$$

is also exact. Furthermore, by Theorem 16.3, the image of $\phi : \mathcal{H}^m(U) \to \mathcal{H}^k(U)$ is closed. Hence, $\mathcal{S}(U)$, as a separated quotient of $\mathcal{H}^k(U)$, inherits a Fréchet space topology, in fact, a Montel space topology. Now suppose that $V \subset U$ is another set in our basis $\mathcal{U}$ and suppose for this set we have an exact sequence

$$\mathcal{H}^p \xrightarrow{\psi} \mathcal{H}^q \longrightarrow \mathcal{S} \longrightarrow 0$$

defined in a neighborhood of $\overline{V}$. Then we may construct the following commutative diagram with exact rows:

$$\begin{array}{ccccccc}
\mathcal{H}^m(U) & \xrightarrow{\phi} & \mathcal{H}^k(U) & \longrightarrow & \mathcal{S}(U) & \longrightarrow & 0 \\
\beta\big\downarrow & & \alpha\big\downarrow & & r_{V,U}\big\downarrow & & \\
\mathcal{H}^p(V) & \xrightarrow{\psi} & \mathcal{H}^q(V) & \longrightarrow & \mathcal{S}(V) & \longrightarrow & 0
\end{array}$$

Where $r_{V,U}$ is the restriction map and α and β are constructed by using the familiar lifting argument (the projectivity of free modules). The maps α and β are module homomorphisms and, hence, are given by matrices of holomorphic functions on V through vector-matrix multiplication. This clearly implies that they are continuous. If $\mathcal{S}(U)$ has the topology it inherits from being a quotient of $\mathcal{H}^k(U)$, then $\mathcal{H}^k(U) \to \mathcal{S}(U)$ is an open map. Also, the map $\mathcal{H}^k(V) \to \mathcal{S}(V)$ is continuous if $\mathcal{S}(V)$ is given the quotient topology. These facts combine to force the map $r_{V,U}$ to be continuous. We draw two conclusions from this:

(1) The topology on $\mathcal{S}(U)$ is independent of the way in which it is expressed as a quotient of a free finite rank $\mathcal{H}(U)$-module; and
(2) if $V \subset U$ are two sets in our basis, then the restriction map $r_{V,U} : \mathcal{S}(U) \to \mathcal{S}(V)$ is continuous.

It is now clear how to define the topology on $\mathcal{S}(U)$ for a general open set U. We cover U by a countable collection $\{W_i\}$ of sets from our basis. Then $f \to \{f|_{W_i}\} : \mathcal{S}(U) \to \prod_i \mathcal{S}(W_i)$ is an injective continuous linear map of $\mathcal{S}(U)$ onto a closed subspace of the Montel space $\prod_i \mathcal{S}(W_i)$. The image is closed because it is the subspace of $\prod_i \mathcal{S}(W_i)$ consisting of $\{g_i\}$ such that $g_i = g_j$ on $W_i \cap W_j$ for all i, j. Since a closed subspace of a Montel space is Montel, this serves to put a Fréchet space structure on $\mathcal{S}(U)$ under which it is actually a Montel space. Note that, by construction, this topology has the property that for each $x \in U$ there is a basic neighborhood $W \subset U$ containing x such that the restriction map $\mathcal{S}(U) \to \mathcal{S}(W)$ is continuous. Now suppose that $V \subset U$ is another open set and $\mathcal{S}(V)$ is given a Fréchet space topology with this same property. Then we claim that the restriction map $r_{V,U} : \mathcal{S}(U) \to \mathcal{S}(V)$ has closed graph and, hence, is continuous. To see this, let $\{(f_n, r_{V,U}(f_n))\}$ be a sequence in the graph which converges to the point (f, g). Then $f_n \to f$ and $r_{V,U}(f_n) \to g$. Now for each point x of V we can choose a basic neighborhood W of x such that the restrictions $\mathcal{S}(U) \to \mathcal{S}(W)$ and $\mathcal{S}(V) \to \mathcal{S}(W)$ are both continuous. This clearly implies that $r_{V,U}(f)|_W = g|_W$. But since this is true for a neighborhood of each point of V, we conclude that $r_{V,U}(f) = g$ and, hence, that the graph of $r_{V,U}$ is closed as required. We draw two conclusions from this:

(1) The Fréchet space topology on $\mathcal{S}(U)$ is uniquely defined by the property that for each basic open set $W \subset U$ the restriction map $\mathcal{S}(U) \to \mathcal{S}(W)$ is continuous; and
(2) if $V \subset U$ are two open sets, then the restriction map $r_{V,U} : \mathcal{S}(U) \to \mathcal{S}(V)$ is continuous.

Thus, we have proved the existence of a Fréchet sheaf structure on $\mathcal{S}$ under which it is a Montel sheaf. Furthermore, it is clear from the construction that a morphism $\mathcal{T} \to \mathcal{S}$ is continuous if and only if it is continuous locally, that is, if and only if $\mathcal{T}(W) \to \mathcal{S}(W)$ is continuous for a neighborhood W of each point of X. This, and the construction of the topology on basic open sets shows that the topology has the property that if $\mathcal{H}^k \to \mathcal{S}$ is a surjective morphism over an open set U, then $\mathcal{H}^k(U) \to \mathcal{S}(U)$ is continuous. The uniqueness is obvious from the construction.

In view of the above theorem, we may, henceforth, assume that every coherent analytic sheaf comes equipped with a structure of a Montel sheaf. A morphism $\phi : \mathcal{S} \to \mathcal{T}$ between two Fréchet sheaves is said to be continuous if $\phi : \mathcal{S}(U) \to \mathcal{T}(U)$ is continuous for each open set U. When is a morphism of coherent analytic sheaves continuous? Always!

16.5 Theorem. *A morphism $\phi : \mathcal{S} \to \mathcal{T}$ of sheaves of $\mathcal{H}$-modules between two coherent analytic sheaves is automatically continuous.*

Proof. As was pointed out in the proof of the previous theorem, a morphism between coherent analytic sheaves is continuous if and only if it is continuous locally. If x is a point of X then we may choose a basic neighborhood W for x of the kind used in the proof of the previous theorem. Then there are surjective morphisms $\alpha : \mathcal{H}^k \to \mathcal{S}$ and $\beta : \mathcal{H}^m \to \mathcal{T}$ defined over U which are also surjective on sections over U. Then the usual lifting argument gives us the λ in the following commutative diagram

$$\begin{array}{ccccc}
\mathcal{H}^k(W) & \xrightarrow{\alpha} & \mathcal{S}(W) & \longrightarrow & 0 \\
\lambda \big\downarrow & & \phi \big\downarrow & & \\
\mathcal{H}^m(W) & \xrightarrow{\beta} & \mathcal{T}(W) & \longrightarrow & 0
\end{array}$$

It follows that ϕ is continuous because β and λ are continuous and α is open. Thus, ϕ is locally continuous and, hence, continuous.

We are now in a position to prove the main theorem in the subject – Cartan's Theorem B. The first step is an approximation theorem:

16.5 Theorem. *If X is a Stein space, $\mathcal{S}$ a coherent analytic sheaf on X and $W \subset X$ is an Oka-Weil subdomain, then the space of restrictions to W of global sections of $\mathcal{S}$ is dense in $\mathcal{S}(W)$.*

Proof. We first prove that this is true in the case where $\mathcal{S}$ is the structure sheaf $\mathcal{H}$. Since W is an Oka-Weil subdomain, there is a holomorphic map $\phi : X \to \mathbb{C}^n$ such that ϕ maps W biholomorphically onto a closed subvariety of the unit polydisc $\Delta(0, 1)$ centered at 0. If f is a holomorphic function on W and K is a compact subset of W then it follows from Corollary 15.18 that f has the form $g \circ \phi$ in a neighborhood of K, where g is a holomorphic function on an open polydisc Δ with closure contained in $\Delta(0, 1)$. If $\{h_j\}$ is a sequence of polynomials converging to g in the topology of uniform convergence on compact subsets of Δ, then $f_j = h_j \circ \phi$ defines a sequence of holomorphic functions on X which converge uniformly to f on K.

In order to prove the theorem for a general coherent analytic sheaf $\mathcal{S}$, we choose a sequence $\{W_n\}$ of Oka-Weil subdomains with $W = W_0$ and $\overline{W}_n \subset W_{n+1}$ for each n. We claim that the image of $\mathcal{S}(W_n)$ under restriction is dense in $\mathcal{S}(W_m)$ if $m < n$. In fact, we can find a surjective morphism $\mathcal{H}^k \to \mathcal{S}$ in a neighborhood of $\overline{W}_n$. Then both $\mathcal{H}^k(W_n) \to \mathcal{S}(W_n)$ and $\mathcal{H}^k(W_m) \to \mathcal{S}(W_m)$ are surjective. Thus, an element f of $\mathcal{S}(W_m)$ can be lifted to an element g of $\mathcal{H}^k(W_m)$ and this can be expressed as the limit of a sequence of restrictions of elements $h_j \in \mathcal{H}^k(W_m)$ by the result of the above paragraph. The image of the sequence $\{h_j\}$ in $\mathcal{S}(W_n)$ will then have the property that its restriction to $\mathcal{S}(W_m)$ converges to f.

To finish the proof, we choose a translation invariant metric ρ_n defining the topology of $\mathcal{S}(W_n)$ for each n. Since the metric ρ_n may be replaced by the metric $\sum_{i=0}^{n} \rho_i$ without changing the topology it defines, we may assume without loss of generality that the sequence of metrics is increasing in the sense that $\rho_m(f) \le \rho_n(f)$ if $m < n$ and $f \in \mathcal{S}(W_m)$. Then if $\epsilon > 0$ and $f \in \mathcal{S}(W) = \mathcal{S}(W_0)$, we choose $g_1 \in \mathcal{S}(W_1)$ such that $\rho_0(f - g_1) < \epsilon/2$.

We then inductively choose a sequence $g_n \in \mathcal{S}(W_n)$ with $\rho_{n-1}(g_{n-1} - g_n) < \epsilon 2^{-n}$. Clearly, for each m, the sequence $\{(g_n)|_{W_m} : n > m\}$ converges in the metric ρ_m to an element $h_m \in \mathcal{S}(W_m)$. Furthermore, $h_m = (h_n)|_{W_m}$ for $m < n$ and $\rho_0(f - h_0) < \epsilon$. Hence, the h_n define a global section $h \in \mathcal{S}(X)$ such that $\rho_0(f - h) < \epsilon$. This completes the proof.

16.6 Cartan's Theorem A. *If X is a Stein space and $\mathcal{S}$ is a coherent analytic sheaf on X, then $\mathcal{S}(X)$ generates $\mathcal{S}_x$ at every point $x \in X$.*

Proof. Let M_x be the submodule of $\mathcal{S}_x$ generated by the global sections of $\mathcal{S}$. Let W be a neighborhood of x which is an Oka-Weil subdomain and which has the property that $\mathcal{S}(W)$ generates $\mathcal{S}_x$. We find a surjective morphism $\phi : \mathcal{H}^k \to \mathcal{S}$ over W with the property that $\phi : \mathcal{H}^k(W) \to \mathcal{S}(W)$ is also surjective. Then the space of sections $f \in \mathcal{H}^k(W)$ such that $\phi(f)_x \in M_x$ is closed by Theorem 16.2. It follows from the open mapping theorem that the set of $g \in \mathcal{S}(W)$ such that $g_x \in M_x$ is also closed. However, this set contains the set of restrictions to W of all elements of $\mathcal{S}(X)$ which by Theorem 16.5 is dense in $\mathcal{S}(W)$. It follows that $M_x = \mathcal{S}_x$ and the proof is complete.

16.7 Cartan's Theorem B. *If X is a Stein space and $\mathcal{S}$ is a coherent analytic sheaf on X, then $\mathcal{S}$ is acyclic on X.*

Proof. With the machinery we have built up, the proof is just like the proof of Theorem 11.5. We use Corollary 15.16 to express X as the union of a sequence of Oka-Weil subdomains $\{W_n\}$ such that $\overline{W}_n \subset W_{n+1}$ and $\mathcal{S}$ is acyclic on each W_n. Then, since the W_n's form a nested sequence, $\{W_n\}$ is a Leray cover of X. We also have that the space of global sections of $\mathcal{S}$ is dense in $\mathcal{S}(W_n)$ for each n. That $\mathcal{S}$ is acyclic on X now follows from a principle that is quite general and which we state in the next lemma.

16.8 Lemma. *Suppose $\mathcal{S}$ is a sheaf on a topological space X. If X is the union of an increasing sequence $\{W_i\}$ of open subsets such that $\mathcal{S}$ is acyclic on each W_i, then*

(a) *$H^p(X, \mathcal{S}) = 0$ for $p > 1$;*
(b) *if, in addition, $\mathcal{S}$ is a Fréchet sheaf and $\mathcal{S}(X)$ is dense in $\mathcal{S}(W_i)$ for each i then $H^1(X, \mathcal{S}) = 0$.*

Proof. Let

$$0 \longrightarrow \mathcal{S} \longrightarrow \mathcal{F}^0 \xrightarrow{\delta^0} \mathcal{F}^1 \xrightarrow{\delta^1} \cdots \longrightarrow \mathcal{F}^p \xrightarrow{\delta^p} \cdots$$

be a flabby resolution of $\mathcal{S}$. Suppose $p > 1$ and $f \in \mathcal{F}^p(X)$ with $\delta f = 0$. Then we will prove by induction that there is a sequence $\{g_n\}$ with $g_n \in \mathcal{F}^{p-1}(W_n)$, $g_n = g_{n-1}$ on W_{n-1} and $\delta g_n = f_n$ on W_n. Clearly, if we can show this then part(a) of the Lemma will be established. since such a sequence determines a global section $g \in \mathcal{F}^{p-1}(X)$ such that $g = g_n$ on W_n and, consequently, $\delta g = f$ on all of X.

Suppose we have managed to construct the sequence $\{g_n\}$ for all $n \le m$. Because $\mathcal{S}$ is acyclic on each W_n, there exists a section $\tilde{g}_{m+1} \in \mathcal{F}^{p-1}(W_{m+1})$ such that $\delta g_{m+1} = f$. However, $\tilde{g}_{m+1} - g_m$ may not be zero on W_m. Whatever it is, it is in the kernel of δ and, hence, there exists $h_{m+1} \in \mathcal{F}^{p-2}(W_m)$ such that $\delta h_{m+1} = \tilde{g}_{m+1} - g_m$ on W_m. Since $\mathcal{F}^{p-2}$ is flabby we may assume that h_{m+1} is actually a section defined on all of X. We then set

$g_{m+1} = \tilde{g}_{m+1} - \delta h_{m+1}$ on W_{m+1}. Clearly, this serves to extend our sequence $\{W_n\}$ to $n = m+1$ and completes the induction. This completes the proof of part(a). Note that this proof does not work when $p = 1$ since, in this case, there is no $\mathcal{F}^{p-2}$.

We now proceed with the proof of part(b). By Leray's Theorem we may compute $H^1(X, \mathcal{S})$ using Čech cohomology for the Leray cover $\{W_n\}$. Our argument will involve an approximation argument for 0-cochains. The space of 0-cochains for $\mathcal{S}$ and the cover $\{W_n\}$ on an open set U is just $\prod_n \mathcal{S}(W_n \cap U)$ and, as a countable product of Fréchet spaces, is itself a Fréchet space. For each j, we let ρ_j be a translation invariant metric on the Fréchet space of 0-Čech cochains on W_j. As in the proof of Theorem 16.5, we may assume that the sequence $\{\rho_j\}$ is increasing in the sense that $\rho_j(g|_{W_j}) \le \rho_k(g)$ if $j < k$ and g is a 0-cochain on W_k. Suppose f is a 1-cocycle for the cover $\{W_n\}$. We inductively construct a sequence $\{g_j\}$ where g_j is a 0-cochain on W_j such that $\delta g_j = f$ on W_j and $\rho_j(g_j - g_{j+1}|_{W_j}) < 2^{-j}$ for each j. Suppose such a sequence $\{g_j\}$ has been constructed for all indices $j < k$. We use the fact that $H^1(W_{k-1}, \mathcal{S}) = 0$ to find a 0-cochain t on W_k such that $\delta t = f$ on W_k. Then,

$$\delta(t - g_{k-1}) = 0$$

in W_{k-1}. This means that $t - g_{k-1}$ is the 0-cochain on W_{k-1} determined by a section $r \in \mathcal{S}(W_{k-1})$. Using the density hypotheses, we may choose a global section that approximates this section as closely as we desire. Thus, there is a global 0-cochain s such that $\delta s = 0$ and $\rho_{k-1}(t - g_{k-1} - s) < 2^{-k+1}$. Then $g_k = t - s$ has the properties that $\delta g_k = f$ on W_k and $\rho_{k-1}(g_k - g_{k-1}) < 2^{-k+1}$. Thus, by induction, we may construct the sequence $\{g_j\}$ as claimed. Now on a given W_k consider the sequence $\{g_j)|_{W_k}\}_{j=k}^{\infty}$. This is a Cauchy sequence in the metric ρ_k since $\rho_k(g_{j+1} - g_j) < \rho_j(g_{j+1} - g_j) < 2^{-j}$ for $j \ge k$. Furthermore, the terms of this sequence differ from the first term by cocycles (since the difference is killed by δ. Thus, the sequence may be regarded as a fixed cochain plus a uniformly convergent sequence of cocycles. It follows that this sequence actually converges in the topology of 0-cochains on W_k and the limit h_k satisfies $\delta h_k = f|_{W_k}$. Furthermore, $h_{k+1} = h_k$ on W_k and, hence, the h_k determine a 0-cochain h on X. Clearly, $\delta h = f$. Thus, every 1-cochain is a coboundary and the proof is complete.

Cartan's Theorem B has a host of corollaries. We list some of these in the next few pages. The first five follow immediately using sheaf theory techniques which are, by now, familiar to us and so we leave their proofs as excercises:

Corollary 16.9. *If X is a Stein space, then every surjective morphism $\mathcal{S} \to \mathcal{T}$ of coherent analytic sheaves induces a surjective morphism $\mathcal{S}(X) \to \mathcal{T}(X)$ on global sections.*

Corollary 16.10. *If Y is a closed subvariety of a Stein space X then every holomorphic function on Y is the restriction of a holomorphic function on X.*

Corollary 16.11. *If X is a Stein space, an if $\{f_i\}$ is a finite set of holomorphic functions on X which does not vanish simultaneously at any point of X then there is a set of holomorphic functions $\{g_i\}$ on X such that $\sum g_i f_i = 1$. In other words, each finitely generated ideal of $\mathcal{H}(X)$ is contained in a maximal ideal of the form $M_x = \{f \in \mathcal{H}(X) : f(x) = 0\}$ for some $x \in X$.*

If K is compact, then by $\mathcal{H}(K)$ we mean the algebra of functions holomorphic in a neighborhood of K, in other words, the sections of the sheaf $\mathcal{H}$ on K.

16.12 Corollary. *Let K be a compact holomorphically convex subset of a Stein space X. Then every coherent analytic sheaf defined in a neighborhood of K is acyclic on K and is generated as a sheaf of $\mathcal{H}$ modules by a finite set of sections over K;*

16.13 Corollary. *Let K be a compact holomorphically convex subset of a Stein space X. Then every maximal ideal of $\mathcal{H}(K)$ is of the form $M_x = \{f \in \mathcal{H}(X) : f(x) = 0\}$ for some $x \in K$.*

Let $\mathcal{H}^*$ denote the sheaf of invertible holomorphic functions under multiplication. In view of the short exact sequence

$$0 \longrightarrow \mathbb{Z} \longrightarrow \mathcal{H} \xrightarrow{f \to \exp(2\pi i f)} \mathcal{H}^* \longrightarrow 0$$

we conclude that:

16.14 Corollary. *If X is a Stein space then*

(a) $H^1(X, \mathbb{Z}) \simeq \mathcal{H}^*(X)/\exp(\mathcal{H}(X))$;
(b) $H^2(X, \mathbb{Z}) \simeq H^1(X, \mathcal{H}^*)$;

Using Corollary 16.14 as in problem 9.9, one can now show that

16.15 Corollary. *If X is a Stein space then the group of isomorphism classes of holomorphic line bundles is isomorphic to the group of isomorphism classes of continuous line bundles which is isomorphic to $H^2(X, \mathbb{Z})$.*

Now suppose X is a connected complex manifold. The rings $\mathcal{H}(U)$ for U a connected open subset of X are all integral domains. Let $\mathcal{M}$ denote the sheaf on X generated by the presheaf which assigns to each connected open set U the quotient field of $\mathcal{H}(U)$. Then $\mathcal{M}$ is called the *sheaf of meromorphic functions* on U. We let $\mathcal{M}^*$ be the sheaf of non-zero elements of $\mathcal{M}$ under multiplication. Clearly, $\mathcal{H}^* \subset \mathcal{M}^*$. The quotient sheaf is denoted $\mathcal{D}$ and called the *sheaf of divisors* on X.

The long exact sequence of sections associated to the short exact sequence

$$0 \longrightarrow \mathcal{H}^* \longrightarrow \mathcal{M}^* \longrightarrow \mathcal{D} \longrightarrow 0$$

combined with Corollary 16.14 yields an exact sequence

$$\Gamma(X, \mathcal{M}^*) \longrightarrow \Gamma(X, \mathcal{D}) \longrightarrow H^1(X, \mathcal{H}^*) \longrightarrow H^1(X, \mathcal{M}^*)$$

An element of $\Gamma(X, \mathcal{D})$ is called a *divisor* and its image in $H^1(X, \mathcal{H}^*) = H^2(X, \mathbb{Z})$ is called its *Chern class*. Now by Corollary 16.15 an element of $H^1(X, \mathcal{H}^*)$ corresponds to a holomorphic line bundle L on X and L corresponds to the zero element of $H^1(X, \mathcal{H}^*)$ if and only if it is the trivial line bundle – that is, if and only if it has a section that is nowhere vanishing. However, if L just has a section that is not identically zero, then it is easy to see that the class in $H^1(X, \mathcal{H}^*)$ corresponding to L is sent to 0 by the map $H^1(X, \mathcal{H}^*) \to H^1(X, \mathcal{M}^*)$ (problem 16.7). However, the holomorphic sections of a holommorphic line bundle form a coherent analytic sheaf. By Cartan's Theorem A, every coherent sheaf on a Stein space is generated by its global sections. We conclude that every holomorphic line bundle on X has a section which is not identically zero. Hence, $H^1(X, \mathcal{H}^*) \to H^1(X, \mathcal{M}^*)$ is the zero map, and we conclude from the above exact sequence and Corollary 16.14 that:

16.16 Corollary. *If X is a Stein manifold, then there is an exact sequence*

$$0 \longrightarrow \Gamma(X, \mathcal{H}^*) \longrightarrow \Gamma(X, \mathcal{M}^*) \longrightarrow \Gamma(X, \mathcal{D}) \longrightarrow H^2(X, \mathbb{Z}) \longrightarrow 0$$

Thus, every element of $H^2(X, \mathbb{Z})$ is the Chern class of some divisor and a divisor has Chern class zero if and only if it is the divisor of a global meromorphic function.

We now turn to the final topic of this section – cohomology of coherent sheaves on compact holomorphic varieties. We cannot expect a compact variety to be a Stein space, since the only global holomorphic functions on such a variety are constants and so the conclusion of Theorem A does not hold. However, we will prove that every coherent analytic sheaf on a compact holomorphic variety has finite dimensional cohomology. The proof depends on a theorem of Schwartz on compact perturbations of surjective bounded linear maps between Fréchet spaces.

A continuous linear map $\phi : X \to Y$ between two topological vector spaces is said to be *compact* if there exists a neighborhood U of 0 in X such that $\phi(U)$ has compact closure in Y.

16.17 Theorem. *If $\mathcal{S}$ is a coherent analytic sheaf on a Stein space X and U is an open set with compact closure in X, then the restriction map $\mathcal{S}(X) \to \mathcal{S}(U)$ is a compact map.*

Proof. We first note that this is true if $\mathcal{S}$ is the structure sheaf $\mathcal{H}$. In fact, $\{f \in \mathcal{H}(X) : \sup_{\overline{U}} |f(x)| < 1\}$ is a neighborhood of zero in $\mathcal{H}(X)$ and its image in $\mathcal{H}(U)$ is bounded and, hence, has compact closure since $\mathcal{H}(U)$ is a Montel space.

Now suppose $\mathcal{S}$ is any coherent sheaf on X. We choose an Oka-Weil subdomain W such that $\overline{U} \subset W$ and such that for some k there is a surjective morphism $\mathcal{H}^k \to \mathcal{S}$ defined over W. The corresponding map on global sections is also surjective by Cartan's Theorem B. Thus, we have the following commutative diagram:

$$\begin{array}{ccccc}
\mathcal{H}^k(W) & \longrightarrow & \mathcal{S}(W) & \longrightarrow & 0 \\
\downarrow & & \downarrow & & \\
\mathcal{H}^k(U) & \longrightarrow & \mathcal{S}(U) & \longrightarrow & 0
\end{array}$$

There is an neighborhood of zero in $\mathcal{H}^k(W)$ whose image in $\mathcal{H}^k(U)$ has compact closure by the result of the previous paragraph. The image of this neighborhood in $\mathcal{S}(W)$ is a neighborhood of zero in $\mathcal{S}(W)$ by the open mapping theorem and its image in $\mathcal{S}(U)$ will have compact closure by the commutativity of the diagram. Thus, $\mathcal{S}(W) \to \mathcal{S}(U)$ is a compact map. However, $\mathcal{S}(X) \to \mathcal{S}(U)$ is the composition of this map with the restriction map $\mathcal{S}(X) \to \mathcal{S}(W)$. Since, the composition of a continuous linear map with a compact map is clearly compact, the proof is complete.

We now prove the Cartan- Serre theorem assuming Schwartz's Theorem. We will end the Chapter with a proof of Schwartz's Theorem.

16.18 Cartan-Serre Theorem. *If X is a compact holomorphic variety and $\mathcal{S}$ is a coherent analytic sheaf, then $H^p(X, \mathcal{S})$ is finite dimensional for all p.*

Proof. We choose a finite open cover $\mathcal{W} = \{W_i\}$ of X consisting of sets which are Stein spaces. Finite intersections of sets in this cover are also Stein spaces (Problem 15.4). We then choose another such cover $\mathcal{U} = \{U_j\}$ which is a refinement of the first cover and, in fact, has the property that for each j there is an integer $\iota(j)$ such that $\overline{U}_j \subset W_{\iota(j)}$. Then for each multi-index $\alpha = (j_0, \dots, j_p)$ the multi-index $\iota(\alpha) = (\iota(j_0), \dots, \iota(j_p))$ has the property that $\overline{U}_\alpha \subset W_{\iota(\alpha)}$ so that the restriction map $\mathcal{S}(W_{\iota(\alpha)}) \to \mathcal{S}(U_\alpha)$ is a compact map. It follows that the refinement morphism $\iota^* : \mathcal{C}^p(\mathcal{W}) \to \mathcal{C}^p(\mathcal{U})$ from the space of Čech p-cochains for $\mathcal{W}$ to the space of Čech p-cochains for $\mathcal{U}$ is a compact map between Fréchet spaces. Since both covers $\mathcal{W}$ and $\mathcal{U}$ are Leray covers for $\mathcal{S}$ the map ι^* induces an isomorphism of cohomology. Hence, if $\mathcal{Z}^p(\mathcal{W})$ and $\mathcal{Z}^p(\mathcal{U})$ are the spaces of Čech p-cocycles for $\mathcal{W}$ and $\mathcal{U}$, then the map

$$f \oplus g \to \delta^{p-1}(f) + \iota^*(g) : \mathcal{C}^{p-1}(\mathcal{W}) \oplus \mathcal{Z}^p(\mathcal{W}) \to Z^p(\mathcal{W})$$

is surjective. Since ι^* is compact, it follows from Schwartz's theorem (Theorem 16.20) that δ^{p-1} has closed image and finite dimensional cokernel. Hence, $H^p(X, \mathcal{S})$ is finite dimensional.

It remains to prove Schwartz's theorem. We first prove a dual version of Schwartz's theorem from which the theorem itself will follow:

16.19 Theorem. *Let X and Y be locally convex topological vector spaces and let $A : X \to Y$ be a continuous linear map which has closed image and is a topological isomorphism onto its image, Let $C : X \to Y$ be a compact continuous linear map. Then $B = A + C$ has finite dimensional kernel K, closed image I and the induced map $B : X/K \to I$ is a topological isomorphism.*

Proof. A and $-C$ agree on the kernel K of B. Thus, $A|_K$ is a topological isomorphism of K onto a subspace of Y, but it is also a compact map. This implies that the image of K in Y under A has a neighborhood of zero with compact closure. However, a topological vector space which is locally compact is necessarily finite dimensional. It follows that K is finite dimensional.

Now there is a closed subspace $L \subset X$ which is complementary to K in X. This follows from the Hahn-Banach theorem. In fact, if $\{x_i\}$ is a basis for the vector space K, then the Hahn-Banach theorem implies that we can find for each i a continuous linear functional f_i on X with $f_i(x_j) = \delta_{ij}$. The intersection of the kernels of the f_i will then be a closed complement for K. If L is such a complement, then $A|_L$ and $C|_L$ are continuous linear maps of L into Y, the first a topological isomorphism onto its image and the second a compact map. Thus, we are back in our original situation except now we have that $A|_L + C|_L$ is injective. Thus, to complete the proof it is enough to prove the theorem in the case where $A + B$ is injective.

We need to prove that $B = A + C$ has closed image and is a topological isomorphism onto its image. Thus, let x_α be a net in X and suppose that $B(x_\alpha)$ converges to $y \in Y$. Since C is compact, there is a continuous seminorm ρ on X such that $U_\rho = \{x \in X : \rho(x) < 1\}$ is a neighborhood which C maps to a set with compact closure in Y. Suppose $\{\rho(x_\alpha)\}$ is bounded, say by M. Then $C(x_\alpha)$ is contained in the set $(M+1)C(U_\alpha)$, which has compact closure in Y. Thus, $C(x_\alpha)$ has a cluster point $z \in Y$ and $A(x_\alpha) = B(x_\alpha) - C(x_\alpha)$ has $y - z$ as a cluster point. Since the image of A is closed, we have that $y - z = Ax$ for a unique $x \in X$. Since A is a topological isomorphism onto its image, we conclude that the original net $\{x_\alpha\}$ has x as a cluster point and $Bx = y$ since B is continuous. Thus, in the case where $\{\rho(x_\alpha)\}$ is bounded, we have that $\{x_\alpha\}$ has a cluster point x and $B(x) = y$.

Now suppose that $\{\rho(x_\alpha)\}$ is unbounded, then, after possibly modifying $\{x_\alpha\}$ to eliminate terms where $\rho(x_\alpha) = 0$, we may consider the net $\{x'_\alpha\}$ defined by

$$x'_\alpha = \frac{x_\alpha}{\rho(x_\alpha)}$$

This net has the property that $B(x'_\alpha)$ converges to zero and $\{\rho(x'_\alpha)\}$ is bounded, in fact $\rho(x'_\alpha) = 1$ for every α. Thus, we are back in the previous case except that y has been replaced by zero. We conclude that $\{x'_\alpha\}$ has a cluster point x and $B(x) = 0$. However, this is only possible if $x = 0$, since we are assuming that B is injective, and zero cannot be a cluster point of a net of elements x'_α with $\rho(x'_\alpha) = 1$. Thus, the first case was the only one possible and we conclude that every net $\{x_\alpha\} \subset X$ for which $B(x_\alpha)$ converges in Y has a cluster point in X. This is enough to imply that the image of B is closed and the inverse of B on that image is continuous. This completes the proof.

It turns out that the previous theorem is the dual of the one we want. In order to prove Schwartz's theorem we must first define a topology on the dual of a Fréchet space and prove an important theorem of Mackey and Arens concerning duality. There are many ways to topologize the dual of a locally convex topological vector space. The useful ways are of the following type:

A *saturated family* of bounded subsets of a topological vector space X is a family which is closed under subsets, multiplication by scalars, finite unions and closed convex balanced hulls. Such a family is said to cover X if X is the union of the sets in the family.

16.20 Definition. *If X is a locally convex topological vector space and κ is a saturated family of bounded subsets of X which covers X, then X^*_κ will denote the space of continuous linear functionals on X with the topology of uniform convergence on sets in κ.*

Clearly X^*_κ is a locally convex topological vector space. A family of seminorms defining the topology is the family of all seminorms of the form ρ_K where K is a set in κ and $\rho_K(f) = \sup\{|f(x)| : x \in K\}$. The family of sets we want to use is the family c of all sets with compact closure. This is not always a saturated family since it is not always true that the closed convex, balanced hull of a compact set is compact. However, for Fréchet spaces we have:

16.21 Theorem. *If X is a Fréchet space, then the closed convex, balanced hull of every compact subset of X is also compact.*

Proof. If D is the closed unit disc in $\mathbb{C}$ and $K \subset X$ is compact, then $D \cdot K$ is the image of the compact set $D \times K$ under the scalar multiplication map and is, hence, compact. The closed convex hull of $D \cdot K$ will be a closed convex balanced set containing K and, hence, to prove the Theorem it suffices to prove that the closed convex hull of a compact set in a Fréchet space is compact.

Recall that a subset of a complete metric space is compact if and only if it is closed and totally bounded. A subset S of a metric space is totally bounded if for each $\epsilon > 0$ there is a finite set F in S so that each point of S is within ϵ of some point of F.

Choose a translation invariant metric ρ defining the topology of X. If K is a compact subset of the Fréchet space X, then K is totally bounded. Hence, given $\epsilon > 0$ there exists a finite set of points $\{y_i\}_{i=1}^n \subset K$ such that each $x \in K$ is within $\epsilon/2$ of some y_i. Then the convex hull L of the set $\{y_i\}$ is the image of the map

$$(s_1, \cdots, s_n) \to \sum s_i y_i : S \to X$$

where S is the simplex $\{(s_1, \cdots, s_n) \in (R^+)^n : \sum s_i = 1\}$. Thus, L is compact. Furthermore, every element of the convex hull of K is within $\epsilon/2$ of a point of L. Since L is compact, it is also totally bounded and we may find a finite set of points $\{x_j\}$ such that every point of L is within $\epsilon/2$ of some x_j. It follows that every point of the convex hull of K is within ϵ of some x_j. Thus, we have proved that the convex hull of K is totally bounded. The closure of a totally bounded set is clearly totally bounded as well. Thus, the closed convex hull of a compact subset of a Fréchet space is totally bounded and, hence, compact.

The above theorem implies that the family c of sets with compact closure in a Fréchet space is a saturated family. The topology on X^* which it determines is the topology of uniform convergence on compact subsets of X. The space X^* with this topology will be denoted X^*_c.

Our final preliminary result before proving Schwartz's Theorem is the following theorem of Mackey-Arens. Note that, for any saturated family κ covering X, each element $x \in X$ determines a continuous linear functional on X^*_κ by $f \to f(x) : X^*_\kappa \to \mathbb{C}$. Thus, X embedds in the second dual $(X^*_\kappa)^*$. In the case where the sets in κ are precompact, every continuous linear functional on X^* has this form.

16.22 Mackey-Arens Theorem. *If X is a locally convex topological vector space and κ is a saturated family of precompact subsets of X which covers X, then every continuous linear functional on X^*_κ is determined by an element of X. Thus, $X = (X^*_\kappa)^*$.*

Proof. If λ is a continuous linear functional on X^*_κ then there is a neighborhood V of zero in X^*_κ such that $|\lambda(f)| < 1$ for all $f \in V$. We may assume that the neighborhood V has the form

$$V = \{f \in X^*_c : \sup_{x \in K} |f(x)| < 1\}$$

for some $K \in \kappa$. Without loss of generality we may assume K is compact convex and bounded since κ is saturated and consists of precompact sets. We regard X as embedded in $(X^*_\kappa)^*$ and we give the later space the weak-* topology – that is, the topology of pointwise convergence as functions on X^*_κ. This may also be described as the topology of uniform convergence on the family σ of sets which are convex balanced hulls of finite sets. On the image of X in $(X^*_c)^*$ the σ topology is weaker than the original topology on X and, hence, K is also compact in this topology. If we knew the result of the theorem in the case of the σ topology, then we could conclude that every continuous linear functional on $(X^*_c)^*$ with the σ topology is given by an element of X^*_c. Let us assume this for the moment. If λ is not an element of K, then it follows from the Hahn-Banach theorem (TVS 5(b)) that there exists a continuous linear functional f on $(X^*_c)^*_\sigma$ such that $|f(\lambda)| > 1$ and $|f(x)| < r < 1$

for all $x \in K$. However, under our assumption about the σ topology, we must have that $f \in X_c^*$ and it then follows that $f \in V$. This is a contradiction, since $|f(\lambda)| = |\lambda(f)| < 1$ if $f \in V$. We conclude that λ does belong to K and, in particular, λ belongs to X.

It remains to prove the theorem in the case where κ is the family σ generated by finite sets. In this case, let λ and V and K be as above. Since, $\kappa = \sigma$ we have that K is the convex balanced hull of a finite set $\{x_i\}_{i=1}^n$. Then the set of $f \in X^*$ such that $f(x_i) = 0$ for $i = 1, \dots, n$ is just $\cap_k k^{-1}V$ and so it follows that λ vanishes on this set. It then follows from elementary linear algebra that λ must be a linear combination of the x_i and, hence, must belong to X.

16.23 Theorem (Schwartz). *Let X and Y be Fréchet spaces and suppose that $A : X \to Y$ is a surjective continuous linear map and $C : X \to Y$ is a compact continuous linear map. Then $B = A + C$ has closed image and finite dimensional cokernel.*

Proof. We consider the duals X_c^* and Y_c^* of X and Y in the topology of uniform convergence on compact subsets. Then the continuous linear map $A : X \to Y$ has a dual $A^* : Y_c^* \to X_c^*$ defined by

$$A^*(f)(x) = f(A(x))$$

Since A is continuous, it maps compact sets to compact sets, from which it follows that A^* is continuous. Similarly, C^* is continuous.

Claim 1. *The linear map $C^* : Y_c^* \to X_c^*$ is compact.*

Since C is compact, there is a zero neighborhood U in X such that $C(U)$ has compact closure K in Y. Then the zero neighborhood $V_K = \{f \in Y_c^* : \sup_{x \in K} |f(x)| < 1\}$ has the property that $C^*(V_K)$ is a family of continuous functions which is uniformly bounded by one in modulus on U. If L is any compact subset of X then $C^*(V_K)$ is uniformly bounded on L, since $L \subset kU$ for some $k > 0$, and $C^*(V_K)$ is equicontinuous on L, since $|f(x) - f(x')| < \epsilon$ if $x - x' \in \epsilon U$ and $f \in C^*(V_K)$. It follows from the Ascoli-Arzela Theorem that $C^*(V_K)$ has compact closure in the space of all functions on X in the topology of uniform convergence on compact subsets of X. However, the space of continuous linear functionals is clearly closed in this topology since X is a metric space. Thus, $C^*(V_K)$ has compact closure in X_c^* and C^* is a compact operator.

Claim 2. *The linear map $A^* : Y_c^* \to X_c^*$ has closed image and is a topological isomorphism onto its image.*

Let $\{f_\alpha\}$ be a net in Y_c^* such that $A^*(f_\alpha)$ converges in X_c^* to g. Let K be a compact subset of Y. The fact that A is an open map implies that K is the image under A of a compact subset L of X (Problem 16.8). We have that $A^*(f_\alpha) = f_\alpha \circ A$ converges to g uniformly on L and this implies that $f\alpha$ converges uniformly on K. Since this is true for every compact set $K \subset Y$, the net $\{f_\alpha)$ converges in the topology of Y_c^* to an element f. Clearly $A^*(f) = g$. Hence, the image of A^* is closed and, on the image, the inverse map is continuous. This establishes Claim 2.

We now have the hypotheses of Theorem 16.19 satisfied for the pair of operators A^* and C^*. We conclude that $B^* = A^* + C^*$ has finite dimensional kernel Z and closed image and induces a toplogical isomorphism from Y_c^*/Z to its image in X_c^*. Now the map $Y_c^*/Z \to X_c^*$

induced by B^* is the dual of the continuous linear map B' obtained by composing B with the map $Y \to Z^*$ defined by $y \to \{f \to f(y)\} \quad \forall f \in Z, y \in Y$. Thus, to finish the proof we need to show that B' is surjective. In other words, the proof will be complete if we can prove:

Claim 3. *If $B : X \to Y$ is a bounded linear map between Fréchet spaces and $B^* : Y_c^* \to X_c^*$ is injective, has closed image and is a topological isomorphism onto its image, then B is surjective.*

The map B induces a continuous linear map $X/\ker B \to Y$ which has as dual the map B^* as a map from Y_c^* to $\{f \in X_c^* : f(x) = 0 \quad \forall x \in \ker B\}$. This map has closed image and so if it is not surjective, then the Hahn-Banach Theorem implies that there is a continuous linear functional F on X^* which vanishes identically on $B^*(Y_c^*)$ and does not vanish identically on $\{f \in X^* : f(x) = 0 \quad \forall x \in \ker B\}$. However, by the Mackey-Arens Theorem (Theorem 16.22), the functional F has to have the form $F(f) = f(x)$ for some $x \in X$. In other words, there is an $x \in X$ such that $B^*(g)(x) = g(B(x))$ for all $f \in Y_c^*$ but $g(x) \neq 0$ for some $f \in X_c^*$ which vanishes on $\ker B$. This is impossible, since such an x would necessarily be in $\ker B$. Thus, B^* is surjective as a map from Y_c^* to the dual of $X/\ker B$. It follows that we may assume without loss of generality that B is injective and B^* is surjective.

Thus, suppose that $\{x_n\}$ is a sequence in x such that $\{B(x_n)\}$ converges to $y \in Y$. Then the set $S = \{B(x_n)\} \cup \{y\}$ is compact. Hence, the set

$$U = \{f \in Y^* : |f(y)| < 1 \quad \forall y \in S\} \subset \{f \in Y^* : |f(B(x_n))| < 1 \quad \forall n\}$$

is a neighborhood of zero in Y_c^* and, hence, its image $B^*(U)$ is a neighborhood of zero in in X_c^*. This implies that there is a compact set $K \subset X$ and a $\delta > 0$ such that

$$\{f \in X^* : |f(x)| < \delta \quad \forall \quad x \in K\} \subset B^*(U) = \{f \in X_c^* : |f(x_n)| < 1 \quad \forall \quad x \in K\}$$

or

$$\{f \in X^* : |f(x)| < 1 \quad \forall \quad x \in \delta^{-1}K\} \subset B^*(U) = \{f \in X_c^* : |f(x_n)| < 1 \quad \forall \quad x \in K\}$$

From the Hahn-Banach theorem (TVS 5(b)), it now follows that the sequence $\{x_n\}$ lies in the compact convex balanced hull of $\delta^{-1}K$ and, hence, lies in a compact subset of X. Then it has a cluster point x and, clearly, $B(x) = y$. Thus, B is surjective. This completes the proof of claim 3 and the proof of the Theorem.

16. Problems

1. Prove that a bounded linear map $\phi : X \to Y$ between two Fréchet spaces is an open map if and only if every convergent sequence in Y lifts under ϕ to a convergent sequence in X.
2. Prove Corollary 16.9.

3. Prove Corollary 16.10.
4. Prove Corollary 16.11.
5. Prove Corollary 16.12.
6. Prove Corollary 16.13.
7. Prove that if the line bundle corresponding to an element of $H^1(X, \mathcal{H}^*)$ has a non-zero section then the element has image zero in $H^1(X, \mathcal{M}^*)$.
8. Prove that if $X \to Y$ is a surjective continuous linear map between Fréchet spaces, then each compact subset of Y is the image of a compact subset of X, Hint: Use the fact that in a complete metric space the compact sets are the closed totally bounded sets.

17. The Borel-Weil-Bott Theorem

In this chapter we give an application in group representation theory of some of the machinery we have developed in previous chapters. The application is a proof of the Borel-Weil-Bott Theorem. This theorem characterizes the cohomology of equivariant holomorphic line bundles on the flag manifold in terms of the irreducible finite dimensional representations of the corresponding semisimple Lie algebra. The proof we give here is due to Milicic. It is not terribly difficult but it does require detailed knowledge of the structure theory for complex semisimple Lie algebras and their finite dimensional representations. We begin with a quick review of this material.

Let $\mathfrak{g}$ be a complex Lie algebra. Associated to $\mathfrak{g}$ is a unique connected, simply connected complex Lie group G having $\mathfrak{g}$ as its Lie algebra. A complex Lie group is a complex manifold and, hence, we may ask that a homomorphism of it into another complex Lie group be holomorphic rather than just continuous and, in particular, we may ask that finite dimensional representations be holomorphic. Any finite dimensional complex linear representation of $\mathfrak{g}$ gives rise to a holomorphic representation of G through exponentiation. Conversely, any holomorphic finite dimensional representation of G gives rise to a complex linear finite dimensional representation of $\mathfrak{g}$ through differentiation. The adjoint action of $\mathfrak{g}$ on itself, denoted by ad, is such a representation of $\mathfrak{g}$. That is, for $\xi, \eta \in \mathfrak{g}$, $\mathrm{ad}_\xi(\eta) = [\xi, \eta]$. The corresponding representation of the group G will be denoted by Ad. It is a holomorphic representation of G as a group of Lie algebra automorphisms of $\mathfrak{g}$.

The *Killing form* for a Lie algebra $\mathfrak{g}$ is the symmetric bilinear form

$$\langle \xi, \eta \rangle = \mathrm{tr}(\mathrm{ad}_\xi \, \mathrm{ad}_\eta)$$

A Lie algebra is said to be semisimple if it is a direct product of simple Lie algebras – that is, non-abelian Lie algebras with no non-trivial ideals. A Lie algebra is semisimple if and only if its Killing form is non-singular.

17.1 Definition. *If $\mathfrak{g}$ is a semisimple Lie algebra, then*

(a) *a Borel subalgebra $\mathfrak{b}$ of $\mathfrak{g}$ is a maximal solvable subalgebra of $\mathfrak{g}$;*
(b) *a Cartan subalgebra $\mathfrak{h}$ is a maximal abelian subalgebra of $\mathfrak{g}$ consisting of elements which act semisimply on $\mathfrak{g}$.*

A semisimple Lie algebra $\mathfrak{g}$ has a Borel subalgebra $\mathfrak{b}$ and every Borel subalgebra $\mathfrak{b}$ contains a Cartan subalgebra $\mathfrak{h}$. Furthermore, if $\mathfrak{n} = [\mathfrak{b}, \mathfrak{b}]$, then $\mathfrak{n}$ is a nilpotent ideal of $\mathfrak{b}$ and, as vector spaces, $\mathfrak{b} = \mathfrak{h} \oplus \mathfrak{n}$.

If $\mathfrak{h}$ is a Cartan subalgebra, then its elements act semisimply on $\mathfrak{g}$, which implies that $\mathfrak{g}$ decomposes as a direct sum of subspaces $\mathfrak{g}_\alpha$ where $\alpha \in \mathfrak{h}^*$ and

$$\mathfrak{g}_\alpha = \{\xi \in \mathfrak{g} : \mathrm{ad}_\eta(\xi) = \alpha(\eta)\xi \quad \forall \eta \in \mathfrak{h}\}.$$

Since $\mathfrak{h}$ is maximal abelian, the space $\mathfrak{g}_0$ is $\mathfrak{h}$ itself. The non-zero elements $\alpha \in \mathfrak{h}^*$ for which $\mathfrak{g}_\alpha$ is non-empty are called *roots*. We denote the set of all roots by Δ. Thus,

$$\mathfrak{g} = \mathfrak{h} \oplus \sum_{\alpha \in \Delta} \mathfrak{g}_\alpha.$$

A simple calculation shows that

$$[\mathfrak{g}_\alpha, \mathfrak{g}_\beta] \subset \mathfrak{g}_{\alpha+\beta}.$$

One can also show that if α and β are roots then $[\mathfrak{g}_\alpha, \mathfrak{g}_\beta] \neq 0$ if and only if $\alpha+\beta$ is either zero or is a root. It turns out that for each root α the spaces $\mathfrak{g}_\alpha$ are one dimensional. It follows that $[\mathfrak{g}_\alpha, \mathfrak{g}_\beta] = \mathfrak{g}_{\alpha+\beta}$ if $\alpha+\beta$ is a root and that $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ is a one-dimensional subspace of $\mathfrak{h}$. In fact, $\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}$ and $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ span a three dimensional Lie algebra isomorphic to $\mathrm{sl}_2(\mathbb{C})$. Proofs of these facts can be found in any elementary treatment of Lie algebra theory.

The fact that nilpotent matrices have trace zero implies that

$$\langle \mathfrak{g}_\alpha, \mathfrak{g}_\beta \rangle = 0 \quad \text{if} \quad \alpha \neq -\beta.$$

Since, $\mathfrak{h} = \mathfrak{g}_0$, this implies that $\mathfrak{h}$ is othogonal to each $\mathfrak{g}_\alpha$ for α a root. This in turn implies that the Killing form remains non-singular when restricted to $\mathfrak{h}$. Thus, the Killing form induces an isomorphism between $\mathfrak{h}$ and its dual $\mathfrak{h}^*$. That is, to each $\lambda \in \mathfrak{h}^*$ there is a unique $\eta_\lambda \in \mathfrak{h}$ such that $\lambda(\xi) = \langle \xi, \eta_\lambda \rangle$. Then $\langle \lambda, \mu \rangle = \langle \eta_\lambda, \eta_\mu \rangle$ defines a non-singular bilinear form on $\mathfrak{h}^*$ which we shall also call the Killing form. If we let $\mathfrak{h}'$ be the real subspace of $\mathfrak{h}^*$ spanned by the set of roots, then one can prove that $\mathfrak{h}'$ is a real form of $\mathfrak{h}^*$ – i. e. is a real subspace with the property that $\mathfrak{h}^* = \mathfrak{h}' \oplus i\mathfrak{h}'$. Furthermore, the restriction of the Killing form to $\mathfrak{h}'$ is positive definite. Again, for proofs we refer the reader to any standard treatment of Lie algebra theory.

A positive root system for a Cartan $\mathfrak{h}$ is a subset $\Delta^+ \subset \Delta$ such that: (i) it is closed under addition in the sense that if $\alpha, \beta \in \Delta^+$, then $\alpha+\beta \in \Delta^+$ provided $\alpha+\beta$ is a root; and (ii) for every root $\alpha \in \Delta$, exactly one of α, $-\alpha$ belongs to Δ^+. Such a system may be constructed by choosing a real hyperplane through 0 in $\mathfrak{h}^*$ which does not meet Δ and then letting Δ^+ be the set of all roots on one side of this hyperplane. If Δ^+ is a positive root system, then so is its complement in Δ and this is usually denoted Δ^-. A system of positive roots determines two Borel subalgebras $\mathfrak{b}^+$ and $\mathfrak{b}^-$ by

$$\mathfrak{b}^+ = \mathfrak{h} \oplus \sum_{\alpha \in \Delta^+} \mathfrak{g}_\alpha, \quad \mathfrak{b}^- = \mathfrak{h} \oplus \sum_{\alpha \in \Delta^-} \mathfrak{g}_\alpha,$$

It is easy to see that if $\mathfrak{b}$ is any Borel subalgebra of $\mathfrak{g}$ and $\mathfrak{h}$ is a Cartan contained in $\mathfrak{b}$ then $\mathfrak{b}$ may be represented as the algebra $\mathfrak{b}^+$ if the positive root system Δ^+ is chosen to be the set of roots whose root spaces are contained in $\mathfrak{b}$. One may also represent $\mathfrak{b}$ as $\mathfrak{b}^-$ by choosing the positive root system to be the complement of this set.

A key result in the theory of semisimple Lie groups is the following theorem which we will not prove:

17.2 Theorem. *Under* Ad *the group G acts transitively on the set of all Borel subalgebras and on the set of all Cartan subalgebras.*

Let B be the normalizer in G of the Borel subalgebra $\mathfrak{b}$ – that is,

$$B = \{g \in G : \mathrm{Ad}_g(\mathfrak{b}) \subset \mathfrak{b}\}.$$

Then B is the subgroup of G corresponding to the subalgebra $\mathfrak{b}$ under the Lie correspondence. A subgroup of this form is called a Borel subgroup of G. Note that B is the stabilizer of $\mathfrak{b}$ under the action of G on the set of Borel subalgebras. Thus, the set X of Borel subalgebras of $\mathfrak{g}$ is in one to one correspondence with the quotient space G/B. Another key result from Lie theory which we will not prove is the following:

17.3 Theorem. *If B is a Borel subgroup of a complex semisimple Lie group G, then G/B is a compact complex manifold.*

Thus, the set X of Borel subalgebras of $\mathfrak{g}$ may be given the structure of a compact complex manifold through its identification with G/B. We call this the *flag manifold* of G. Note also that G acts on X – through left multiplication if X is represented as the coset space G/B or, equivalently, through Ad if X is represented as the set of Borel subalgebras of $\mathfrak{g}$. This action is, of course, transitive and is holomorphic in the sense that $g \times x \to gx : G \times X \to X$ is a holomorphic map.

We are interested in studying the finite dimensional representations of $\mathfrak{g}$. Let (π, V) be such a representation. By Lie theory, π exponentiates to a holomorphic representation of the Lie group G on V which we also denote by π. If $\mathfrak{h}$ is a Cartan subalgebra of $\mathfrak{g}$, and H is the connected Lie subgroup of G corresponding to $\mathfrak{h}$, then H is a connected complex abelian Lie group; hence, it must be a complex torus – that is, a product of copies of the punctured plane $\mathbb{C}^*$. The corresponding product of unit circles is a maximal compact subgroup T of H. The complex vector space V may be given a T-invariant inner product (by integrating any inner product over the compact group T). Then (π, V) is a finite dimensional unitary representation of T. It follows that V may be written as a direct sum of subspaces on which the representation of T is irreducible and each of these subspaces consists of eigenvectors for the operators $\pi(t)$ for $t \in T$. It follows that they are also eigenvectors for the operators $\pi(\xi)$ for $\xi \in \mathfrak{h}$. In other words,

17.4 Theorem. *The restriction of the representation π to $\mathfrak{h}$ is completely reducible.*

From this it follows that V is a direct sum of subspaces V_λ where $\lambda \in \mathfrak{h}^*$ and

$$V_\lambda = \{v \in V : \pi(\xi)v = \lambda(\xi)v \quad \forall \xi \in \mathfrak{h}\}$$

The elements $\lambda \in \mathfrak{h}^*$ for which $V_\lambda \neq 0$ are called the *weights* of the representation π. Since the representation is finite dimensional, there can be only finitely many of them. How does the rest of the Lie algebra $\mathfrak{g}$ act on V? We get an idea by using the root space decomposition. Thus if $\alpha \in \Delta$ and $\xi \in \mathfrak{g}_\alpha$, then a calculation shows that $\pi(\xi)V_\lambda \subset V_{\lambda+\alpha}$. Now, if the representation π is irreducible, then for any weight λ the weight space V_λ must generate V. It follows that one must be able to obtain all the weights from a given one by successively adding roots.

17.5 Definition. *For a finite dimensional representation (π, V), a Cartan $\mathfrak{h}$, a positive root system Δ^+ and weights λ and μ, we say $\mu < \lambda$ if $\lambda - \mu$ is a sum of positive roots. A weight is called a highest weight if it is maximal relative to this relation. If a weight λ has the property that, for each root α, either $\lambda + \alpha$ or $\lambda - \alpha$ is not a weight for (π, V), then λ will be called an extremal weight.*

If λ is a highest weight, then λ is an extremal weight since $\lambda + \alpha$ fails to be a weight for every positive root α. We also have $\pi(g_\alpha)V_\lambda = 0$ for every $\alpha \in \Delta^+$ since $V_{\lambda+\alpha} = 0$. It follows from this and the commutation relations in $\mathfrak{g}$ that the $\mathfrak{g}$-invariant subspace of V generated by any non-zero vector $v \in V_\lambda$ is contained in the span of v and the spaces V_μ for $\mu < \lambda$. This implies the following theorem:

17.6 Theorem. *For an irreducible finite dimensional representation (π, V) of $\mathfrak{g}$, a Cartan subalgebra $\mathfrak{h}$ and a system of positive roots Δ^+, there is a unique highest weight $\lambda \in \mathfrak{h}^*$. Furthermore, V_λ is one dimensional.*

Note that this also implies that, for an irreducible finite dimensional representation there is exactly one extremal weight corresponding to each positive root system.

As before, let $\mathfrak{h}'$ be the real subspace of $\mathfrak{h}^*$ spanned by the roots. Equivalently, $\mathfrak{h}'$ is the real subspace of $\mathfrak{h}^*$ consisting of those weights λ for which $\langle \lambda, \alpha \rangle$ is real for every $\alpha \in \Delta^+$. Each $\alpha \in \Delta^+$ defines a hyperplane $\{\mu \in \mathfrak{h}' : \langle \mu, \alpha \rangle = 0\}$ in $\mathfrak{h}'$. The complement in $\mathfrak{h}'$ of the union of these hyperplanes has finitely many components and these are open sets called *Weyl chambers*. That is, a Weyl chamber is a non-empty subset of $\mathfrak{h}'$ of the form $\{\lambda \in \mathfrak{h}' : \epsilon(\alpha)\langle \lambda, \alpha \rangle > 0 \quad \forall \alpha \in \Delta^+\}$, where ϵ is a function from Δ^+ to $\{1, -1\}$ (not every such function defines a non-empty set). The positive Weyl chamber is $\{\lambda \in \mathfrak{h}' : \langle \lambda, \alpha \rangle > 0 \quad \forall \alpha \in \Delta^+\}$.

For $\alpha \in \Delta^+$, the operator s_α on $\mathfrak{h}'$ defined by

$$s_\alpha(\lambda) = \lambda - 2\frac{\langle \alpha, \lambda \rangle}{\langle \alpha, \alpha \rangle}\alpha$$

is reflection through the hyperplane determined by α. The Weyl group W for $\mathfrak{g}$ is the group of transformations of $\mathfrak{h}'$ generated by the set of reflections s_α.

It is clear that the Weyl chambers are in one to one correspondence with the possible choices of positive root systems. That is, for each Weyl chamber there is exactly one positive root system for which it is the positive Weyl chamber. It is also easy to prove that the Weyl group acts transitively on the set of Weyl chambers.

We define a norm in $\mathfrak{h}'$ by setting

$$||\lambda|| = \sqrt{\langle \lambda, \lambda \rangle}.$$

It turns out, though we shall not prove it here, that each element of the Weyl group arises from a transformation of $\mathfrak{h}$ of the form Ad_g, where $g \in G$ is in the normalizer of $\mathfrak{h}$. In fact W is isomorphic to the normalizer of $\mathfrak{h}$ in G mod the centralizer of $\mathfrak{h}$ in G. It follows from this and the G invariance of the Killing form, that the Weyl group acts as a group of isometries relative to the above norm.

We set

$$\rho = \frac{1}{2}\sum_{\alpha \in \Delta^+} \alpha$$

This element plays a special role throughout the theory of semisimple Lie algebras.

17.7 Theorem. *Let (π, V) be a finite dimensional irreducible representation of $\mathfrak{g}$, $\mathfrak{h} \subset \mathfrak{g}$ a Cartan subalgebra, Δ^+ a positive root system, $\Lambda \subset \mathfrak{h}^*$ the set of weights for (π, V) and $\lambda \in \Lambda$ the highest weight. Then*

(a) $\dfrac{2\langle \mu, \alpha \rangle}{\langle \alpha, \alpha \rangle} \in \mathbb{Z} \quad \forall \alpha \in \Delta, \mu \in \Lambda$;

(b) *Λ is closed under the action of the Weyl group;*

(c) *if* $\alpha \in \Delta^+$ *then* $\langle \lambda, \alpha \rangle \geq 0$ *and* $\langle \rho, \alpha \rangle > 0$*;*
(d) $||\mu|| \leq ||\lambda|| \quad \forall \mu \in \Lambda$*;*
(e) $\mu \in \Lambda$ *is extremal if and only if* $||\mu|| = ||\lambda||$ *and the Weyl group acts transitively on the set of extremal weights;*
(f) $||\mu + \rho|| < ||\lambda + \rho||$ *for every* $\mu \in \Lambda$ *distinct from* λ*.*

Proof. For a given root $\alpha \in \Delta$, consider the subalgebra $\mathfrak{l}_\alpha$ spanned by $\mathfrak{g}_\alpha$, $\mathfrak{g}_{-\alpha}$ and $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$. This is a copy of $\mathrm{sl}_2(\mathbb{C})$. If $\mu \in \Lambda$, the spaces $V_{\mu - n\alpha}$, for n an integer, span a subspace of V which is a finite dimensional representation of $\mathfrak{l}_\alpha$. However, we know that the weights for a finite dimensional irreducible representation of $sl_2(\mathbb{C})$ are symmetric about the origin. Since every finite dimensional representation decomposes as a direct sum of irreducibles, the weights for any finite dimensional representation of $sl_2(\mathbb{C})$ are symmetric about the origin. A calculation shows that an element of the Cartan subalgebra $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ of our copy of $sl_2(\mathbb{C})$ acts on each space V_ν as a fixed scalar multiple of $\langle \nu, \alpha \rangle$. It follows from symmetry about the origin that $\langle \lambda - m\alpha, \alpha \rangle = -\langle \lambda, \alpha \rangle$ for some integer m and, from this, that

$$\frac{2\langle \lambda, \alpha \rangle}{\langle \alpha, \alpha \rangle} = m.$$

This completes the proof of part (a). It also proves part (b) since it implies that the weight $\lambda - m\alpha$ is $s_\alpha(\lambda)$, from which it follows that the set of weights is closed under the action of the Weyl group. The first inequality of part (c) also follows from this argument applied to the case where $\mu = \lambda$, since then only non-negative integers n yield non-zero subspaces $V_{\lambda - n\alpha}$. Thus, the integer m above must be non-negative in this case, which implies that $\langle \lambda, \alpha \rangle \geq 0$. The proof of the second inequality in part (c) is more complicated. It requires a development of the properties of systems of simple roots. We leave this development and the resulting proof of the fact that ρ is in the positive chamber to the exercises (Problems 4 – 9).

If μ is any weight in Λ, it has the form $\mu = \lambda - \nu$ where ν is a sum of positive roots. If μ also satisfies $\langle \mu, \alpha \rangle \geq 0$ for all $\alpha \in \Delta^+$, then we have

$$\langle \mu, \mu \rangle = \langle \lambda, \lambda \rangle - \langle \nu, \lambda \rangle - \langle \mu, \nu \rangle \leq \langle \lambda, \lambda \rangle$$

Thus, $||\mu|| \leq ||\lambda||$ for every weight μ in the closure of the positive chamber. However, this implies $||\mu|| \leq ||\lambda||$ for all $\mu \in \Lambda$ since every such μ may be brought into the closure of the positive chamber by applying a Weyl group transformation and Weyl group tranformations are isometries. This completes the proof of part (d);

Suppose that $\mu \in \Lambda$ is not an extremal weight. Then there is a root α such that $\mu + \alpha$ and $\mu - \alpha$ are both roots. Then

$$||\mu||^2 + ||\alpha||^2 = \frac{1}{2}(||\mu + \alpha||^2 + ||\mu - \alpha||^2) \leq ||\lambda||^2$$

by part (d) and this implies that $||\mu|| < ||\lambda||$. Thus, only the extremal weights can have norm equal to $||\lambda||$. The proof of part (e) will be complete if we can show that the Weyl group acts transitively on the set of extremal weights, since this will imply that they all have norm equal to that of λ. Using a Weyl group transformation, we can move any extremal weight into the closure of the positive chamber and it will still be an extremal weight. Thus, we will have completed the proof of (e) if we show that λ is the only extremal root in the closure of the positive chamber. Thus, suppose that μ is extremal and lies in the closure of the positive chamber. Then for each $\alpha \in \Delta^+$ we have $\langle \mu, \alpha \rangle \geq 0$ and either $\mu + \alpha$ or $\mu - \alpha$ is not a root. However, the $sl_2(\mathbb{C})$ argument of the first paragraph shows that it must be $\mu + \alpha$ that fails to be a root if $\langle \mu, \alpha \rangle \geq 0$. Hence, μ is the highest root λ and the proof of part(e) is complete.

If μ is any weight in Λ then

$$\begin{aligned}\langle \mu + \rho, \mu + \rho \rangle &= \langle \mu, \mu \rangle + 2\langle \mu, \rho \rangle + \langle \rho, \rho \rangle \leq \langle \lambda, \lambda \rangle + 2\langle \mu, \rho \rangle + \langle \rho, \rho \rangle \\ &= \langle \lambda + \rho, \lambda + \rho \rangle - 2\langle \lambda - \mu, \rho \rangle < \langle \lambda + \rho, \lambda + \rho \rangle\end{aligned}$$

by part (c) and the fact that $\lambda - \mu$ is a sum of positive roots. This proves part (f).

The elements of $\mathfrak{h}^*$ that satisfy the condition in part(a) of the above theorem are called *integral weights.*

17.8 Definition. *We say that a weight $\lambda \in \mathfrak{h}'$ is dominant relative to a system of positive roots if $\langle \alpha, \lambda + \rho \rangle > 0$ for every positive root α. Thus, a weight λ is dominant if and only if $\lambda + \rho$ belongs to the positive Weyl chamber.*

We know that for every finite dimensional irreducible representation there is a unique highest weight and it is easy to see from part(c) of Theorem 17.7 that highest weights are dominant. In fact, using the theory of Verma modules one can prove that the finite dimensional irreducible representations of a complex semisimple Lie algebra are classified by their highest weights. We won't prove it here, but the theorem that does this is the following:

17.9 Theorem. *Let $\mathfrak{h} \subset \mathfrak{g}$ be a Cartan subalgebra and $\Delta^+ \subset \mathfrak{h}^*$ a system of positive roots. Then each dominant integral weight λ is the highest weight for a unique (up to isomorphism) finite dimensional irreducible representation of $\mathfrak{g}$.*

We now introduce the objects of study in the Borel-Weil-Bott Theorem – the G-equivariant holomorphic line bundles. Each of these is constructed from an integral weight in $\mathfrak{h}^*$ by an induction process which we shall describe below. In what follows, $\mathcal{H}$ will denote the sheaf of holomorphic functions on either G or X. If it is not clear from the context which is meant, we shall use the notation ${}_G\mathcal{H}$ or ${}_X\mathcal{H}$. We also fix, for the remainder of the discussion, a Cartan subalgebra $\mathfrak{h}$, a system of positive roots Δ^+ and the Borel subalgebra $\mathfrak{b} = \mathfrak{b}^-$ spanned by $\mathfrak{h}$ and the negative root spaces. The corresponding Borel subgroup of G will be denoted B. We represent the flag manifold X as G/B.

Let (σ, W) be a finite dimensional holomorphic representation of B – equivalently, a finite dimensional complex linear representation of $\mathfrak{b}$. The space $G \times_B W$ is constructed from $G \times W$ by identifying points which lie in the same orbit of the B-action described by $b \times (g \times w) \to gb^{-1} \times \sigma(b)w$. Projection on the first factor of $G \times W$, followed by the projection $\rho : G \to G/B$, induces a well defined holomorphic projection $G \times_B W \to X$. The inverse image of the typical point under this projection is $gB \times_B W$, which has a well defined vector space structure under which it is isomorphic to W. In other words, $G \times_B W$ is a holomorphic vector bundle over X with fiber W. Furthermore, G acts on $G \times_B W$ through action on the left in the first factor. By a G-equivariant holomorphic vector bundle over X we mean a complex manifold M with a holomorphic G action $G \times M \to M$, a holomorphic projection $\gamma : M \to X$ which commutes with the G actions on M and X and a complex vector space structure on $M_x = \gamma^{-1}(x)$ for each $x \in X$ such that $m \to gm : M_x \to M_{gx}$ is linear for each $g \in G$ and $x \in X$. In this sense $G \times_B W$ is a G-equivariant holomorphic vector bundle over X.

17.10 Definition. *Let (σ, W) be a finite dimensional holomorphic representation of a Borel subgroup $B \subset G$, then we will call $G \times_B W$ the induced bundle over X determined by (σ, W) and denote it by $\mathrm{I}(\sigma)$. The sheaf of holomorphic sections of this vector bundle will be denoted $\mathcal{I}(\sigma)$.*

Note that, from the construction of $\mathrm{I}(\sigma)$, we may characterize the sheaf $\mathcal{I}(\sigma)$ in the following way:

$$\mathcal{I}(\sigma)(U) = \{f \in \mathcal{H}(\rho^{-1}(U), W) : f(gb^{-1}) = \sigma(b)f(g) \quad \forall b \in B\}$$

where $\rho : G \to G/B = X$ is the projection and $\mathcal{H}(\rho^{-1}(U), W)$ denotes the space of holomorphic W valued functions on $\rho^{-1}(U)$.

If (π, V) is a finite dimensional holomorphic representation of G then we may speak of the trivial G equivariant holomorphic vector bundle over X with fiber V. This is the complex manifold $X \times V$ with projection $X \times V \to X$ just the projection on the first factor and G action given by $g \times (x \times v) \to gx \times \pi(g)v$. Note that if we consider the restriction σ of π to B, then we have a representation of B and we may consider its induced bundle $\mathrm{I}(\sigma) = X \times_B V$. In fact the map

$$g \times v \to \rho(g) \times \pi(g)v : G \times V \to X \times V$$

induces a G-equivariant vector bundle isomorphism $G \times_B V \to X \times V$. Thus, the trivial G-equivariant bundle $X \times V$ and the induced bundle $\mathrm{I}(\sigma)$ are isomorphic. This is part of what is to be shown in the next theorem.

17.11 Theorem. *The induction functor $(\sigma, W) \to \mathrm{I}(\sigma)$ is an equivalence of categories from the category of finite dimensional holomorphic representations of B to the category of G-equivariant holomorphic vector bundles over X. Furthermore, the bundle $\mathrm{I}(\sigma)$ is trivial if and only if the representation σ is the restriction to B of a holomorphic representation of G.*

Proof. Let $x_0 = B/B$ be the point of X corresponding to B. Then B is the stabilizer of x_0. If $M \to X$ is any G-equivariant vector bundle over X, then the fiber M_{x_0} over x_0 is invariant under B and, hence, is a finite dimensional representation of B. The resulting correspondence $M \to M_{x_0}$ is a functor from G-equivariant vector bundles over X to finite dimensional holomorphic representations of B. It is clear from the definitions that $(\sigma, W) \to \mathrm{I}(\sigma)$ followed by $M \to M_{x_0}$ is isomorphic to the identity. On the other hand, if M is a G-equivariant holomorphic vector bundle, then we define a map $G \times M_{x_0} \to M$

by $g \times m \to gm$. This clearly induces a map $G \times_B M_{x_0} \to M$ which is an isomorphism of G-equivariant vector bundles. Thus, the composition $M \to M_{x_0}$ followed by $(\sigma, W) \to \mathrm{I}(\sigma)$ is also isomorphic to the identity. This completes the proof of the first statement of the theorem. We have already proved the second statement in one direction. We leave the other direction as as exercise (Problem 17.1).

Note that, for any finite dimensional holomorphic representation (σ, W) of B, the sheaf of holomorphic sections $\mathcal{I}(\sigma)$ of $\mathrm{I}(\sigma)$ is a coherent analytic sheaf. It follows from the Cartan-Serre Theorem (Theorem 16.18) that the cohomology $H^p(X, \mathcal{I}(\sigma))$ is finite dimensional for each p. Because of the G-equivariance, G acts in a holomorphic fashion on $H^0(X, \mathcal{I}(\sigma))$. It is also true, but less obvious, that G acts in a holomorphic fashion on each cohomology space $H^p(X, \mathcal{I}(\sigma))$ (Problem 17.2). Of course, it is not at all clear at this point which of these cohomology groups, if any, are non-vanishing.

We now turn to the special case of induction of one dimensional representations of B. Each integral $\lambda \in \mathfrak{h}^*$ determines a character e^λ on H by $\mathrm{e}^\lambda(\exp(\xi)) = \exp(\lambda(\xi))$. We can extend the character e^λ on H to a character (which we also call e^λ) on B by letting it be the identity on N. Thus, e^λ defines a holomorphic representation of B on the one dimensional vector space $\mathbb{C}$. Each one dimensional holomorphic representation of B arises in this way. If we apply the induction functor, this yields a G-equivariant holomorphic line bundle $\mathrm{I}(\mathrm{e}^\lambda)$. We denote the sheaf of holomorphic sections of this line bundle by $\mathcal{H}_\lambda$. Thus, if U is an open subset of X and $\rho : G \to G/B = X$ is the projection, then

$$\mathcal{H}_\lambda(U) = \{f \in \mathcal{H}(\rho^{-1}(U)) : f(gb^{-1}) = \mathrm{e}^\lambda(b)f(g)\}$$

This raises several questions:

(1) For which integral weights λ and for which integers p is $H^p(X, \mathcal{H}_\lambda)$ non-zero?
(2) For which integral weights λ and for which integers p is $\mathcal{H}_\lambda(X)$ an irreducible representation of $\mathfrak{g}$?
(3) Does every irreducible finite dimensional representation of $\mathfrak{g}$ arise this way?

The first step in answering these questions is to note a relationship between the sheaves $\mathcal{H}_\lambda$ and the Serre sheaves $\mathcal{H}(k)$ on projective space. This is done in Theorem 17.13 below. First, we prove a technical lemma.

17.12 Lemma. *If $\mathfrak{g}$ is a semisimple Lie algebra, then $\mathfrak{g}$ has a faithful finite dimensional irreducible representation.*

Proof. If $\mathfrak{g}$ is simple (is non-abelian with no non-trivial ideals) then the adjoint representation of $\mathfrak{g}$ on itself is faithful, since the kernel of $g \to \mathrm{ad}_g$ must be 0, and irreducible, since an invariant subspace is an ideal.

Now if $\mathfrak{g}$ is not simple, we express it as a product $\prod \mathfrak{g}_i$ of simple Lie algebras $\mathfrak{g}_i$ and find a faithful irreducible representation (π_i, V_i) for each of them. Then the tensor product $\otimes_i \pi_i$ will be a faithful irreducible representation of $\mathfrak{g} = \prod_i \mathfrak{g}_i$ on $\otimes_i V_i$.

17.13 Theorem. *Let (π, V) be a faithful finite dimensional irreducible representation of $\mathfrak{g}$ with highest weight λ and let $P(V^*)$ be the projective space of the dual V^* of V, then there is a closed holomorphic embedding $\phi : X \to P(V^*)$ such that $\mathcal{H}_\lambda = \phi^{-1}\mathcal{H}(1)$.*

Proof. Since λ is a highest weight, the weight space V_λ is a complement to the span of the subspaces $\pi(\mathfrak{g}_\alpha)W$ for $\alpha \in \Delta^-$. Thus, if W_λ is the space of all elements of V^* which vanish on this span, then there is a non-zero element $w \in W_\lambda$. Let (π^*, V^*) denote the dual of the representation (π, V). Clearly, (π^*, V^*) is also a faithful, irreducible representation. For the corresponding representation of G, we have

$$(\pi^*(g)w)(v) = w(\pi(g^{-1})v) \quad \forall w \in V^*$$

It follows that W_λ is the lowest weight space of (π^*, V^*) and has weight $-\lambda$. Also $\pi^*(\mathfrak{g}_\alpha)W_\lambda = 0$ if $\alpha \in \Delta^-$. On the other hand, $\pi^*(\mathfrak{g}_\alpha)W_\lambda \neq 0$ for each $\alpha \in \Delta^+$. For if $\mathfrak{g}_\alpha$ and $\mathfrak{g}_{-\alpha}$ both annihilate W_λ, then it follows from the commutation relations among the root spaces and from the fact that W_λ generates V^*, that $\mathfrak{g}_\alpha$ and $\mathfrak{g}_{-\alpha}$ both annihilate all of V^*. This is impossible since π^* is a faithful representation. Thus, the Borel subalgebra $\mathfrak{b} = \mathfrak{h} + \sum_{\alpha\in\Delta^-} \mathfrak{g}_\alpha$ is the stabilizer of W_λ in $\mathfrak{g}$. That is,

$$\mathfrak{b} = \{\xi \in \mathfrak{g} : \pi(\xi)W_\lambda \subset W_\lambda\}.$$

It follows from this that the corresponding Borel subgroup B is the stabilizer of the one dimensional subspace W_λ under the action π^* of G on V^*. Thus, under the action of G on the complex projective space $P(V^*)$ of V^* induced by π^*, the point $p_0 \in P(V^*)$ determined by W_λ has B as stabilizer. Hence, the orbit of p_0 under this action of G is compact, hence closed, and is a copy of $X = G/B$. Since the action of G on $P(V^*)$ is holomorphic, it follows from the complex implicit function theorem that the orbit is a compact, complex submanifold of $P(V^*)$. Thus, we have a closed holomorphic embedding $\phi : X \to P(V^*)$.

Now a typical section of the sheaf $\mathcal{H}(1)$ on an open subset of $P(V^*)$ is determined by a homogeneous function of degree one on the corresponding open subset of V^* – that is, by a linear functional on V^*, i. e. a vector $v \in V$. If U is an open subset of X, then the pullback of this section to U may be identified with the function f on $\rho^{-1}(U) \subset G$ defined by $g \to \pi^*(g)w(v)$, where w is a non-zero element of W_λ. This function satifies

$$f(gb^{-1}) = \pi^*(gb^{-1})w(v) = e^\lambda(b)\pi^*(g)w(v) = e^\lambda(b)f(g)$$

and, hence, determines a section of $\mathcal{H}_\lambda$. It follows that $\phi^{-1}\mathcal{H}(1) = \mathcal{H}_\lambda$.

The above result shows why we choose to depart from the standard convention in the theory of group representations which assigns to a Borel $\mathfrak{b}$ the system of positive roots for which $\mathfrak{b} = \mathfrak{b}^+$. Instead, we assign to $\mathfrak{b}$ the system of positive roots for which $\mathfrak{b} = \mathfrak{b}^-$. With this convention, the λ of the above theorem is a highest weight rather than a lowest weight. This is a better convention because then highest weights induce positive line bundles rather than negative line bundles. This is the convention we will use throughout this section.

As an immediate corollary of the above theorem, we have:

17.14 Corollary. *The flag manifold X is a non-singular projective variety.*

The Casimir operator is a particular element of the center $\mathcal{Z}(\mathfrak{g})$ of the enveloping algebra $U(\mathfrak{g})$. It may be described as

$$\Omega = \sum \eta_i^2 + \sum_{\alpha\in\Delta^+} (\xi_\alpha\xi_{-\alpha} + \xi_{-\alpha}\xi_\alpha) = \sum \eta_i^2 + \sum_{\alpha\in\Delta^+} (2\xi_{-\alpha}\xi_\alpha + [\xi_\alpha, \xi_{-\alpha}])$$

where the η_i form a self dual basis for $\mathfrak{h}$ relative to the Killing form, so that $\langle \eta_i, \eta_j \rangle = \delta_{ij}$. The elements $\xi_\alpha \in \mathfrak{g}_\alpha$ are chosen so that $\langle \xi_\alpha, \xi_{-\alpha} \rangle = 1$. For each α the element $\theta_\alpha = [\xi_\alpha, \xi_{-\alpha}]$ belongs to $\mathfrak{h}$ and has the property that $\lambda(\theta_\alpha) = \langle \lambda, \alpha \rangle$ for each $\lambda \in \mathfrak{h}^*$, where $\langle , \rangle$ also stands for the form on $\mathfrak{h}^*$ which is dual to the Killing form. Thus, we have

$$\Omega = \sum \eta_i^2 + \sum_{\alpha \in \Delta^+} \theta_\alpha + 2 \sum_{\alpha \in \Delta^+} \xi_{-\alpha} \xi_\alpha.$$

Let $\rho \in \mathfrak{h}^*$ be one half the sum of the positive roots.

17.15 Theorem. *On a finite dimensional irreducible representation (π, V) of highest weight λ, the Casimir operator acts as the scalar $\langle \lambda, \lambda + 2\rho \rangle$.*

Proof. Since π is irreducible, $\pi(\Omega)$ must be a scalar times the identity operator. We calculate the scalar by applying $\pi(\Omega)$ to a highest weight vector. Thus, let v_λ be a highest weight vector in V, then $\pi(\xi_\alpha) v_\lambda = 0$ for each $\alpha \in \Delta^+$. It follows that

$$\pi(\Omega) v_\lambda = \sum \lambda(\eta_i)^2 + \sum_{\alpha \in \Delta^+} \lambda(\theta_\alpha)$$

From the definition of the dual form on $\mathfrak{h}^*$, it follows that $\sum \lambda(\eta_i)^2 = \langle \lambda, \lambda \rangle$. We also have $\lambda(\theta_\alpha) = \langle \lambda, \alpha \rangle$. Thus,

$$\pi(\Omega) v_\lambda = \langle \lambda, \lambda \rangle + \sum_{\alpha \in \Delta^+} \langle \lambda, \alpha \rangle = \langle \lambda, \lambda \rangle + \langle \lambda, 2\rho \rangle = \langle \lambda, \lambda + 2\rho \rangle$$

This completes the proof.

17.16 Theorem. *On the sheaf $\mathcal{H}_\lambda$, the Casimir acts as the scalar $\langle \lambda, \lambda + 2\rho \rangle$.*

Proof. The action ℓ of $U(\mathfrak{g})$ on $\mathcal{H}_\lambda$ is the infinitesimal form of the action ℓ of G on the bundle $\mathrm{I}(\mathrm{e}^\lambda)$. Thus, let f be a local section of $\mathcal{H}_\lambda$ defined in a neighborhood any point $x \in X$. Let B be Borel subgroup which is the stabilizer of x. Then we may represent X as G/B and x as the identity coset of B. Choose a Cartan $\mathfrak{h} \subset \mathfrak{b}$ and a positive root system such that the corresponding Borel subalgebra $\mathfrak{b}$ is $\mathfrak{b}^-$. Then for $b \in B$

$$\ell(\xi_{-\alpha}\xi_\alpha) f(b) = \frac{d}{dt} \ell(\xi_\alpha) f(\exp(-t\xi_{-\alpha}) b)|_{t=0} = \frac{d}{dt} e^\lambda (b^{-1} \exp(t\xi_{-\alpha}))|_{t=0} \ell(\xi_\alpha) f(e) = 0$$

and

$$\begin{aligned} \ell(\eta) f(b) =& \frac{d}{dt} f(\exp(-t\eta) b)|_{t=0} = \frac{d}{dt} e^\lambda (b^{-1} \exp(t\eta))|_{t=0} f(e) \\ &= \lambda(\eta) e^\lambda (b^{-1}) f(e) = \lambda(\eta) f(b) \end{aligned}$$

for $\eta \in \mathfrak{h}$. That $\ell(\Omega) f(b) = \langle \lambda, \lambda + 2\rho \rangle f(b)$ now follows as in the previous theorem from the identity

$$\Omega = \sum \eta_i^2 + \sum_{\alpha \in \Delta^+} \theta_\alpha + 2 \sum_{\alpha \in \Delta^+} \xi_{-\alpha} \xi_\alpha.$$

This says that the sections $\ell(\Omega)f$ and $\langle\lambda,\lambda+2\rho\rangle f$ of the line bundle corresponding to $\mathcal{H}(\lambda)$ agree at the point x. However, since x was arbitrary, this completes the proof.

Note that

$$\langle\lambda,\lambda+2\rho\rangle = \langle\lambda+\rho,\lambda+\rho\rangle - \langle\rho,\rho\rangle = ||\lambda+\rho||^2 - ||\rho||^2$$

If ω is a weight of the form $\omega = w\rho - \rho$ for some Weyl group element $w \in W$, then

$$||\omega+\rho||^2 - ||\rho||^2 = ||w\rho||^2 - ||\rho||^2 = 0$$

since the Weyl group is a group of isometries of $\mathfrak{h}^*$. Thus, the Casimir acts as the zero operator on $\mathcal{H}_\omega$ for all $\omega \in \{w\rho - \rho : w \in W\}$.

We can now begin the proof of the Borel-Weil-Bott theorem. The proof described here is due to Milicic.

17.17 Lemma. *Let (π, V) be a finite dimensional irreducible representation of $\mathfrak{g}$ with highest weight λ. Let w be an element of W and set $\omega = w\rho - \rho$. Then, for all p,*

(a) *$H^p(X, \mathcal{H}_\omega)$ is a trivial $\mathfrak{g}$ module; and*
(b) *as $\mathfrak{g}$ modules, $H^p(X, \mathcal{H}_{w\lambda+\omega}) \simeq H^p(X, \mathcal{H}_\omega) \otimes V$.*

Proof. By the remarks preceding this theorem, we know that the Casimir acts on $\mathcal{H}_\omega$ as the zero operator. However, the only finite dimensional representations with this property are the trivial representations (this follows from Theorem 17.15, which implies that any highest weight of such a representation is the zero weight). This proves part(a).

We consider the representation corresponding to (π, V) of G on V and denote by σ its restriction to B. Then the induced bundle $\mathrm{I}(\sigma)$ is trivial by problem 17.1 and, in fact, is just the trivial G-equivariant bundle $X \times V$. Thus, its sheaf of sections $\mathcal{I}(\sigma)$ is just $\mathcal{H} \otimes V$. On the other hand, we may *twist* σ by tensoring it with the one dimensional representation determined by the character e^ω to obtain a representation σ_ω of B. Then the corresponding G-equivariant bundle is $\mathrm{I}(\sigma_\omega) = \mathrm{I}(\mathrm{e}^\omega) \otimes \mathrm{I}(\sigma)$. The corresponding sheaf of sections is $\mathcal{I}(\sigma_\omega) = \mathcal{H}_\omega \otimes V$. Here the $\mathfrak{g}$ action is the tensor product of the natural action on $\mathcal{H}_\omega$ with the action on V given by the representation σ. It follows that, as $\mathfrak{g}$-modules,

$$H^p(X, \mathcal{I}(\sigma_\omega)) = H^p(X, \mathcal{H}_\omega) \otimes V.$$

We now have that, as a $\mathfrak{g}$-module, $H^p(X, \mathcal{I}(\sigma_\omega))$ is the tensor product of a trivial $\mathfrak{g}$-module, $H^p(X, \mathcal{H}_\omega)$ and the module (π, V). In particular, this implies that the Casimir acts on $H^p(X, \mathcal{I}(\sigma_\omega))$ as the scalar $\langle\lambda, \lambda+2\rho\rangle$ by Theorem 17.15.

We next construct a filtration $\{V^p\}$ of V by B-submodules. We set $V^0 = V$. Now the highest weight space of V is V_λ and this space is one dimensional. The sum of the weight spaces V_μ for $\mu \neq \lambda$ is a B-submodule V^1 of codimension one in V. It also must have a highest weight ν and a highest weight space V_ν. The sum of the weight spaces in V^1 other than V_ν together with any codimension one subspace of V_ν yields a B-submodule V^2 of codimension one in V^1. Proceeding in this way, one constructs a decreasing filtration $\{V^p\}$ of V by B-submodules such that V^p/V^{p+1} is a one dimensional B-module for each p. Since, on a one dimensional B-module, B necessarily acts by a character e^ν for some $\nu \in \mathfrak{h}^*$, we have a sequence $\{\nu^p\}$ of elements of $\mathfrak{h}^*$ such that B acts as e^{ν^p} on V^p/V^{p+1}.

The elements ν^p are necessarily weights which occur in the representation V and, hence, are weights dominated by λ. When we tensor by the one dimensional representation of B with character e^ω, we still have the same filtration but each weight ν^p is changed to $\mu^p = \nu^p + \omega$.

It follows from Theorem 17.11 that the sheaf $\mathcal{I}(\sigma_\omega)$ is filtered by a sequence of subsheaves $\{\mathcal{I}(\sigma^p)\}$, where σ^p is σ_ω restricted to V^p and the subquotient $\mathcal{I}(\sigma^p)/\mathcal{I}(\sigma^{p+1})$ is isomorphic to $\mathcal{H}_{\mu^p}$. By Theorem 17.16, the Casimir Ω acts on $\mathcal{H}_{\mu^p}$ as the scalar $\langle \mu^p, \mu^p + 2\rho\rangle$. From this it follows that $\prod_p(\Omega - \langle \mu^p, \mu^p + 2\rho\rangle)$ acts as zero on $\mathcal{I}(\sigma_\omega)$. This and Problem 17.3 then imply that $\mathcal{I}(\sigma_\omega)$ decomposes as a direct sum of subsheaves which are the Ω- eigenspaces for the eigenvalues $\langle \mu^p, \mu^p + 2\rho\rangle)$.

Note that

$$\mu^p + \rho = \nu^p + \omega + \rho = \nu^p + w\rho = w(w^{-1}\nu^p + \rho)$$

so that each $\mu^p + \rho$ is of the form $w(\nu + \rho)$, where ν is a weight for the representation (π, V). It follows from Theorem 17.7(f) that either $||\mu^p + \rho|| < ||\lambda + \rho||$ or $\mu^p + \rho = w(\lambda + \rho)$. In the latter case, $\mu^p + \rho = w\lambda + w\rho - \rho + \rho = w\lambda + \omega + \rho$, so that $\mu^p = w\lambda + \omega$. Since $\langle \mu, \mu + 2\rho\rangle = ||\mu + \rho||^2 - ||\rho||^2$ for any μ, this implies that the only μ^p for which $\langle \mu^p, \mu^p + 2\rho\rangle = \langle \lambda, \lambda + 2\rho\rangle$ is $\mu^p = w\lambda + \omega$. From this it follows that the summand of $\mathcal{I}(\sigma_\omega)$ on which $\Omega - \langle \lambda, \lambda + 2\rho\rangle$ vanishes must be a copy of $\mathcal{H}_{w\lambda+\omega}$ since this sheaf appears with multiplicity one as a subquotient of $\mathcal{I}(\sigma_\omega)$. Then

$$H^p(X, \mathcal{I}(\sigma_\omega)) = H^p(X, \mathcal{H}_{w\lambda+\omega}) \oplus H^p(X, \mathcal{J})$$

Where $\mathcal{J}$ is a summand of $\mathcal{I}(\sigma_\omega)$ on which $\Omega - \langle \lambda, \lambda + 2\rho\rangle$ is non-vanishing. However, we proved above that the Casimir Ω acts as the scalar $\langle \lambda, \lambda + 2\rho\rangle$ on the left side of this equality. It acts as this same scalar on the first term on the right side but it acts on the second term as an operator with eigenvalues all distinct from $\langle \lambda, \lambda + 2\rho\rangle$. It follows that $H^p(X, \mathcal{J}) = 0$ for all p and

$$H^p(X, \mathcal{H}_\omega) \otimes V = H^p(X, \mathcal{I}(\sigma_\omega)) = H^p(X, \mathcal{H}_{w\lambda+\omega})$$

for all p. This completes the proof.

The above is a very strong result. In particular, it implies the following:

17.18 Corollary. *The set of integers p such that $H^p(X, \mathcal{H}_\mu)$ is non-vanishing is constant as $\mu + \rho$ varies over the integral weights in a given Weyl chamber.*

Proof. If $\mu + \rho$ belongs to the Weyl chamber which is the image of the positive chamber under $w \in W$, then there is a dominant root λ such that $\mu + \rho = w(\lambda + \rho)$. Thus, $\mu = w\lambda + \omega$ where $\omega = w\rho - \rho$. It then follows from the previous lemma that $H^p(X, \mathcal{H}_\mu)$ is non-vanishing if and only if $H^p(X, \mathcal{H}_\omega)$ is non-vanishing. Since, ω depends only on the chamber determined by w, the proof is complete.

The Borel-Weil Theorem follows easily from the Lemma 17.17 and the above corollary applied to the case where w is the identity, so that $w\lambda + \omega = \lambda$.

17.19 Borel-Weil Theorem. *If λ is a dominant integral weight, then as $\mathfrak{g}$-modules*

(a) $H^0(X, \mathcal{H}_\lambda)$ *is the irreducible finite dimensional representation with highest weight λ; and*
(b) $H^p(X, \mathcal{H}_\lambda) = 0$ *for $p \neq 0$.*

Proof. Let (π, V) be the finite dimensional irreducible of highest weight λ. By Lemma 17.17 with $w = \mathrm{id}$, we have

$$H^p(X, \mathcal{H}_\lambda) = H^p(X, \mathcal{H}) \otimes V$$

This immediately implies part(a) since $H^0(X, \mathcal{H})$ is the space $\mathbb{C}$ of constants by virtue of the fact that X is compact and connected. Also, the set of p for which $H^p(X, \mathcal{H}_\lambda)$ is non-vanishing is independent of λ as long as λ is dominant. Thus, part(b) will be proved if we can show that there exists a dominant weight μ such that $H^p(X, \mathcal{H}_\mu) = 0$ for all $p > 0$. However, if V is a faithful finite dimensional irreducible $\mathfrak{g}$-module, then Theorem 17.13 implies that the pullback of $\mathcal{H}(1)$ under the embedding $\phi : X \to P(V^*)$ is $\mathcal{H}_\nu$, where ν is the highest weight of V. It follows that the pullback of $\mathcal{H}(k) = \otimes^k \mathcal{H}(1)$ is $\mathcal{H}_{k\nu}$. Thus, the cohomology of $\mathcal{H}_{k\nu}$ on X is the cohomology of $\mathcal{H}_{\phi(X)} \otimes \mathcal{H}(k)$ on $P(V^*)$, where $\mathcal{H}_{\phi(X)}$ is the structure sheaf of the subvariety $\phi(X)$ extended by zero to $P(V^*)$. Since $\mathcal{H}_{\phi(X)}$ is coherent, $\mathcal{H}_{\phi(X)} \otimes \mathcal{H}(k)$ is acyclic for sufficiently large k by Theorems 13.18 and 14.12. It follows that $\mathcal{H}_{k\nu}$ is acyclic on X for sufficiently large k. This completes the proof.

The Borel-Weil-Bott Theorem, computes $H^p(X, \mathcal{H}_\lambda)$ for all integral λ – not just the dominant ones. The next step in accomplishing this is to consider the case of singular λ – that is, λ for which $\lambda + \rho$ is not in a Weyl chamber but lies in a wall.

17.20 Theorem. *If λ is an integral weight for which $\lambda + \rho$ lies in a wall, then*

$$H^p(X, \mathcal{H}_\lambda) = 0$$

for all p.

Proof. For this argument we need to know something about the entire center $\mathcal{Z}(\mathfrak{g})$ of the enveloping algebra $U(\mathfrak{g})$, not just the Casimir element. We begin with a brief review of the relevant information about $\mathcal{Z}(\mathfrak{g})$.

Modulo the right ideal generated by $\mathfrak{b}^+$ each element of $\mathcal{Z}(\mathfrak{g})$ is equivalent to an element in $U(\mathfrak{h})$. This correspondence defines a homomorphism $\gamma : \mathcal{Z}(\mathfrak{g}) \to U(\mathfrak{h})$ called the Harish-Chandra homomorphism. Since $\mathfrak{h}$ is abelian, the algebra $U(\mathfrak{h})$ is may be regarded as the algebra of polynomial functions on the complex vector space $\mathfrak{h}^*$. The Harish-Chandra homomorphism is an injective homomorphism with image equal to the set of polynomials f such that $f(\mu - \rho)$ is invariant under the action of the Weyl group. It follows that every complex homomorphism of $\mathcal{Z}(\mathfrak{g})$ has the form $\theta_\lambda = \delta_\lambda \circ \gamma$ where $\lambda \in \mathfrak{h}^*$ and $\delta_\lambda(f) = f(\lambda)$ for each polynomial f on $\mathfrak{h}^*$. Elements λ_1 and λ_2 determine the same complex homomorphism (i. e. $\theta_{\lambda_1} = \theta_{\lambda_2}$) if and only if $\lambda_1 + \rho$ and $\lambda_2 + \rho$ belong to the same Weyl group orbit. Thus, the set of complex homomorphisms of $\mathcal{Z}(\mathfrak{g})$ is in one to one correspondence with the set of Weyl group orbits. If V is a $\mathfrak{g}$-module for which $\mathcal{Z}(\mathfrak{g})$ acts as scalars, then it acts via a complex homomorphism θ_λ. In this case, V is said to have *infinitesimal character* θ_λ.

Now for the proof of the Theorem. If $\lambda + \rho$ belongs to a wall, then its Weyl group orbit consists entirely of elements which belong to walls. If $H^p(X, \mathcal{H}_\lambda) \neq 0$ for some p, then it will be a $\mathfrak{g}$-module with infinitesimal character θ_λ by Theorem 17.16. However, there is no finite dimensional representation with infinitesimal character θ_λ since every finite dimensional irreducible representation has infinitesimal character θ_μ for μ dominant and the Weyl group orbit of a dominant weight does not meet any walls.

We need one more lemma before proving the Borel-Weil-Bott Theorem:

17.21 Lemma. *Let (π, V) be a finite dimensional irreducible representation of $\mathfrak{g}$ and let Λ be its set of weights. Let Δ^+ be a system of positive roots and let α be an element of Δ^+. If $\delta \in \mathfrak{h}'$ satisfies $\langle \delta, \alpha \rangle = 0$ and $\langle \delta, \beta \rangle > 0$ for $\beta \in \Delta^+$ distinct from α, then the maximal value of $||\delta + \nu||$ for $\nu \in \Lambda$ is achieved at exactly two points, $\nu = \mu$ and $\nu = s_\alpha(\mu)$, where μ is the highest weight in Λ.*

Proof. Since $||\nu||^2 = \langle \nu, \nu \rangle$ is a convex function of ν in $\mathfrak{h}'$, the maximum clearly can occur only at weights $\lambda + \mu$ for which μ is an extremal weight in Λ. Given two extremal weights μ and ν, we have $||\mu|| = ||\nu||$ by Theorem 17.7 and, hence,

$$||\delta + \mu||^2 - ||\delta + \nu||^2 = 2\langle \delta, \mu - \nu \rangle$$

If μ is the highest weight in Λ then $\mu - \nu$ is a sum of positive roots and so $||\delta + \mu||^2 - ||\delta + \nu||^2 > 0$ except in the case where $\mu - \nu$ involves only the root α, i. e. has the form $n\alpha$ for some n. The only extremal weight of this form is $s_\alpha(\mu)$.

Two Weyl chambers are adjacent if their closures have a wall in common. The reflection through that wall will then interchange the two chambers. If β is a positive root defining the wall (that is, the wall is $\{\lambda \in \mathfrak{h}' : \langle \lambda, \beta \rangle = 0\}$), then $\langle \lambda, \beta \rangle$ and $\langle s_\beta(\lambda), \beta \rangle$ will have opposite signs, while $\langle \lambda, \alpha \rangle$ and $\langle s_\beta(\lambda), \alpha \rangle$ will have the same sign for other positive roots α. Thus, for λ in one of the two chambers there will be exactly one more negative number in the set $\{\langle \lambda, \alpha \rangle : \alpha \in \Delta^+\}$. The distance from a Weyl chamber (or one of its elements) to the positive chamber is the minimal number of such wall crossings needed to pass from the positive chamber to the given chamber. Thus, it is the number of negative numbers in the set $\{\langle \lambda, \alpha \rangle : \alpha \in \Delta^+\}$ for λ in the chamber. The length of a Weyl group element w is the distance from $w\rho$ to the positive chamber.

17.22 Borel-Weil-Bott Theorem. *Let λ be an integral weight. Then*

(a) *if $\lambda + \rho = w(\mu + \rho)$ for a dominant weight μ and $w \in W$ of length d, then $H^d(X, \mathcal{H}_\lambda)$ is isomorphic to the irreducible $\mathfrak{g}$ module of highest weight μ;*
(b) *$H^p(X, \mathcal{H}_\lambda)$ vanishes in all other cases.*

Proof. We already know from Theorem 17.20 that $H^p(X, \mathcal{H}_\lambda) = 0$ if $\lambda + \rho$ lies in a wall. Thus, we may assume that $\lambda + \rho$ lies in a Weyl chamber – i. e. that $\lambda = w\mu + \omega$ for a dominant integral weight μ, where $\omega = w\rho - \rho$. By Theorem 17.17 we have

$$H^p(X, \mathcal{H}_\lambda) = H^p(X, \mathcal{H}_\omega) \otimes V$$

where V is the irreducible finite dimensional representation of highest weight μ. Thus, the theorem is true for a given λ with $\lambda + \rho$ in the chamber determined by w if and only if it is true for the weight $\omega = w\rho - \rho$ – that is, if and only if $H^p(X, \mathcal{H}_\omega) = 0$ unless p is the length of w, in which case $H^p(X, \mathcal{H}_\omega) = \mathbb{C}$. Of course, this means that the theorem is true for all weights λ with $\lambda + \rho$ in a given chamber if and only if it is true for one such weight.

We prove the theorem by induction on the distance d from our chamber to the positive chamber. The case $d = 0$ is the Borel-Weil theorem. Thus, we suppose the theorem is true for Weyl chambers at distance $d-1$ from the positive chamber and consider a chamber C at a distance d, obtained by applying w of length d to the positive chamber. We choose an integral weight μ which is in the positive chamber and is the highest weight of an irreducible finite dimensional representation (ψ, W). Then $w\mu$ is in the our chamber C. Now C is adjacent to a chamber C' at distance $d-1$ from the positive chamber. Let α be the positive root defining the wall separating the two chambers. Then $w' = S_\alpha \circ w$ is the Weyl group element of length $d-1$ which maps the positive chamber to C'. We set $\eta = w'(\mu) \in C'$.

As in the proof of Lemma 17.17, let σ be the representation ψ restricted to the Borel B and let σ_τ be the tensor product of σ with the one dimensional representation of B determined by an integral weight τ. Here we choose τ to be a weight that satisfies

$$\langle \tau + \rho, \alpha \rangle = 0, \quad \langle \tau + \rho, \beta \rangle > 0 \quad \text{for} \quad \beta \in \Delta^+, \beta \neq \alpha$$

For example, $\tau = w\rho + w'\rho - \rho$ has this property. We then consider the induced bundle $\mathrm{I}(\sigma_\tau)$ and its sheaf of sections $\mathcal{I}(\sigma_\tau)$. As in the proof of Lemma 17.17, $\mathrm{I}(\sigma_\tau) = \mathrm{I}(\mathrm{e}^\tau) \otimes \mathrm{I}(\sigma)$ as G-equivariant bundles, $\mathcal{I}(\sigma_\tau) = \mathcal{I}(\mathrm{e}^\tau) \otimes W$ as sheaves of $\mathfrak{g}$-modules, and for each p

$$H^p(X, \mathcal{I}(\sigma_\tau)) = H^p(X, \mathcal{H}_\tau) \otimes W$$

as $\mathfrak{g}$-modules. Since, $\tau + \rho$ is in a wall, it follows from Theorem 17.20 that

$$H^p(X, \mathcal{H}_\tau) = 0 \quad \text{and, hence,} \quad H^p(X, \mathcal{I}(\sigma_\tau)) = 0$$

for all p.

Also, as in the proof of Lemma 17.17, the sheaf $\mathcal{H}_\tau$ decomposes into eigenspaces for the action of the Casimir Ω and the possible eigenvalues are of the form $\langle \nu + \tau, \nu + \tau + 2\rho \rangle = ||\nu + \tau + \rho||^2 - ||\rho||^2$ where ν is a weight of the representation (ψ, W). Now we apply Lemma 17.21 in the case where the positive root system is the one for which $\eta + \rho$ is in the positive chamber and δ is $\tau + \rho$. It implies that the maximal value for the expression $||\nu + \tau + \rho||^2 - ||\rho||^2$, as ν ranges over the weights of ψ, is achieved only for $\nu = \eta$ and $\nu = s_\alpha(\eta)$.

We have

$$s_\alpha(\eta) = \eta - 2\frac{\langle \alpha, \eta \rangle}{\langle \alpha, \alpha \rangle}\alpha$$

and, since $\langle \alpha, \eta \rangle > 0$, this implies that $s_\alpha(\eta) < \eta$ in our ordering of integral weights. We may define a B-submodule W' of W to be the span of all weight spaces for weights $\nu < \eta$. If W'' is the quotient W/W', then we have a short exact sequence of B-modules

$$0 \to W' \to W \to W'' \to 0$$

with W' containing the weight space for weight $s_\alpha(\eta)$ and W'' containing the weight space for weight η. If we tensor this sequence with the one dimensional representation with character e^τ and then induce, we are led to a corresponding short exact sequence of sheaves of $\mathfrak{g}$-modules:

$$0 \to \mathcal{A} \to \mathcal{I}(\sigma_\tau) \to \mathcal{B} \to 0$$

If we then apply the projection onto the Ω-eigenspace for eigenvalue $t = ||\eta + \tau + \rho||^2 - ||\rho||^2 = ||s_\alpha(\eta) + \tau + \rho||^2 - ||\rho||^2$, we obtain a short exact sequence of sheaves:

$$0 \to \mathcal{A}_t \to \mathcal{I}(\sigma_\tau)_t \to \mathcal{B}_t \to 0$$

From the construction, it is clear that $\mathcal{B}_t = \mathcal{H}(\eta+\tau)$ and $\mathcal{A}_t = \mathcal{H}(s_\alpha(\eta)+\tau)$. Since $\mathcal{I}(\sigma_\tau)$ has vanishing cohomology in all degrees, the same thing is true of its direct summand $\mathcal{I}(\sigma_\tau)_t$. From the long exact sequence of cohomology, we conclude that

$$H^{p+1}(X, s_\alpha(\eta) + \tau) \simeq H^p(X, \eta + \tau)$$

for all p. Since $\tau + \rho$ is in the wall determined by α, it is fixed by s_α and so, since η is in the chamber C', it follows that $\eta + \tau + \rho$ is also in C'. Then

$$s_\alpha(\eta) + \tau + \rho = s_\alpha(\eta + \tau + \rho)$$

is the corresponding element of the chamber C. Since we have assumed the theorem true for all $\lambda + \rho$ in C' and, hence, for $\lambda + \rho = \eta + \tau + \rho$, the above identity for cohomology shows that the theorem is also true when $\lambda + \rho$ is the element $s_\alpha(\eta+\tau+\rho)$ of the chamber C. As noted above, the theorem is true for all $\lambda+\rho$ in a chamber if it is true for one. This completes the proof.

17. Problems

1. Let (σ, W) be a finite dimensional holomorphic representation of a Borel subgroup B. Prove that the induced bundle $\mathrm{I}(\sigma)$ on X is a trivial G-equivariant vector bundle if and only if σ is the restriction to B of a holomorphic representation of G.
2. Prove that if (σ, W) is a finite dimensional holomorphic representation of B then there is a holomorphic action of G on $H^p(X, \mathcal{I}(\sigma))$ for each p.
3. Suppose that A is a linear operator on a vector space V (not necessarily finite dimensional) and suppose that $\lambda_1, \ldots, \lambda_n$ are distinct scalars such that $\prod_i (A - \lambda_i) = 0$. Then prove that V decomposes as a direct sum of eigenspaces of A with eigenvalues $\lambda_1, \ldots, \lambda_n$.
4. Let α and β be roots. Use the fact that $\dfrac{2\langle\beta,\alpha\rangle}{\langle\alpha,\alpha\rangle}$ and $\dfrac{2\langle\alpha,\beta\rangle}{\langle\beta,\beta\rangle}$ are integers (Theorem 17.7) to prove that if $||\beta|| \le ||\alpha||$ then $\dfrac{2\langle\beta,\alpha\rangle}{\langle\alpha,\alpha\rangle}$ can only have the values $0, 1$, or -1.
5. Use the result of Problem 4 to prove that if α and β are non-proportional roots and $\langle\beta,\alpha\rangle > 0$ then $\alpha - \beta$ is a root.

6. Let Δ^+ be a system of positive roots. A simple element of Δ^+ is one which is not a sum of two elements of Δ^+. Use the result of Problem 5 to prove that if α and β are distinct simple elements of Δ^+ then $\langle \beta, \alpha \rangle < 0$.
7. Prove that the set $\{\alpha_i\}$ of simple elements of Δ^+ is a basis for the real vector space spanned by the roots. Furthermore, each element of Δ^+ has an expansion in this basis with coeficients which are all positive. This basis is called the *system of simple roots* generating Δ^+.
8. Let $\{\alpha_i\}$ be the system of simple roots generating Δ^+ as in Problem 7. Choose a basis $\{\alpha_i'\}$ for $\mathfrak{h}^*$ bi-orthogonal to $\{\alpha_i\}$ relative to the Killing form (i. e. choose $\{\alpha_i'\}$ so that $\langle \alpha_i, \alpha_j' \rangle = \delta_{ij}$). Prove that a weight is in the positive Weyl chamber if and only if its expansion in terms of the basis $\{\alpha_i'\}$ has all positive coeficients.
9. Prove that $\langle \rho, \alpha \rangle > 0$ for every positive root α.

References

1. Atiyah, M. F. and Macdonald, I. G., *Introduction to Commutative Algebra*, Addison-Wesley, 1969.
2. Bourbaki, N., *Groupes et algèbras de Lie*, Masson, 1975.
3. Bredon, G. E., *Sheaf Theory*, McGraw–Hill, 1967.
4. Cartan, H. and Serre, J. P., *Un théorèm de finitude concernant les variétiés analytiques compactes*, C. R. **237** (1953), 128–130.
5. Gunning, R. C., *Introduction to Holomorphic Functions of Several Complex Variables*, Wadsworth and Brooks/Cole, 1990.
6. Gunning, R. C. and Rossi, H., *Analytic Functions of Several Complex Variables*, Prentice-Hall, 1965.
7. Hartshorne, R., *Algebraic Geometry*, Springer–Verlag, 1977.
8. Humphreys, J. E., *Introduction to Lie Algebras and Representation Theory*, Springer–Verlag, 1970.
9. Hungerford, T. W., *Algebra*, Holt, Rinehart and Winston, 1974.
10. Kashiwara, M., Schapira, P., *Sheaves on Manifolds*, Springer Verlag, 1990.
11. Miličić, D., *A proof of the Borel–Weil theorem*, unpublished notes.
12. Rudin, W., *Functional Analysis*, McGraw – Hill, 1973.
13. Serre, J. P., *Geométrie algébrique et géométrie analytique*, Ann. Inst. Fourier **6** (1956), 1–42.
14. Serre, J. P., *Faisceaux algébriques cohérents*, Ann. of Math. **61** (1955), 197–278.
15. Schaefer, H. H., *Topological Vector Spaces*, Macmillan, 1966.
16. Zariski, O and Samuel, P., *Commutative Algebra*, Van Nostrand, 1960.