This document contains some of the details from 2.1-2.5 of Hartshorne - I've filled in a lot of the details that I've been curious. The word 'ring' means commutative unitary.

Let $\mathcal{F}$ be a presheaf of abelian groups (it could just as well be rings - from here on in, I just say a sheaf) over a topological space X.

Theorem 1. *There exists a sheaf $\mathcal{F}^+$ on X and a morphism $\phi : \mathcal{F} \to \mathcal{F}^+$ which satisfies the following universal property for all sheaves $\mathcal{H}$ and morphisms $\lambda : \mathcal{F} \to \mathcal{H}$.*

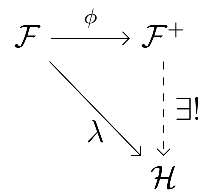

Proof. Let $U \subseteq X$ be open. Define $\mathcal{F}^+(U)$ to be the group (ring, module, ...) of functions $s : U \to \cup_{p\in U}\mathcal{F}_p$ which satisfy two requirements. The first requirement is that the functions be 'sections', ie we require that $s(q) \in \mathcal{F}_q$ for all $q \in U$. The second requirement is that s is modelled, at least locally, on an element of $\mathcal{F}(U)$. More precisely, for any $p \in U$, we require that there exists an open subset $p \in V \subseteq U$ and a section $t \in \mathcal{F}(V)$ so that for all $q \in Q$ we have '$s = t$', ie that $s(q)$ equals the image of t inside $\mathcal{F}_q$. We use the obvious restriction maps to make $\mathcal{F}^+$ into a presheaf.

Now, this gives a sheaf on X: let U be a fixed open subset and let $\{W_\alpha\}_{\alpha\in\mathcal{A}}$ be a fixed open cover of U. Suppose I give you a section $s \in \mathcal{F}^+(U)$ which satisfies $s|_{W_\alpha} = 0 \in \mathcal{F}^+(W_\alpha)$ for all α: then it is immediate that $s = 0$. Now suppose I give you sections $s_\alpha \in \mathcal{F}^+(W_\alpha)$ for all $\alpha \in \mathcal{A}$ which agree on all double intersections. Then you define $s : U \to \cup_{p\in U}\mathcal{F}_p$ by $q \mapsto s_\alpha(q)$ for $q \in W_\alpha$. We must show that $s \in \mathcal{F}^+(U)$. The first condition is obviously satisfied, so we verify the second: fix $p \in U$: for some α we have $p \in W_\alpha$. Pick an open neighbourhood $p \in W \subseteq W_\alpha$ so that there exists $t \in \mathcal{F}(W)$ which satisfies $t = s_\alpha$ on W. Then $t|_{W\cap U} \in \mathcal{F}(W \cap U)$ does the trick for our s.

So it is indeed a sheaf. Now we describe the morphism $\phi : \mathcal{F} \to \mathcal{F}^+$. Fix an open subset U of X. Let $t \in \mathcal{F}(U)$. We define $\phi_U(t) : U \to \cup_{p\in U}\mathcal{F}_p$ by $q \mapsto t_q \in \mathcal{F}_q$ - here, t_q means the image of t under the natural map $\mathcal{F}(U) \to \mathcal{F}_q$. You may check for yourself that $\{\phi_U\}_{U\subseteq X}$ is a morphism - the only thing to check is that the squares induced by a nest of open sets $W \subseteq U$ commute.

Now we verify the universal property. Fix a sheaf $\mathcal{H}$ and a morphism $\lambda : \mathcal{F} \to \mathcal{H}$. Fix an open subset U of X. Consider the diagram (now of groups)

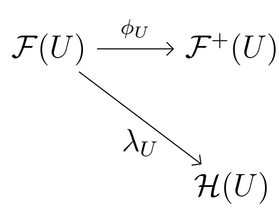

Define $\psi_U : \mathcal{F}^+(U) \to \mathcal{H}(U)$ as follows. Fix an element $s \in \mathcal{F}^+(U)$. By definition of $\mathcal{F}^+$, for each $p \in U$ we may find an open neighbourhood $p \in W_p \subseteq X$ so that there exists $t^p \in \mathcal{F}(W_p)$ which satisfy $(t^p)_q = s(q)$ for all $q \in W_p$. Consider the elements

$$\lambda_{W_p}(t^p) \in \mathcal{H}(W_p)$$

I claim that these elements agree on all double intersections $W_p \cap W_q$. Since λ is a morphism, it suffices to prove that $t^p|_{W_p\cap W_q} = t^q|_{W_p\cap W_q}$, and this is true because $(t^p)_x = s(x)$ for all $x \in W_p$.

Since $\mathcal{H}$ is a sheaf, the collection $\left\{\lambda_{W_p}(t^p)\right\}_{p\in U}$ glue to give a unique section $r \in \mathcal{H}(U)$. Define $\psi_U(s) = r$ - it may be verified that this r does not depend on our choice of W_p and of t^p.

First, we check that the diagram commutes: let $t \in \mathcal{F}(U)$. Then $\phi_U(t)$ is the map $U \to \cup_{q\in U}\mathcal{F}_q$ defined by $p \mapsto t_p$. We calculate ψ_U of this element. For an arbitrary element of $\mathcal{F}^+(U)$, we must pick an open neighbourhood of each point to calculate ψ_U. For this element however, we may pick the single element $t \in \mathcal{F}(U)$. Following through the definition of ψ_U, we must glue the 'data' $\lambda_U(t) \in \mathcal{H}(U)$ to obtain an element of $\mathcal{H}(U)$ - but the glueing is already done for us because we worked with a single element of $\mathcal{F}(U)$! So the diagram commutes.

Next, note that $\{\psi_U\}_{U\subseteq X}$ is indeed a sheaf map - we must verify for an open nest $V \subseteq U$ the commutativity of the diagram

$$\begin{array}{ccc} \mathcal{F}^+(U) & \xrightarrow{\psi_U} & \mathcal{H}(U) \\ \downarrow & & \downarrow \\ \mathcal{F}^+(V) & \xrightarrow{\psi_V} & \mathcal{H}(V) \end{array}$$

But this is easy to see - if the data $t^p \in \mathcal{F}(W_p)$ works for ψ_U, then the data $t^p|_{W_p\cap V}$ works for ψ_V.

We now note that for all $q \in X$ we have that $\mathcal{F}_q \cong \mathcal{F}^+_q$, and this isomorphism is canonical via $\phi_q : [(W,a)] \mapsto [(W,\phi_W(a))]$. It is injective: suppose that $[W,\phi_W(a)] = [V,\phi_V(b)]$. Then there exists an open subset $U \subseteq V \cap W$ so that $\phi_W(a)|_U = \phi_V(b)|_U$. Since ϕ is a morphism, this means that $\phi_U(a|_U) = \phi_U(b|_U)$. Each of the elements of $\mathcal{F}^+(U)$ are functions $U \to \cup_{p\in U}\mathcal{F} + p$; saying that they are equal means they agree at every point. Their evaluation at a point $r \in X$ is defined by $\phi_U(a|_U)(r) = a_r$ - in our current notation we might also write this as $[U,a] \in \mathcal{F}_r$. This immediately implies that $[U,a] = [U,b] \in \mathcal{F}_q$, ie that $[W,a] = [V,b]$.

Also, the map surjects: let $[W,s] \in \mathcal{F}^+_q$ - this means that $W \subseteq X$ is open and that $q \in W$ and that $s \in \mathcal{F}^+(W)$. By definition of $\mathcal{F}^+$, there exists an open subset $p \in V \subseteq W$ and an element $a \in \mathcal{F}(V)$ so that for all $p \in V$ we have that $s(p) = a_p$. It is immediate that $[W,s] = [V,s|_V] = [V,\phi_V(a)] = \phi_q([V,a])$.

From this we now verify the *uniqueness* of the universal property. For a fixed morphism $\lambda : \mathcal{F} \to \mathcal{H}$, suppose $\Psi : \mathcal{F}^+ \to \mathcal{H}$ were another solution (besides our solution ψ) to the diagram. At a point $p \in X$ we have the following group diagram

$$\begin{array}{ccc} \mathcal{F}_p & \xrightarrow{\phi_p :\cong} & \mathcal{F}^+_p \\ & \searrow^{\lambda_p} & \downarrow{\scriptstyle \psi_p,\,\Psi_p} \\ & & \mathcal{H}_p \end{array}$$

Since ϕ_p is an isomorphism, it follows that we have $\psi_p = \Psi_p$. Since $\mathcal{F}^+, \mathcal{H}$ are sheaves and morphisms are determined by their behaviour at stalks, we immediately obtain that $\psi = \Psi$.

□

Call $\mathcal{F}^+$ the **sheafification** of $\mathcal{F}$, or the **sheaf associated to the presheaf** $\mathcal{F}$.

Corollary 1. *The pair* $(\mathcal{F}^+, \phi)$ *is unique up to unique isomorphism.*

Proof. It satisfies the universal property. □

Let $f : X \to Y$ be a map of topological spaces, and let $\mathcal{G}$ be a sheaf on Y. We define the **inverse image sheaf** to be the sheaf on X associated to the presheaf

$$U \mapsto \varinjlim_{f(U)\subseteq V} \mathcal{G}(V)$$

Here, the limit is taken over open subsets V of Y which approximate (ie contain) $f(U)$; a smaller open subset better approximates $f(U)$. Denote this sheaf by $f^{-1}\mathcal{G}$. In particular, for an open subset $U \subseteq X$ and any morphism from the inverse image presheaf to a sheaf $\mathcal{H}$ we have the following universal property.

$$\begin{array}{ccc} \varinjlim_{f(U)\subseteq V} \mathcal{G}(V) & \longrightarrow & f^{-1}\mathcal{G}(U) \\ & \searrow & \downarrow \exists! \\ & & \mathcal{H}(U) \end{array}$$

Lemma 2. *Let $\mathcal{F}$ be a sheaf on X and let $\mathcal{H}$ be a sheaf on Y. Let $f : X \to Y$ be a map of topological spaces. Then there are canonical maps $\mathcal{H} \to f_*f^{-1}\mathcal{H}$ and $f^{-1}f_*\mathcal{F} \to \mathcal{F}$.*

Proof. This is just unwinding the definitions. We first construct $\mathcal{H} \to f_*f^{-1}\mathcal{H}$. Let $U \subseteq Y$ be open; we are looking for a map

$$\mathcal{H}(U) \to f_*f^{-1}\mathcal{H}(U) = f^{-1}\mathcal{H}(f^{-1}(U))$$

It suffices to exhibit a map to the presheaf that $f^{-1}\mathcal{H}$ comes from; for we simply compose this map with ϕ_U. Our map

$$\mathcal{H}(U) \to \varinjlim_{f(f^{-1}(U)\subseteq V} \mathcal{H}(V)$$

just comes from the basic inclusion $f(f^{-1}(U)) \subseteq U$: we simply map $a \mapsto [(U, a)]$. These maps obviously communte with restrictions along an open nest.

Next we exhibit the map $f^{-1}f_*\mathcal{F} \to \mathcal{F}$. By the universal property of $f^{-1}(f_*\mathcal{F})$, it suffices to (for an open subset $U \subseteq X$) to exhibit a map

$$\varinjlim_{f(U)\subseteq V} f_*\mathcal{F}(V) \to \mathcal{F}(U)$$

The domain is $\varinjlim_{f(U)\subseteq V} \mathcal{F}(f^{-1}(V))$. An element of this is a class $[(W, a)]$ with $W \subseteq Y$ open and $f(U) \subseteq W$ and $a \in \mathcal{F}(f^{-1}(W))$. Thus we have that $U \subseteq f^{-1}(W)$; this means our map is $[(W, a)] \mapsto a|_U$. Again, this gives a sheaf map. □

We say that f^{-1} is the **left adjoint** of f_* and that f_* is the **right adjoint** of f^{-1}. This is justified by the following result.

Theorem 3. *Let $f : X \to Y$ be a map and let $\mathcal{F}$ be a sheaf on X and let $\mathcal{H}$ be a sheaf on Y. There is a canonical bijection (of sets)*

$$\mathrm{Hom}_X(f^{-1}\mathcal{H}, \mathcal{F}) = \mathrm{Hom}_Y(\mathcal{H}, f_*\mathcal{F})$$

Proof. To a X-map $T : f^{-1}\mathcal{H} \to \mathcal{F}$ we associate the Y-map $\mathcal{H} \to f_*f^{-1}\mathcal{H} \stackrel{f_*T}{\to} f_*\mathcal{F}$ where the first map comes from the lemma. It is obvious what the map f_*T is: for $U \subseteq Y$ open, we assign $(f_*T)_U = T_{f^{-1}(U)}$.

To the Y-map $R : \mathcal{H} \to f_*\mathcal{F}$ we associate the X-map $f^{-1}\mathcal{H} \stackrel{f^{-1}R}{\to} f^{-1}(f_*\mathcal{F}) \to \mathcal{F}$ - here the second map comes from the lemma, and $f^{-1}R$ is as follows. Let $U \subseteq X$ be open. By the universal property, it suffices to exhibit a map

$$\lim_{\longrightarrow f(U)\subseteq V} \mathcal{H}(V) \to f^{-1}f_*\mathcal{F}(U)$$

and via the sheafification map, it suffices to exhibit a map

$$\lim_{\longrightarrow f(U)\subseteq V} \mathcal{H}(V) \to \lim_{\longrightarrow f(U)\subseteq V} f_*\mathcal{F}(V) = \lim_{\longrightarrow f(U)\subseteq V} \mathcal{F}(f^{-1}(V))$$

The natural map is $[(W,a)] \mapsto [(W, R_W(a))]$. It remains to verify that these constructions are inverses - I don't have the heart to do this.

□

Fix a ring A, and let $X = \operatorname{Spec}(A)$ be the set of prime ideals. Define a topology on X by saying that the closed subsets of X are precisely the sets $\{\mathfrak{p} \in X : I \subseteq \mathfrak{p}\}$, where $I \subseteq A$; call this set $V(I)$. Since $V(I) = V((I)_A)$, we may assume that I is an ideal.

Lemma 4. *Defining the closed sets to be of the form $V(I)$ as above forms a topology on X, which we call the* ***Zariski Topology****. Furthermore, for two ideals I, J we have that $V(I) \subseteq V(J)$ if and only if $\sqrt{J} \subseteq \sqrt{I}$.*

Proof. It is easy to verify that $V(IJ) = V(I) \cup V(J)$ and that $V(\sum_i I_i) = \cap_i V(I_i)$ for an arbitrary family of ideals $\{I_i\}$. Furthermore, $\emptyset = V(A)$ and $X = V(0)$. The second statement boils down to the fact from commutative algebra that for any ideal $J \subseteq A$ we have that $\sqrt{J}$ equals the intersection of all prime ideals which contain J. For completeness we prove this.

One direction is trivial. Conversely, suppose that $x \notin \sqrt{J}$. Let $S = \{x^n : n \geq 0\}$. Since $J \cap S = \emptyset$, the ideal JA_S of A_S is proper, and is therefore contained inside a maximal ideal M of A_S. The pullback $\mathfrak{m}$ of M to A is thus a prime of A which contains J. It is easy to see that $x \notin \mathfrak{m}$. □

Lemma 5. *A basis for the topology on $X = \operatorname{Spec} A$ are the sets $D(f) = \{\mathfrak{p} \in X : f \notin \mathfrak{p}\}$ as f ranges over A.*

Proof. Let $U \subseteq X$ be open, and let $\mathfrak{q} \in U$. Since U is open, it is equal to $X - V(I)$ for some ideal I of A. Thus I is not a subset of $\mathfrak{q}$; so we may pick $a \in I$ with $a \notin \mathfrak{q}$. I claim that $D(a) \subseteq U$: if $\mathfrak{p} \in D(a)$, then $a \notin \mathfrak{p}$, which means I is not a subset of $\mathfrak{p}$, which means $\mathfrak{p} \in U$. Since $\mathfrak{q} \in D(a)$, we are done. □

Lemma 6. *Let $X = \operatorname{Spec} A$, and let f_α be a collection of elements of A. Let $X = \cup_\alpha D(f_\alpha)$ if and only if $(f_\alpha)_\alpha = A$. In particular, X is quasicompact.*

Proof. Suppose $1 = \sum_\alpha a_\alpha f_\alpha$. Let $\mathfrak{p} \in X$; if each f_α lives in $\mathfrak{p}$, then so does 1 which is impossible. Does $f_\beta \notin \mathfrak{p}$ for some β. Thus $\mathfrak{p} \in D(f_\beta)$. Conversely, suppose (f_α) is proper; then pick a prime ideal P containing (f_α). Since every $f_\alpha \in P$, we have that $P \notin D(f_\alpha)$ for every α.

Finally, cover X by an open cover; we may assume each element of this open cover is $D(f)$ for some $f \in A$. Then the finitely many f which generate the unit ideal also cover X. □

We now define a sheaf of rings on X; it is called the **structure sheaf**. Let $U \subseteq X$ be open. We define $\mathcal{O}(U)$ to be the set of functions $s : U \to \bigsqcup_{\mathfrak{p} \in U} A_{\mathfrak{p}}$ which satisfy the following two conditions. The first is that they are sections, ie that $s(\mathfrak{q}) \in A_{\mathfrak{q}}$ for all $\mathfrak{q} \in U$.

The second condition is less obvious, but is equally important. Suppose we are given $a, s \in A$ so that for all $\mathfrak{p} \in U$ we have that $s \notin \mathfrak{p}$. Then there is a natural map from $U \to \bigsqcup_{\mathfrak{p} \in U} A_{\mathfrak{p}}$, namely $\mathfrak{q} \mapsto a/s \in A_{\mathfrak{q}}$ where here a/s denotes its image under the localisation map $A \to A_{\mathfrak{q}}$. We require that our maps s be locally modelled on this condition. More precisely, for each $\mathfrak{p} \in U$ we require there be an open neighbourhood $\mathfrak{p} \in V \subseteq U$ and $b, f \in A$ so that $f \notin \mathfrak{q}$ for every $\mathfrak{q} \in V$. Then we insist that over V, the function s is modelled by b/f as above, ie we have that $s(\mathfrak{q}) = b/f \in A_{\mathfrak{q}}$ for all $q \in V$.

This is a presheaf of rings under pointwise operations and usual function restriction. It is a sheaf because it is constructed exactly as we constructed the sheafification.

Theorem 7. *The structure sheaf $\mathcal{O}$ of $X = \operatorname{Spec}(A)$ relates in the following way to various localisations of A.*

(1) For every $\mathfrak{p} \in X$ we have that $\mathcal{O}_{\mathfrak{p}} \cong A_{\mathfrak{p}}$.

(2) For every $f \in A$ we have that $\mathcal{O}(D(f)) \cong A_f$. In particular, $\mathcal{O}(X) \cong A$.

Proof. For the first claim, define a map $\phi : \mathcal{O}_{\mathfrak{p}} \to A_{\mathfrak{p}}$ by $[V, s] \mapsto s(\mathfrak{p})$. This is well defined because our 'restriction' maps are just function restrictions, and it is clearly a ring map. An element of $A_{\mathfrak{p}}$ can be represented as a quotient a/f with $a \in A$ and $f \in A - \mathfrak{p}$. Let $U \subseteq X$ denote the open set $X - V(f)$, and let $t \in \mathcal{O}(U)$ denote the map corresponding to a, f. Then $\phi : [U, t] \mapsto a/f \in A_{\mathfrak{p}}$, so the map is surjective.

Now let $s, t \in \mathcal{O}(U)$ have the same image at $\mathfrak{p}$. By shrinking U, we may assume that s corresponds to a/f and t corresponds to b/g. Thus there exists $h \notin \mathfrak{p}$ so that $h(ga - fb) = 0 \in A$. It follows that $s = t$ inside every local ring $A_{\mathfrak{q}}$ where $f, g, h \notin \mathfrak{q}$. The set of all such q is the open neighbourhood $V = X - (V(f) \cup V(g) \cup V(h)) = X - V(fgh)$ of $\mathfrak{p}$. So $s = t$ in an open neighbourhood of $\mathfrak{p}$, proving that ϕ is injective.

For the second part of the theorem, define a function $\psi : A_f \to \mathcal{O}(D_f)$ by sending a/f^n to the element s of $\mathcal{O}(D_f)$ corresponding to a/f^n. First, we show that ψ is injective. Suppose $\psi(a/f^n) = \psi(b/f^m)$. This means that for all $\mathfrak{q} \in D_f$ we have that $a/f^n = b/f^m \in A_{\mathfrak{q}}$. For a fixed $p \in D_f$ we may find an element $h \in A - \mathfrak{p}$ so that $h(f^m a - f^n b) = 0 \in A$. Letting $\mathfrak{a}$ denote the annihilator of $f^m a - f^n b$, we have that $h \in \mathfrak{a}$. Since $h \notin \mathfrak{p}$, we see that $\mathfrak{a}$ is not a subset of $\mathfrak{p}$. So we have that $\mathfrak{p} \in D_f$ implies that $\mathfrak{a}$ is not a subset of p, which implies that $p \notin V(\mathfrak{a})$. Thus we have that $V(\mathfrak{a}) \cap D_f = \emptyset$. Thus $V(\mathfrak{a})$ is contained in the (set-theoretic) complement of D_f, which is $V(f)$. In particular, the earlier lemma shows that $\sqrt{f} \subseteq \sqrt{\mathfrak{a}}$, which shows that $f^l(f^m a - f^n b) = 0$ for some $l \geq 1$. This proves that $a/f^n = b/f^m \in A_f$, so ψ is injective.

Now, we show that ψ is onto. Let $s \in \mathcal{O}(D(f))$. Our first reduction is to simplify what we may assume about the definition of the structure sheaf. By its definition, we may cover $D(f)$ with open sets V_i so that $s|_{V_i}$ is represented by a_i/g_i, where $a_i, g_i \in A$ and $g_i \notin \mathfrak{q}$ for all $\mathfrak{q} \in V_i$. In other words, $V_i \subseteq D(g_i)$. Since the sets $D(h)$ form a basis for the Zariski topology as h ranges over A, we may, making our open sets V_i smaller if neccesary, assume that $V_i = D(h_i)$. Fix an i for now. By the lemma this gives us that $\sqrt{h_i} \subseteq \sqrt{g_i}$, with says that for some $n \geq 0$ and $c \in A$ we have that $h_i^n = cg_i$, which implies that $a_i/g_i = ca_i/h_i^n$. Relabelling h_i^n as h_i (certainly $D(h) = D(h^n)$) and ca_i as a_i, we may thus assume that $D(f)$ is covered by open sets of the form $D(h_i)$, and that $s|_{D(h_i)}$ corresponds to a_i/h_i on this open subset.

Next, we reduce to the case of considering only finitely many h_i. We have that $D(f) \subseteq \cup_i D(h_i)$, which says that $\cap_i (V(h_i)) = V(\sum_i (h_i)) \subseteq V(f)$. By the lemma this implies that $f^n \in \sum_i (h_i)$ for some $n \geq 0$, and hence is a *finite* A-linear combination of the h_i. This finite set of the $D(h_i)$ in fact covers the $D(f)$.

Since we have that $D(h_i) \cap D(h_j) = D(h_i h_j)$, it must be the case that inside $A_{h_i h_j}$ we have the equality $a_i/h_i = a_j/h_j$; this is because s is represented by both of them on this open subset, and we know that $\psi : A_{h_i h_j} \to \mathcal{O}(D(h_i h_j))$ is injective. Thus for some $n \geq 0$ we have that

$$(h_i h_j)^n (h_j a_i - h_i a_j) = 0$$

We are only dealing with finitely many h_i; thus we may pick n so that it works simultaneously for all pairs $(a_i/h_i, a_j/h_j)$. Relabel h_i^{n+1} as h_i and $h_i^n a_i$ as a_i. We still have the representation of s on $D(h_i)$ as a_i/h_i, and now we may assume that $h_j a_i = h_i a_j$ for all i, j.

For $b_i \in A$, write $f^n = \sum_i b_i a_i$; we previously showed this is possible. Define the element $a \in A$ as $a = \sum_i b_i a_i$. For any j we then have

$$\begin{aligned} h_j a &= \sum_i b_i a_i h_j \\ &= \sum_i b_i h_i a_j \\ &= f^n a_j \end{aligned}$$

Thus, on $D(h_i)$ we may also represented s by a/f^n. It is immediate that $\psi(a/f^n) = s$, proving the theorem. □

Definition 1. *Let X be a space with a sheaf of rings $\mathcal{O}$. We say that $(X, \mathcal{O})$ is a **locally ringed space** if for all $p \in X$ we have that $\mathcal{O}_p$ is a local ring. A **morphism** of locally ringed spaces $(X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ consists of a map $f : X \to Y$ and a map $f^\# : \mathcal{O}_Y \to f_* \mathcal{O}_X$ so that for all $p \in X$ we have that $f^\#_P : \mathcal{O}_{Y,f(p)} \to \mathcal{O}_{X,p}$ is local.*[1] *By this map I mean as follows: as V ranges over the open neighbourhoods of $f(P)$ we have that $f^{-1}(V)$ ranges over (some of) the open neighbourhoods of P. Take limits; this latter term naturally maps to the stalk $\mathcal{O}_{X,P}$.*

Lemma 8. *Let $\phi : A \to B$ be a ring map, and let $f : \operatorname{Spec} B \to \operatorname{Spec} A$ be defined by $\mathfrak{q} \mapsto \phi^{-1}(\mathfrak{q})$. Then f is continuous.*

Proof. Let $\mathfrak{a} \subseteq A$ be an ideal. We prove that $f^{-1}(V(\mathfrak{a}) = V(\phi(\mathfrak{a}))$, which says that the inverse image of a closed set under f is closed. Let $\mathfrak{q} \in f^{-1}(V(\mathfrak{a})$. Then $f(\mathfrak{q}) \in V(\mathfrak{a})$, which says that $\mathfrak{a} \subseteq f(\mathfrak{q})$. Applying ϕ yields that $\phi(\mathfrak{a}) \subseteq \phi(\phi^{-1}(\mathfrak{q})) \subseteq \mathfrak{q}$ where the final containment is a set-theoretic fact. Therefore $\mathfrak{q} \in V(\phi(\mathfrak{a}))$.

Conversely, let $\mathfrak{q} \in V(\phi(\mathfrak{a}))$ - we must show that $\mathfrak{q} \in f^{-1}(V(\mathfrak{a}))$, ie that $f(\mathfrak{q}) \in V(\mathfrak{a})$, ie that $\mathfrak{a} \subseteq f(\mathfrak{q})$. Since $\mathfrak{q} \in V(\phi(\mathfrak{a}))$ we have that $\phi(\mathfrak{a}) \subseteq \mathfrak{q}$: applying f says that $f(\phi(\mathfrak{a})) \subseteq f(\mathfrak{q})$. The left hand side is $\phi^{-1}\phi(\mathfrak{a})$, which by general set theory contains $\mathfrak{a}$. □

Proposition 9. *Every ring map $\phi : A \to B$ induces a morphism of ringed spaces $(f, f^\#)$: $(\operatorname{Spec} B, \mathcal{O}_{\operatorname{Spec} B}) \to (\operatorname{Spec} A, \mathcal{O}_{\operatorname{Spec} A})$.*

Proof. The continuous map f is given by the previous lemma. For notation, given any ring R we denote $\mathcal{O}_{\operatorname{Spec} R}$ by $\mathcal{O}_R$. We describe the morphism $f^\# : \mathcal{O}_A \to f_* \mathcal{O}_B$. To this end, let $V \subseteq \operatorname{Spec} A$ be open. Let $s \in \mathcal{O}_A(V)$ - it is a map

$$V \to \bigsqcup_{\mathfrak{q} \in V} A_{\mathfrak{q}}$$

[1] A map of local rings is called **local** if the inverse image of the maximal ideal is the maximal ideal

which satisfies (1) $s(\mathfrak{p}) \in A_\mathfrak{p}$ for all $\mathfrak{p} \in V$ and (2) for all $\mathfrak{p} \in V$ there exists an open neighbourhood $W(\mathfrak{p}) = W$ of $\mathfrak{p}$ and $a, f \in A$ with $f \notin \mathfrak{m}$ for all $\mathfrak{m} \in W$ so that $s|_W$ corresponds to $a/f : W \to \bigsqcup_{\mathfrak{m}\in W} A_\mathfrak{m}$. We must associate to s a section $t \in \mathcal{O}_B(f^{-1}(V))$ - that is a map

$$f^{-1}(V) \to \bigsqcup_{\mathfrak{m}\in f^{-1}(V)} B_\mathfrak{m}$$

which satisfies the right properties. Since we are ranging over $\mathfrak{m} \in f^{-1}(V)$, we definitely have $f(\mathfrak{m}) \in V$, and associated to this is the local homomorphism $\phi_\mathfrak{m} : A_{f(\mathfrak{m})} \to B_\mathfrak{m}$. Define t to be the composition

$$f^{-1}(V) \to V \to \bigsqcup_{f(\mathfrak{m})\in V} A_{f(\mathfrak{m})} \to \bigsqcup_{\mathfrak{m}\in f^{-1}(V)} B_\mathfrak{m}$$

where the first map is f, the second map is s, and the final map is the union of all the $\phi_\mathfrak{m}$. This map obviously satisfies that $t(\mathfrak{q}) \in B_\mathfrak{q}$ for all $q \in f^{-1}(V)$. For the second condition, let $\mathfrak{p} \in f^{-1}(V)$. Pick a neighbourhood W of $f(p)$ in V so that $s|_W$ corresponds to a/f, where $a \in A$ and $f \notin \mathfrak{m}$ for all $\mathfrak{m} \in W$. It is routine to verify that $\phi(f) \notin \mathfrak{q}$ for all $\mathfrak{q} \in f^{-1}(W)$, and that t corresponds to $\phi(a)/\phi(f)$ on the neighbourhood $f^{-1}(W)$ of $\mathfrak{p}$.

Finally, we must show that the induced stalk map $f^\#$ on the stalks are local homomorphisms. For $q \in \operatorname{Spec} B$, consider the map

$$f^\#_\mathfrak{q} : \mathcal{O}_{A,f(\mathfrak{q})} \to \mathcal{O}_{B,\mathfrak{q}}$$

Unwinding the definition of this map and using our known isomorphisms, this map is seen to be precisely the local homomorphism

$$A_{f(\mathfrak{q})} \to B_\mathfrak{q}$$

More precisely, an element $a/g \in A_{f(\mathfrak{q})}$ corresponds to the class $[D_A(g), a/g] \in \mathcal{O}_{A,f(\mathfrak{q})}$ from theorem 7, and our stalk map sends this to the class $[f^{-1}(D_A(f)), \phi(a)/\phi(g)]$, and this class corresponds to $\phi(a)/\phi(g)$, again from theorem 7, showing that the stalk map is precisely our beloved local homomorphism. □

It is clear that if we start with a ring map $\phi : A \to B$, then the induced map on global sections from the morphism of the previous lemma is exactly ϕ again. What is less clear is the converse.

Lemma 10. *Every morphism $(f, f^\#) : (\operatorname{Spec} B, \mathcal{O}_B) \to (\operatorname{Spec} A, \mathcal{O}_A)$ 'comes from' a ring map $A \to B$ as in the previous lemma.*

Proof. Let $f^\#_{\operatorname{Spec} A} = \phi : A \to B$ denote the map on global sections. For $\mathfrak{p} \in \operatorname{Spec} B$, we thus have a commutative diagram

$$\begin{array}{ccc} A & \xrightarrow{\phi} & B \\ \downarrow & & \downarrow \\ A_{f(\mathfrak{p})} & \xrightarrow{f^\#_\mathfrak{p}} & B_\mathfrak{p} \end{array}$$

I claim that $\phi^{-1}(\mathfrak{p}) = f(\mathfrak{p})$, ie that the continuous map f is precisely the map induced from ϕ from a few lemmas ago. To see this, note that $(f^\#_\mathfrak{p})^{-1}(\mathfrak{p}B_\mathfrak{p}) = f(\mathfrak{p})A_{f(\mathfrak{p})}$ since the map is a local homomorphism. Pulling back $f(\mathfrak{p})A_{f(\mathfrak{p})}$ along the first diagnol map yields the ideal $f(\mathfrak{p})$ of A. The other direction: pulling back $\mathfrak{p}B_\mathfrak{p}$ via the second vertical map yields the ideal $\mathfrak{p}$ of B, and pulling it back via ϕ yields $\phi^{-1}(\mathfrak{p})$. Since the diagram commutes, these ideals are equal. So f is induced from the map ϕ.

We must now show that $f^\#$ is also induced from ϕ - but this is true because they induce the same map on stalks. $\square$

Definition 2. *Let $(X, \mathcal{O})$ be a scheme. We say that X is* ***locally noetherian*** *if it can be covered by open subsets $U_\alpha \cong \operatorname{Spec}(A_\alpha)$ where A_α is noetherian. We call X* ***noetherian*** *if it is locally noetherian and quasi-compact.*

Why about arbitrary affine open subsets? The next theorem answers this question. First, we need an algebraic lemma.

Lemma 11. *Let $f_1, ..., f_r \in A$ with $\phi_i : A \to A_{f_i}$ denoting localisation. Suppose that $(f_1, ..., f_r)$ generate the unit ideal, and let $\mathfrak{a}$ be an ideal of A. Then*

$$\mathfrak{a} = \cap_i \phi_i^{-1}(\phi_i(\mathfrak{a})A_{f_i})$$

Proof. The left hand side is contained in the right hand side. Let $b \in A$ be contained in the right hand side. Then $\phi_i(b) = a_i/f_i^{n_i}$ for some $a_i \in \mathfrak{a}$ and $n_i \geq 0$. Possibly replacing $a_i f^l$ with a_i, we may assume that all of the n_i are all equal to say n. This means that, in the ring A, we have that $f_i^{m_i}(bf_i^n - a_i) = 0$ for some $m_i > 0$. Again, we may take all of the m_i equal to some fixed m. Let $N = n + m$; we thus have that $f_i^N b = f_i^m a_i \in \mathfrak{a}$. Since $(f_1, ..., f_r) = A$, we also have that $(f_1^N, ..., f_r^N) = A$: simply write 1 an an A-linear combination of the f_i and raise it to a huge power. So, for appropriate $c_i \in A$ we have that $1 = \sum_i c_i f_i^N$. Therefore $b = \sum_i c_i b f_i^N \in \mathfrak{a}$, which proves the lemma. $\square$

Theorem 12. *Let X be locally noetherian, and let $U \cong \operatorname{Spec} A$ be open. Then A is noetherian.*

Proof. We make several claims. First: let B be noetherian and let $f \in B$. Then B_f is noetherian. To see this, suppose that $J \subseteq B_f$ is an ideal; denote by ϕ the localisation map $B \to B_f$. Then $I = \phi^{-1}(J)$ is an ideal of B, which is generated by $b_1, ..., b_d$ since B is noetherian. I claim that J is generated by $b_1/1, ..., b_d/1$. Given $x \in J$, we may write it as c/f^n. Then, in B, we write $c = a_1 b_1 + ... + a_d b_d$. Since $\phi(c) = c/1 = f^n(c/f^n)$, we see that $\phi(c) \in J$. Hence $\phi(c)/f^n \in J$, which is exactly x. The claim is proved.

The next claim is as follows: let $U \subseteq X$ be open. Then we may write U as a union of the spectra of noetherian rings. Proceed as follows. Since X is locally noetherian, we may cover X by $W_\alpha \cong \operatorname{Spec} A_\alpha$ with each A_α noetherian. Then $U \cap W_\alpha$ is an open subset of W_α; it is thus equal to a union $\cup_{\lambda \in \Lambda_\alpha} D(f_\lambda)$. Since for any ring R we have a bijection between the primes of R_f and the primes of R that miss $\{f^n : n \geq 0\}$, we may identify the open subset $D(f)$ of $\operatorname{Spec} R$ with the affine scheme $\operatorname{Spec}(R_f)$. Thus we may write

$$U = \cup_\alpha \cup_{\lambda \in \Lambda_\alpha} \operatorname{Spec}((C_\alpha)_{f_\lambda})$$

By the first claim, each of these rings is noetherian, and so the second claim is proved. Furthermore, since U is an affine scheme, it is quasi-compact, hence finitely many of the $\operatorname{Spec}((C_\alpha)_{f_\lambda})$ cover U.

Next, consider the following set-up. We have an affine scheme $X = \operatorname{Spec} A$ and $U = \operatorname{Spec} B$ an open subset of X. We may pick some $f \in A$ so that $D_A(f) \subseteq U$. The inclusion $U \subseteq X$ corresponds to a ring homomorphism $\phi : A \to B$; let $\overline{f}$ be the image of f under this map. I claim that $A_f \cong B_{\overline{f}}$. Indeed, the element $f \in A$ is viewed as a function $X \to \cup_{\mathfrak{q} \in X} A_\mathfrak{q}$. Since $U \subseteq X$ is open, the stalk $\mathcal{O}_{\operatorname{Spec} B, \mathfrak{p}}$, which is equal to $B_\mathfrak{p}$, is isomorphic to the stalk $\mathcal{O}_{\operatorname{Spec} A, \phi^{-1}(\mathfrak{p})}$, which is equal to $A_{\phi^{-1}(\mathfrak{p})}$. The isomorphism is given by the localisation of ϕ at $\mathfrak{p}$. It follows that $\overline{f}$, when viewed as a function

$U \to \cup_{\mathfrak{q} \in \operatorname{Spec} B} B_{\mathfrak{q}}$, is simply the restriction of f to U. This implies that $D_A(f) = D_B(\overline{f})$ as schemes, which says that $A_f \cong B_{\overline{f}}$.

Thus our cover of U by the $\operatorname{Spec}((C_\alpha)_{f_\lambda})$ corresponds to a (finite) covering of $\operatorname{Spec} A$ by the spectrum of the noetherian rings $\operatorname{Spec} A_{\overline{f_\lambda}}$.

Therefore, the problem has been reduced to proving the following. Let A be a ring and let $f_1, ..., f_r$ generate the unit ideal. Let $\phi_i : A \to A_{f_i}$ denote localisation, and suppose that each A_{f_i} is noetherian. Show that A is noetherian. Let

$$\mathfrak{a}_1 \subseteq \mathfrak{a}_2 \subseteq ...$$

be an ideal chain. Then we obtain, for each i, an ideal chain

$$\phi_i(\mathfrak{a}_1)A_{f_i} \subseteq \phi_i(\mathfrak{a}_2)A_{f_i} \subseteq ...$$

Since each A_{f_i} is noetherian, each of these chains must stabalize. Then, by the lemma, when the 'longest' A_{f_i} chain stabalizes, so does the original chain, so A is noetherian. □