Globalizing modules

rapha • 5 Mar 2025

Given an A-module M, we may build a sheaf of $\mathscr{O}$-modules associated to M on the affine scheme $\operatorname{Spec} A$ (with structure sheaf $\mathscr{O}$). Often, this sheaf of modules is constructed by first defining it on a basis, and then using the fact that a sheaf on a basis uniquely defines a sheaf on the whole space.

In this post, I want to present a more "hands on" way of constructing that sheaf. This is part of my quest to "reconcile" with *elementary* constructions (i.e. those using elements instead of universal properties).

* * *

As mentioned in the introduction, we will build a sheaf $\widetilde{M}$, said to be the sheaf **associated to** M on the affine scheme $\operatorname{Spec} A$. This construction plays an important role in algebraic geometry because it's a "local model" for quasicoherent sheaves.

Before we begin, we need the following technical definition. Let U be some open set in $\operatorname{Spec} A$, and let $(s_\mathfrak{p})_{\mathfrak{p}\in U}$ be an element of $\prod_{\mathfrak{p}\in U} M_\mathfrak{p}$. We say that $(s_\mathfrak{p})_{\mathfrak{p}\in U}$ is a **system of compatible germs** if, for every $\mathfrak{p}\in U$, there exists an element $f\in A$ having $\mathfrak{p}\in D(f)$ and $D(f)\subseteq U$, together with an element $t\in M_f$ such that, for every $\mathfrak{q}\in D(f)$, we have $t_\mathfrak{q}=s_\mathfrak{q}$. Here $t_\mathfrak{q}$ is the further localization of t via the canonical $M_f\to M_\mathfrak{q}$.

Now, we use this notion to define our object of interest. Let U be an open set of $\operatorname{Spec} A$. To this open set we associate the data $\widetilde{M}(U)$, which is defined to be the subset of all systems of compatible germs in $\prod_{\mathfrak{p}\in U} M_\mathfrak{p}$.

Notice that $\widetilde{M}(U)$ is an abelian group by adding two systems component-wise, using the abelian group structure on each $M_\mathfrak{p}$. Moreover, we can give $\widetilde{M}(U)$ the structure of an $\mathscr{O}(U)$-module in the following way. For a scalar $a\in\mathscr{O}(U)$ and a system $(s_\mathfrak{p})_{\mathfrak{p}\in U}\in\widetilde{M}(U)$, we define

$$a\cdot(s_\mathfrak{p})_{\mathfrak{p}\in U}=(a_\mathfrak{p}\cdot s_\mathfrak{p})_{\mathfrak{p}\in U},$$

using the fact that $\mathscr{O}_{\mathfrak{p}} \cong A_{\mathfrak{p}}$, and that each $M_{\mathfrak{p}}$ is an $A_{\mathfrak{p}}$-module. We have a couple of things to check before we can claim this actually works:

- First of all, we need to verify the resulting system is of compatible germs. This is not too hard. Let $\mathfrak{p} \in U$, and $f \in A$ and $t \in M_f$ be given for $\mathfrak{p}$ as in the definition of a system of compatible germs. There's a canonical A_f-module structure on M_f, which allows us to talk about $t' = a_f \cdot t$. Here a_f is the image of a via the restriction map $\mathscr{O}(U) \to \mathscr{O}(D(f)) \cong A_f$. It doesn't take much work to see that localizing further using $M_f \to M_{\mathfrak{q}}$ preserves the action on the module, in the sense that $(t')_{\mathfrak{q}} = a_{\mathfrak{q}} \cdot s_{\mathfrak{q}}$ for all $\mathfrak{q} \in D(f)$ (hint: the localization map $M_f \to M_{\mathfrak{q}}$ corresponds to the obvious map $M \otimes A_f \to M \otimes A_{\mathfrak{q}}$; write the appropriate commutative diagram). This shows everything works, and we have a system of compatible germs.

- Then, we need to check that this respects the axioms for a module action. This is so because "taking the germ", i.e. sending a to $a_{\mathfrak{p}}$, is a ring homomorphism for each $\mathfrak{p}$.

If $V \subseteq U$ is an inclusion of open sets, we may define the obvious restriction $\widetilde{M}(U) \to \widetilde{M}(V)$, which simply sends a system $(s_{\mathfrak{p}})_{\mathfrak{p} \in U}$ to the smaller system $(s_{\mathfrak{p}})_{\mathfrak{p} \in V}$. The only subtlety here is that it is not obvious the smaller system is of compatible germs. However, that's not too bad to show (hint: the distinguished opens form a base for the topology on $\operatorname{Spec} A$, and $D(g) \subseteq D(f)$ implies the existence of a canonical localization map $M_f \to M_g$). Notice that this restriction map respects the module structure, in the sense that acting by $\mathscr{O}(U)$ followed by restriction gives the same result as restriction followed by acting by $\mathscr{O}(V)$.

From the definition of the restriction maps as litteral restrictions, it's obvious that we have a presheaf on $\operatorname{Spec} A$. We want to show it is a sheaf. It's clear that the identity axiom holds: if some section restricts locally everywhere to the zero system, that means each component of the system is zero, so the system as a whole is a bunch of zeroes.

We check the gluability axiom. Let U be an open set, and let $\{U_i\}_{i \in I}$ be a family of open sets that cover U. Suppose we also have a family of systems $\{(s^i_{\mathfrak{p}})_{\mathfrak{p} \in U_i}\}_{i \in I}$ with each $(s^i_{\mathfrak{p}})_{\mathfrak{p} \in U_i} \in \widetilde{M}(U_i)$. Suppose further that these systems agree on overlaps, in the sense that for each $(i, j) \in I^2$, we have

$$(s^i_{\mathfrak{p}})_{\mathfrak{p} \in U_i \cap U_j} = (s^j_{\mathfrak{p}})_{\mathfrak{p} \in U_i \cap U_j}.$$

To glue all of these together, the obvious choice is to set

$$s_{\mathfrak{p}} = s^{c(\mathfrak{p})}_{\mathfrak{p}}$$

for each $\mathfrak{p} \in U$, where $c : U \to I$ is some choice function such that each $\mathfrak{p}$ lies in the chosen open set $U_{c(\mathfrak{p})}$. Because the systems agree on overlaps, this definition is independent of the actual choice function. It is also quite clear that this defines a system of compatible germs over U, and that moreover its restriction to each U_i gives back the corresponding system $(s^i_{\mathfrak{p}})_{\mathfrak{p} \in U_i}$.

So we have a sheaf! In fact, a sheaf of $\mathscr{O}$-modules, since as we've remarked above the restriction maps respect the sheaf of ring's action. Notice that we haven't really *used* the fact each section of this sheaf is a system of compatible germs. We have only demonstrated that this property carries over our constructions. But this property will be useful in characterizing this sheaf via its stalks, which is the object of the next section.

The stalks of $\widetilde{M}$

For each $\mathfrak{p} \in \operatorname{Spec} A$, the stalk $\widetilde{M}_{\mathfrak{p}}$ is an $\mathscr{O}_{\mathfrak{p}}$-module in the following way. Recall that each element in $\widetilde{M}_{\mathfrak{p}}$ is an equivalence class $[U, s]$ of pairs, with U an open neighborhood around $\mathfrak{p}$ and $s \in \widetilde{M}(U)$, and where two pairs (U, s) and (U', s') are equivalent if and only if there exists V an open neighborhood of $\mathfrak{p}$ with $V \subseteq U \cap U'$ such that $s|_V = s'|_V$.

Given two pairs $[U, s]$ and $[V, t]$ in $\widetilde{M}_{\mathfrak{p}}$, we may first find a small enough open W around $\mathfrak{p}$ such that $W \subseteq U \cap V$, and define, using the $\mathscr{O}(W)$-module structure on $\widetilde{M}(W)$:

$$[U, s] + [V, t] = [W, s|_W + t|_W].$$

This is well-defined, as one can check quickly. Similarily, given $[V, a] \in \mathscr{O}_{\mathfrak{p}}$ and $[U, s] \in \widetilde{M}_{\mathfrak{p}}$, we again find a small enough open W and define

$$[V, a] \cdot [U, s] = [W, a|_W \cdot s|_W].$$

Everything works well to give to $\widetilde{M}_{\mathfrak{p}}$ the structure of a $\mathscr{O}_{\mathfrak{p}}$-module, as claimed.

We already know that the stalk at $\mathfrak{p}$ of the structure sheaf $\mathscr{O}$ is isomorphic to the localization $A_{\mathfrak{p}}$. In fact, a similar thing can be said of the sheaf $\widetilde{M}$: it is isomorphic to $M_{\mathfrak{p}}$, as we now show. Let $[U, s]$ be an element of $\widetilde{M}_{\mathfrak{p}}$. The element s is a system of compatible germs over U, which we may write as $(s_{\mathfrak{p}})_{\mathfrak{p} \in U}$. Let $\phi([U, s])$ be the element $s_{\mathfrak{p}} \in M_{\mathfrak{p}}$. This makes a well-defined function

$$\phi : \widetilde{M}_{\mathfrak{p}} \to M_{\mathfrak{p}}.$$

This function is in fact an $\mathscr{O}_{\mathfrak{p}}$-linear map (or, which is the same thing, an $A_{\mathfrak{p}}$-linear map).

The map ϕ is injective: suppose $\phi([U, s]) = 0$. Because s is a system of compatible germs, we can find some $f \in A$ and some $t \in M_f$ such that $\mathfrak{p} \in D(f) \subseteq U$, and for each $\mathfrak{q} \in D(f)$, we have $t_\mathfrak{q} = s_\mathfrak{q}$. We may write t as a fraction m/f^n where $m \in M$ and $n \geq 0$ is some integer. Then, because the image of t in $M_\mathfrak{p}$ is zero, there must exist some $s \in A \setminus \mathfrak{p}$ such that

$$s \cdot m = 0 \tag{1}$$

in M. It is clear that $D(sf)$ is an open neighborhood of $\mathfrak{p}$ which is contained in $D(f)$. Moreover, for any $\mathfrak{q} \in D(sf)$, the localization $M_f \to M_\mathfrak{q}$ factorizes through M_{sf}. Now, the image of t in M_{sf} is

$$f^{-1} \cdot m = \frac{s^n \cdot m}{(sf)^n} = 0$$

which is equal to zero by equation (1) above. Hence $t_\mathfrak{q} = 0$ for every $\mathfrak{q} \in D(sf)$. This shows that $[U, s] = [D(d), 0]$, so ϕ is injective.

The map ϕ is also surjective. Pick some element $m/f \in M_\mathfrak{p}$. Now we consider the fraction $s = m/f$ as an element of M_f. It's clear that the system $(s_\mathfrak{q})_{\mathfrak{q} \in D(f)}$ is of compatible germs, and that $\phi([D(f), s]) = m/f$. Hence ϕ is surjective, as we wanted.

Therefore ϕ is an isomorphism!

Remark: the isomorphism ϕ is the canonical one identifying $M_\mathfrak{p}$ with $\widetilde{M}_\mathfrak{p}$ as a colimit, since we've implicitely used morphisms $\widetilde{M}(U) \to M_\mathfrak{p}$ in the construction of ϕ. In other words, this ϕ is what we get when we apply the universal property of $\widetilde{M}_\mathfrak{p}$ as a colimit; here we defined it in an elementary fashion.

Behavior over distinguished open sets

Just as we were expecting each stalk to "be" the localization at a point, just by the way we defined things, we now expect $\widetilde{M}(D(f))$ to be the smaller localization M_f. Moreover, for any inclusion of distinguished open sets $D(g) \subseteq D(f)$, we expect that the isomorphisms $\widetilde{M}(D(f)) \cong M_f$ make the following diagram commute:

$$\begin{array}{ccc} \widetilde{M}(D(f)) & \xrightarrow{\cong} & M_f \\ \downarrow & & \downarrow \\ \widetilde{M}(D(g)) & \xrightarrow[\cong]{} & M_g \end{array}$$

However, this is better seen via an alternative construction.

Another construction

Another construction is given as an exercise in *FoAG* (Ravi Vakil). For any distinguished open set $D(f)$, let $\widetilde{M}(D(f))$ be the localization of M at the multiplicative submonoid of the functions that do not vanish outside of $D(f)$, i.e. those $g \in A$ such that $V(g) \subseteq V(f)$. This definition obviously depends only on the *set* $D(f)$ and not on f the function itself; this point is the main technical advantage in defining $\widetilde{M}(D(f))$ this way instead of directly stating that it is M_f. We still have an isomorphism between $\widetilde{M}(D(f))$ and M_f, however:

Lemma. For any section $f \in A$, there is a canonical isomorphism of A_f-modules

$$\widetilde{M}(D(f)) \cong M_f.$$

Here, the isomorphism being "canonical" means that $\widetilde{M}(D(f))$ *is* the localization of M at f, in the sense that both $\widetilde{M}(D(f))$ and M_f verify the same universal property.

Proof. Recall that for any multiplicative submonoid $S \subseteq A$, we have a canonical isomorphism of $S^{-1}A$-modules between $S^{-1}M$ and $M \otimes_A S^{-1}A$, under which m/s corresponds to $m \otimes (1/s)$. (As usual, "canonical" means "coming from an universal property", the "universal property" in this case being the tensor product's).

Let now S be the multiplicative submonoid of the functions $g \in A$ having $V(g) \subseteq V(f)$. Because $f \in S$, the section f is in particular an invertible element of $S^{-1}A$. Hence, the universal property of localization gives a canonical ring homomorphism $\theta : A_f \to S^{-1}A$ which sends a/f^n to a/f^n (the inverse of f in $S^{-1}A$ may be written as $1/f$ since $f \in S$). On the other hand, any section $g \in S$ verifies $f \in \sqrt{(g)}$, so the element g is invertible in A_f. Again by the universal property, there exists a canonical ring homomorphism in the other direction $S^{-1}A \to A_f$, and it must be the inverse of θ by usual abstract nonsense. Therefore θ is an isomorphism of rings (in fact, an isomorphism of A_f-algebras).

We give $S^{-1}M$ the structure of an A_f-module using θ in the obvious way. We obtain a chain of isomorphisms of A_f-modules

$$M_f \cong M \otimes_A A_f \xrightarrow{\mathrm{id} \otimes \theta} M \otimes_A S^{-1}A \cong S^{-1}M = \widetilde{M}(D(f))$$

under which $m/f^n \in M_f$ corresponds to $m/f^n \in \widetilde{M}(D(f))$, and in the other direction, the element $m/g \in \widetilde{M}(D(f))$ with $ag = f^n$ for some integer $n \geq 1$ corresponds to $am/f^n \in M_f$. ■

Given an inclusion of distinguished open sets $D(g) \subseteq D(f)$, we want to define a restriction morphism from $\widetilde{M}(D(f))$ to $\widetilde{M}(D(g))$. Since we have $V(f) \subseteq V(g)$, by definition the section f is invertible in $\widetilde{M}(D(g))$ (with inverse $1/f$). Of course, strictly speaking, we should say more formally that the map defined by $m \mapsto fm$ is an automorphism of the A_g-module $\widetilde{M}(D(g))$, instead of saying that "f is invertible in $\widetilde{M}(D(g))$". In reality, f is invertible *in* A_g, making A_g into an A_f-algebra. Even better, the universal property of localization for modules gives us an A_f-linear map

$$\widetilde{M}(D(f)) \to \widetilde{M}(D(g))$$

which we call the restriction morphism. Since it's obtained via a universal property, everything is functorial. Moreover, recall that $\mathscr{O}(D(f)) \cong A_f$ in a canonical way, i.e. they verify the same universal property, and so each A_f-module *is* an $\mathscr{O}(D(f))$-module; also, the ring homomorphism $A_f \to A_g$ alluded to is precisely the restriction map for the sheaf of rings $\mathscr{O}$, and the fact it makes $\widetilde{M}(D(g))$ into an A_f-module means precisely that for any $h \in \mathscr{O}(D(f))$ and any $x \in \widetilde{M}(D(f))$, we have

$$(h \cdot x)|_{D(g)} = h|_{D(g)} \cdot x|_{D(g)}.$$

This makes $\widetilde{M}$ into a presheaf of $\mathscr{O}$-modules on the distinguished base.

The restriction $\widetilde{M}(D(f)) \to \widetilde{M}(D(g))$ *is* the localization $M_f \to M_g$, in the sense that the following diagram commutes:

$$\begin{array}{ccc} \widetilde{M}(D(f)) & \xrightarrow{\cong} & M_f \\ \downarrow & & \downarrow \\ \widetilde{M}(D(g)) & \xrightarrow{\cong} & M_g \end{array}$$

Lemma. The presheaf $\widetilde{M}$ is a sheaf on the distinguished base.

Proof. We start with the base identity axiom. Suppose $\{D(f_i)\}_{i \in I}$ is some covering of a distinguished open $D(f)$; by quasicompacity, we may take I to be a finite set $\{1, 2, \ldots, n\}$. Suppose m/g is some section of $\widetilde{M}$ over $D(f)$ such that $(m/g)|_{D(f_i)} = 0$ for each $1 \leq i \leq n$. To verify the identity axiom, it suffices to show $m = 0$ in M_f, given that $m|_{D(f_i)} = 0$ for each i.

As a section over $D(f)$, m restricts to an element $a_i m / f_i^{k_i}$ in M_{f_i}, and this element is zero by hypothesis. Therefore, using the fact I is a finite set, we may choose a large enough integer N such that, for each $1 \leq i \leq n$, it holds in M that $(f_i a_i)^N m = 0$.

Under the identification of $D(f)$ with $\operatorname{Spec} A_f$, each $D(f_i)$ corresponds to $D(f_i/1)$ where $f_i/1$ is the image of f_i in A_f. By laziness, we immediately stop writing "$/1$". Because the sets $D(f_i)$ cover $\operatorname{Spec} A_f$, the elements $(f_i a_i)^N$ generate the whole ring A_f (hint: $D(f_i) = D(a_i f)$). In particular, we may write 1 as a linear combination:

$$1 = h_1(f_1 a_1)^N + h_2(f_2 a_2)^N + \cdots + h_n(f_n a_n)^N,$$

where each h_i is an element of A_f. Because $(f_i a_i)^N m = 0$, we find that

$$m = h_1(f_1 a_1)^N m + h_2(f_2 a_2)^N m + \cdots + h_n(f_n a_n)^N m = 0.$$

This shows $m = 0$ in M_f, so the identity axiom is verified.

Now, we show the base gluability axiom. Fix an arbitrary covering $\{D(f_i)\}_{i \in I}$ of $D(f)$, which may be infinite. Suppose we have a collection of sections $\{m_i/f_i^{k_i}\}_{i \in I}$ with $m_i/f_i^{k_i} \in \widetilde{M}(D(f_i))$ such that these sections all "agree on overlaps", that is,

$$(m_i/f_i^{k_i})|_{D(f_i f_j)} = (m_j/f_j^{k_j})|_{D(f_i f_j)}$$

for every $(i, j) \in I^2$.

We break the proof in two parts: when I is finite, and when it is not. First, suppose I is the finite set $\{1, 2, \ldots, n\}$. To simplify notation, set $g_i = f_i^{k_i}$ (notice that $D(g_i) = D(f_i)$). Each section m_i/g_i restricts to the element $g_j m_i / g_i g_j$ in $M_{g_i g_j}$. The overlap condition then says that for each $(i, j) \in I^2$, there is some integer $k_{ij} \geq 0$ such that

$$(g_i g_j)^{k_{ij}}(g_j m_i - g_i m_j) = 0$$

holds in M. Using the finiteness of I, we may in fact pick a single integer N such that

$$(g_i g_j)^N (g_j m_i - g_i m_j) = 0.$$

We further simplify the situation by setting $b_i = g_i^N m_i$ and $h_i = g_i^{N+1}$. We have $D(f_i) = D(h_i)$ and moreover $m_i/f_i^{k_i} = b_i/h_i$ in $\widetilde{M}(D(f_i))$. The previously displayed equation becomes

$$h_j b_i = h_i b_j. \tag{1}$$

Since $D(h_i)$ covers $\operatorname{Spec} A_f$, we may express 1 as a linear combination

$$1 = r_1 h_1 + r_2 h_2 + \cdots + r_n h_n \tag{2}$$

where each $r_i \in A_f$. We finally define the section that will be the "gluing" of the sections we started with. It's the section r over $D(f)$ defined as

$$r = r_1 b_1 + r_2 b_2 + \cdots + r_n b_n.$$

Notice that for every $i \in I$,

$$\begin{aligned} h_i r &= r_1 h_i b_1 + r_2 h_i b_2 + \cdots + r_n h_i b_n \\ &= r_1 h_1 b_i + r_2 h_2 b_i + \cdots + r_n h_n b_i && \text{(by eqn. 1)} \\ &= (r_1 h_1 + r_2 h_2 + \cdots + r_n h_n) b_i \\ &= b_i. && \text{(by eqn. 2)} \end{aligned}$$

Therefore, r becomes b_i/h_i when restricted to $D(h_i)$, which shows that r is the gluing of our original sections $m_i/f_i^{k_i}$.

We still have to show the gluability axiom holds when I is infinite. In that case, use quasicompactness of $D(f)$ to choose a finite subset $J \subseteq I$ such that $\{D(f_j)\}_{j \in J}$ covers $D(f)$, and do the same construction as before to obtain a section r over $D(f)$. Now pick some new index i which is not in J, and redo the same construction with $J \cup \{i\}$ to obtain yet again a section r' over $D(f)$. But r and r' both restrict to the same sections over the covering afforded by J. By the identity axiom, we must have $r = r'$, so that in particular r restricted to $D(f_i)$ is $m_i/f_i^{k_i}$. This means we can take r, obtained by gluing only a finite number of chosen sections, to play the role of the gluing of the whole collection of sections. This shows the gluing axiom holds. ■

As always, when we have a sheaf on a base, there's a unique canonical way of extending it to a sheaf on the whole space. In fact, this extended sheaf, if defined using the concept of system of compatible germs, is exactly how we defined $\widetilde{M}$ in the first section of this article.