ON ∞-TOPOI

JACOB LURIE

Let X be a topological space and G an abelian group. There are many different definitions for the cohomology group $\mathrm{H}^n(X, G)$; we will single out three of them for discussion here. First of all, one has the singular cohomology $\mathrm{H}^n_{\mathrm{sing}}(X, G)$, which is defined as the cohomology of a complex of G-valued singular cochains. Alternatively, one may regard $\mathrm{H}^n(\bullet, G)$ as a representable functor on the homotopy category of topological spaces, and thereby define $\mathrm{H}^n_{\mathrm{rep}}(X, G)$ to be the set of homotopy classes of maps from X into an Eilenberg-MacLane space $K(G, n)$. A third possibility is to use the sheaf cohomology $\mathrm{H}^n_{\mathrm{sheaf}}(X, \underline{G})$ of X with coefficients in the constant sheaf $\underline{G}$ on X.

If X is a sufficiently nice space (for example, a CW complex), then all three of these definitions agree. In general, however, all three give different answers. The singular cohomology of X is constructed using continuous maps from simplices Δ^k into X. If there are not many maps *into* X (for example if every path in X is constant), then we cannot expect $\mathrm{H}^n_{\mathrm{sing}}(X, G)$ to tell us very much about X. Similarly, the cohomology group $\mathrm{H}^n_{\mathrm{rep}}(X, G)$ is defined using maps from X into a simplicial complex, which (ultimately) relies on the existence of continuous real-valued functions on X. If X does not admit many real-valued functions, we should not expect $\mathrm{H}^n_{\mathrm{rep}}(X, G)$ to be a useful invariant. However, the sheaf cohomology of X seems to be a good invariant for arbitrary spaces: it has excellent formal properties in general and sometimes yields good results in situations where the other approaches do not apply (such as the étale topology of algebraic varieties).

We shall take the point of view that the sheaf cohomology of a space X gives the right answer in all cases. We should then ask for conditions under which the other definitions of cohomology give the same answer. We should expect this to be true for singular cohomology when there are many continuous functions *into* X, and for Eilenberg-MacLane cohomology when there are many continuous functions *out of* X. It seems that the latter class of spaces is much larger than the former: it includes, for example, all paracompact spaces, and consequently for paracompact spaces one can show that the sheaf cohomology $\mathrm{H}^n_{\mathrm{sheaf}}(X, G)$ coincides with the Eilenberg-MacLane cohomology $\mathrm{H}^n_{\mathrm{rep}}(X, G)$. One of the main results of this paper is a generalization of the preceding statement to non-abelian cohomology, and to the case where the coefficient system G is not necessarily constant.

Classically, the non-abelian cohomology $\mathrm{H}^1(X, G)$ of X with coefficients in a possibly non-abelian group G is understood as classifying G-torsors on X. When X is paracompact, such torsors are again classified by homotopy classes of maps from X into an Eilenberg-MacLane space $K(G, 1)$. Note that the group G and the space $K(G, 1)$ are essentially the same piece of data: G determines $K(G, 1)$ up to homotopy equivalence, and conversely G may be recovered as the fundamental group of $K(G, 1)$. To make this canonical, we should say that specifying G is equivalent to specifying the space $K(G, 1)$ *together with a base point*; the space $K(G, 1)$ alone only determines G up to inner automorphisms. However, inner automorphisms of G induce the identity on $\mathrm{H}^1(X, G)$, so that $\mathrm{H}^1(X, G)$ is really a functor which depends only on $K(G, 1)$. This suggests the proper coefficients for non-abelian cohomology are not groups, but "homotopy types" (which we regard as purely combinatorial entities, represented perhaps by simplicial sets). We may define the non-abelian cohomology $\mathrm{H}_{\mathrm{rep}}(X, K)$ of X with coefficients in any simplicial complex K to be the collection of homotopy classes of maps from X into K. This leads to a good notion whenever X is paracompact. Moreover, we gain a great deal by allowing the case where K is not an Eilenberg-MacLane space. For example, if $K = \mathrm{BU} \times \mathbb{Z}$ is the classifying space for complex K-theory and X is a compact Hausdorff space, then $\mathrm{H}_{\mathrm{rep}}(X, K)$ is the usual complex K-theory of X, defined as the Grothendieck group of the monoid of isomorphism classes of complex vector bundles on X.

When X is not paracompact, we are forced to seek a better way of defining $\mathrm{H}(X, K)$. Given the apparent power and flexibility of sheaf-theoretic methods, it is natural to look for some generalization of sheaf cohomology, using as coefficients "sheaves of homotopy types on X." In other words, we want a theory of ∞-stacks (in groupoids) on X, which we will henceforth refer to simply as *stacks*. One approach to this theory is provided by the Joyal-Jardine homotopy theory of simplicial presheaves on X. According to this approach, if K is a simplicial set, then the cohomology of X with coefficients in K should be defined as $\mathrm{H}_{JJ}(X, K) = \pi_0(\mathcal{F}(X))$, where $\mathcal{F}$ is a fibrant replacement for the constant simplicial presheaf with value K on X. When K is an Eilenberg-MacLane space $K(G, n)$, then this agrees with the sheaf-cohomology group (or set) $\mathrm{H}^n_{\mathrm{sheaf}}(X, G)$. It follows that if X is paracompact, then $\mathrm{H}_{JJ}(X, K) = \mathrm{H}_{\mathrm{rep}}(X, K)$ whenever K is an Eilenberg-MacLane space.

However, it turns out that $\mathrm{H}_{JJ}(X, K) \neq \mathrm{H}_{\mathrm{rep}}(X, K)$ in general, even when X is paracompact. In fact, one can give an example of a compact Hausdorff space for which $\mathrm{H}_{JJ}(X, BU \times \mathbb{Z})$ is not equal to the complex K-theory of X. We shall proceed on the assumption that $K(X)$ is the "correct" answer in this case, and give an alternative to the Joyal-Jardine theory which computes this answer. Our alternative is distinguished from the Joyal-Jardine theory by the fact that we require our stacks to satisfy a descent condition only for coverings, rather than for arbitrary hypercoverings. Aside from this point we proceed in the same way, setting $\mathrm{H}(X, K) = \pi_0(\mathcal{F}'(X))$, where $\mathcal{F}'$ is the stack which is obtained by forcing the "constant prestack with value K" to satisfy this weaker form of descent. In general, $\mathcal{F}'$ will not satisfy descent for hypercoverings, and consequently it will not be equivalent to the simplicial presheaf $\mathcal{F}$ used in the definition of H_{JJ}.

The resulting theory has the following properties:

- If X is paracompact, $\mathrm{H}(X, K)$ is the set of homotopy classes from X into K.
- If X is a paracompact space of finite covering dimension, then our theory of stacks is equivalent to the Joyal-Jardine theory. (This is also true for certain inductive limits of finite dimensional spaces, and in particular for CW complexes.)
- The cohomologies $\mathrm{H}_{JJ}(X, K)$ and $\mathrm{H}(X, K)$ always agree when K is "truncated", for example when K is an Eilenberg-MacLane space. In particular, $\mathrm{H}(X, K(G, n))$ is equal to the usual sheaf cohomology $\mathrm{H}^n_{\mathrm{sheaf}}(X, G)$.

In addition, our theory of ∞-stacks enjoys good formal properties which are not always shared by the Joyal-Jardine theory; we shall summarize the situation in §2.10. However, the good properties of our theory do not come without their price. It turns out that the essential difference between stacks (which are required to satisfy descent only for ordinary coverings) and hyperstacks (which are required to satisfy descent for arbitrary hypercoverings) is that the former can fail to satisfy the Whitehead theorem: one can have, for example, a pointed stack (E, η) for which $\pi_i(E, \eta)$ is a trivial sheaf for all $i \geq 0$, such that E is not "contractible" (for the definition of these homotopy sheaves, see §2.8).

In order to make a thorough comparison of our theory of stacks on X and the Joyal-Jardine theory of hyperstacks on X, it seems desirable to fit both of them into some larger context. The proper framework is provided by the notion of an ∞-topos, which is intended to be an ∞-category that "looks like" the ∞-category of ∞-stacks on a topological space, just as an ordinary topos is supposed to be a category that "looks like" the category of sheaves on a topological space. For any topological space X (or, more generally, any topos), the ∞-stacks on X comprise an ∞-topos, as do the ∞-hyperstacks on X. However, it is the former ∞-topos which enjoys the more universal position among ∞-topoi related to X (see Proposition 2.7.4).

Let us now outline the contents of this paper. We will begin in §1 with an informal review of the theory of ∞-categories. There are many approaches to the foundation of this subject, each having its own particular merits and demerits. Rather than single out one of those foundations here, we shall attempt to explain the ideas involved and how to work with them. The hope is that this will render this paper readable to a wider audience, while experts will be able to fill in the details missing from our exposition in whatever framework they happen to prefer.

Section 2 is devoted to the notion of an ∞-topos. We will begin with an intrinsic characterization (analogous to Giraud's axioms which characterize ordinary topoi: see [5]), and then argue that any ∞-category satisfying our axioms actually arises as an ∞-category of "stacks" on *something*. We will then show that any ∞-topos determines an ordinary topos in a natural way, and vice versa. We will also explain the relationship between stacks and hyperstacks.

Section 3 relates the ∞-topos of stacks on a paracompact space with some notions from classical homotopy theory. In particular, we prove that the "non-abelian cohomology" determined by the ∞-topos associated to a paracompact space agrees, in the case of constant coefficients, with the functor $H_{\text{rep}}(X, K)$ described above.

Section 4 is devoted to the dimension theory of ∞-topoi. In particular, we shall define the *homotopy dimension* of an ∞-topos, which simultaneously generalizes the covering dimension of paracompact topological spaces and the Krull dimension of Noetherian topological spaces. We also prove a generalization of Grothendieck's vanishing theorem for cohomology on a Noetherian topological space of finite Krull dimension.

This paper is intended to be the first in a series devoted to the notion of an ∞-topos. In the future we plan to discuss proper maps of ∞-topoi, the proper base change theorem, and a notion of "elementary ∞-topos" together with its relationship to mathematical logic.

I would like to thank Max Lieblich, whose careful reading of this document and many suggestions have improved the exposition considerably. I would also like to thank Bertrand Toën, with whom I shared several stimulating conversations about the subject matter of this paper.

1. ∞-Categories

1.1. **General Remarks.** Throughout this paper, we will need to use the language of ∞-categories. Since there are several approaches to this subject in the literature, we devote this first section to explaining our own point of view. We begin with some general remarks directed at non-experts; seasoned homotopy-theorists and category-theorists may want to skip ahead to the next section.

An ordinary category consists of a collection of objects, together with sets of morphisms between those objects. The reader may also be familiar with the notion of a 2-category, in which there are not only morphisms but also *morphisms between the morphisms*, which are called 2-morphisms. The vision of higher categories is that one should be able to discuss n-categories for any n, in which one has not only objects, morphisms, and 2-morphisms, but k-morphisms for all $k \leq n$. Finally, in some sort of limit one should obtain a theory of ∞-categories where one has morphisms of all orders.

There are many approaches to realizing this vision. We might begin by defining a 2-category to be a "category enriched over categories." In other words, one considers a collection of objects together with a *category* of morphisms $\text{Hom}(A, B)$ for any two objects A and B, and composition *functors* c_{ABC} : $\text{Hom}(A, B) \times \text{Hom}(B, C) \to \text{Hom}(A, C)$ (to simplify the discussion, we shall ignore identity morphisms for a moment). These functors are required to satisfy an associative law, which asserts that for any quadruple (A, B, C, D) of objects, $c_{ACD} \circ (c_{ABC} \times 1) = c_{ABD} \circ (1 \times c_{BCD})$ as functors

$$\text{Hom}(A, B) \times \text{Hom}(B, C) \times \text{Hom}(C, D) \to \text{Hom}(A, D)$$

This leads to the definition of a *strict* 2*-category.*

At this point, we should object that the definition of a strict 2-category violates one of the basic philosophical principles of category theory: one should never demand that two functors be "equal." Instead one should postulate the existence of a natural transformation between two functors. This means that the associative law should not take the form of an equation, but of additional structure: a natural isomorphism $\gamma_{ABCD} : c_{ACD} \circ (c_{ABC} \times 1) \simeq c_{ABD} \circ (1 \times c_{BCD})$. We should also demand the natural transformations γ_{ABCD} be functorial in the quadruple (A, B, C, D), and that they satisfy certain higher associativity conditions. After formulating the appropriate conditions, we arrive at the definition of a *weak* 2*-category*, or simply a 2*-category.*

Let us contrast the notions of "strict 2-category" and "weak 2-category." The former is easier to define, since we do not have to worry about the higher associativity conditions satisfied by the transformations γ_{ABCD}. On the other hand, the latter notion seems more natural if we take the philosophy of category theory seriously. In this case, we happen to be lucky: the notions of "strict 2-category" and "weak 2-category" turn out to be equivalent. More precisely, any weak 2-category can be replaced by an "equivalent" strict 2-category. The question of which definition to adopt is therefore an issue of aesthetics.

Let us plunge onward to 3-categories. Following the above program, we may define a *strict* 3-*category* to consist of a collection of objects together with strict 2-categories $\mathrm{Hom}(A, B)$ for any pair of objects A and B, together with a strictly associative composition law. Alternatively, we could seek a definition of *weak* 3-*category* by allowing $\mathrm{Hom}(A, B)$ to be only a weak 2-category, requiring associativity only up to natural 2- isomorphisms, which satisfy higher associativity laws up to natural 3-isomorphisms, which in turn satisfy still higher associativity laws of their own. Unfortunately, it turns out that these notions are *not* equivalent.

Both of these approaches have serious drawbacks. The obvious problem with weak 3-categories is that an explicit definition is extremely complicated (see [12], where a definition is worked out along these lines), to the point where it is essentially unusable. On the other hand, strict 3-categories have the problem of not being the correct notion: most of the weak 3-categories which occur in nature (such as the fundamental 3-groupoid of a topological space) are not equivalent to strict 3-categories. The situation only gets worse (from either point of view) as we pass to 4-categories and beyond.

Fortunately, it turns out that major simplifications can be introduced if we are willing to restrict our attention to ∞-categories in which most of the higher morphisms are invertible. Let us henceforth use the term (∞, n)-category to refer to ∞-categories in which all k-morphisms are invertible for $k > n$. It turns out that $(\infty, 0)$-categories (that is, ∞-categories in which *all* morphisms are invertible) have been extensively studied from another point of view: they are the same thing as "spaces" in the sense of homotopy theory, and there are many equivalent ways of describing them (for example, using CW complexes or simplicial sets).

We can now proceed to define an $(\infty, 1)$-category: it is something which has a collection of objects, and between any two objects A and B there is an $(\infty, 0)$-category, or "space," called $\mathrm{Hom}(A, B)$. We require that these morphism spaces are equipped with an associative composition law. As before, we are faced with two choices as to how to make this precise: do we require associativity on the nose, or only "up to homotopy" in some sense? Fortunately, it turns out not to matter: as was the case with 2-categories, any $(\infty, 1)$-category with a coherently associative multiplication can be replaced by an equivalent $(\infty, 1)$-category with a strictly associative multiplication.

In this paper, we will deal almost exclusively with $(\infty, 1)$-categories (with a few exceptional appearances of $(\infty, 2)$-categories, for which we will not require any general theory). *Unless otherwise specified, all* ∞-*categories will be assumed to be* $(\infty, 1)$-*categories.*

There are a number of models for these ∞-categories: categories enriched over simplicial sets, categories enriched over topological spaces, Segal categories ([17]), simplicial sets satisfying a weak Kan condition (called quasi-categories by Joyal, see [6] and [7]). All of these approaches give essentially the same notion of ∞-category, so it does not matter which approach we choose.

1.2. ∞-Categories.

The first ingredient needed for the theory of ∞-categories is the theory of ∞-groupoids, or "spaces": namely, homotopy theory. Our point of view on this subject is that exists some abstract and purely combinatorial notion of a "homotopy type" which is ∞-categorical in nature, and therefore slippery and difficult to define in terms of sets. However, there are various ways to "model" homotopy types by classical mathematical objects such as CW complexes or simplicial sets. Let us therefore make the definition that a *space* is a CW complex. We shall agree that when we discuss a space X, we shall refer only to its homotopy-theoretic properties and never to specific features of some particular representative of X.

We can now say what an ∞-category is:

Definition 1.2.1. An ∞-*category* $\mathcal{C}$ consists of a collection of objects, together with a space $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ for any pair of objects $X, Y \in \mathcal{C}$. These Hom-spaces must be equipped with an associative composition law

$$\mathrm{Hom}_{\mathcal{C}}(X_0, X_1) \times \mathrm{Hom}_{\mathcal{C}}(X_1, X_2) \times \ldots \mathrm{Hom}_{\mathcal{C}}(X_{n-1}, X_n) \to \mathrm{Hom}_{\mathcal{C}}(X_0, X_n)$$

(defined for all $n \geq 0$).

Remark 1.2.2. It is customary to use the compactly generated topology on

$$\mathrm{Hom}_{\mathcal{C}}(X_0, X_1) \times \ldots \times \mathrm{Hom}_{\mathcal{C}}(X_{n-1}, X_n),$$

rather than the product topology, so that the product remains a CW complex. This facilitates comparisons with more combinatorial versions of the theory, but the issue is not very important.

Remark 1.2.3. Of the numerous possible definitions of ∞-categories, Definition 1.2.1 is the easiest to state and to understand. However, it turns out to be one of the hardest to work with explicitly, since many basic ∞-categorical constructions are difficult to carry out at the level of topological categories. We will not dwell on these technical points, which are usually more easily addressed using the more sophisticated approaches.

Let us now see how to work with the ∞-categories introduced by Definition 1.2.1. Note first that any ∞-category $\mathcal{C}$ determines an ordinary category $h\,\mathcal{C}$ having the same objects, but with $\mathrm{Hom}_{h\,\mathcal{C}}(X,Y) = \pi_0 \mathrm{Hom}_{\mathcal{C}}(X,Y)$. The category $h\,\mathcal{C}$ is called the *homotopy category* of $\mathcal{C}$ (or sometimes the *derived category* of $\mathcal{C}$). To some extent, working in the ∞-category $\mathcal{C}$ is like working in its homotopy category $h\mathcal{C}$: up to equivalence, $\mathcal{C}$ and $h\,\mathcal{C}$ have the same objects and morphisms. The difference between $h\,\mathcal{C}$ and $\mathcal{C}$ is that in $\mathcal{C}$, one must not ask about whether or not morphisms are *equal*; instead one should ask whether or not one can find a path from one to the other. One consequence of this difference is that the notion of a commutative diagram in $h\,\mathcal{C}$, which corresponds to a *homotopy commutative* diagram in $\mathcal{C}$, is quite unnatural and usually needs to be replaced by the more refined notion of a *homotopy coherent* diagram in $\mathcal{C}$.

To understand the problem, let us suppose that $F : \mathcal{I} \to h\,\mathcal{S}$ is a functor from an ordinary category $\mathcal{I}$ into the homotopy category of spaces $\mathcal{S}$. In other words, F assigns to each object $x \in \mathcal{I}$ a space Fx, and to each morphism $\phi : x \to y$ in $\mathcal{I}$ a continuous map of spaces $F\phi : Fx \to Fy$ (well-defined up to homotopy), such that $F(\phi \circ \psi)$ is homotopic to $F\phi \circ F\psi$ for any pair of composable morphisms ϕ, ψ in $\mathcal{I}$. In this situation, it may or may not be possible to *lift* F to an actual functor $\widetilde{F}$ from $\mathcal{I}$ to the ordinary category of topological spaces, such that $\widetilde{F}$ induces a functor $\mathcal{I} \to h\,\mathcal{S}$ which is equivalent to F. In general there are obstructions to both the existence and the uniqueness of the lifting $\widetilde{F}$, even up to homotopy. To see this, we note that $\widetilde{F}$ determines extra data on F: for every composable pair of morphisms ϕ and ψ, $\widetilde{F}(\phi \circ \psi) = \widetilde{F}\phi \circ \widetilde{F}\psi$, which means that $\widetilde{F}$ gives a *specified* homotopy $h_{\phi,\psi}$ between $F(\phi \circ \psi)$ and $F\phi \circ F\psi$. We should imagine that the functor F to the homotopy category $h\,\mathcal{S}$ is a first approximation to $\widetilde{F}$; we obtain a second approximation when we take into account the homotopies $h_{\phi,\psi}$. These homotopies are not arbitrary: the associativity of composition gives a relationship between $h_{\phi,\psi}, h_{\psi,\theta}, h_{\phi,\psi\circ\theta}$ and $h_{\phi\circ\psi,\theta}$, for a composable triple of morphisms (ϕ, ψ, θ) in $\mathcal{I}$. This relationship may be formulated in terms of the existence of a certain higher homotopy, which is once again canonically determined by $\widetilde{F}$. To obtain the next approximation to $\widetilde{F}$, we should take these higher homotopies into account, and formulate the associativity properties that *they* enjoy, and so forth. A *homotopy coherent* diagram in $\mathcal{C}$ is, roughly speaking, a functor $F : \mathcal{I} \to h\,\mathcal{C}$, together with all of the extra data that would be available if there we could lift F to get some $\widetilde{F}$ which was functorial "on the nose."

An important consequence of the distinction between homotopy commutativity and homotopy coherence is that the appropriate notions of *limit* and *colimit* in $\mathcal{C}$ do not coincide with the notion of *limit* and *colimit* in $h\,\mathcal{C}$ (in which limits and colimits typically do not exist). The appropriately defined limits and colimits in $\mathcal{C}$ are typically referred to as *homotopy limits* and *homotopy colimits*, to avoid confusing them ordinary limits and colimits inside of some category of models for $\mathcal{C}$. We will try to avoid using categories of models at all and work in an "invariant" fashion with an ∞-category $\mathcal{C}$. In particular, the terms *limit* and *colimit* in an ∞-category $\mathcal{C}$ will *always* mean homotopy limit and homotopy colimit. We shall never speak of any other kind of limit or colimit.

Example 1.2.4. Suppose given a collection of objects $\{X_\alpha\}$ in an ∞-category $\mathcal{C}$. The *product* $X = \prod_\alpha X_\alpha$ (if it exists) is characterized by the following universal property: $\mathrm{Hom}_{\mathcal{C}}(Y,X) \simeq \prod_\alpha \mathrm{Hom}_{\mathcal{C}}(Y,X_\alpha)$. Passing to connected components, we see that we also have $\mathrm{Hom}_{h\,\mathcal{C}}(Y,X) \simeq \prod_\alpha \mathrm{Hom}_{h\,\mathcal{C}}(Y,X_\alpha)$. Consequently, any product in $\mathcal{C}$ is also a product in $h\,\mathcal{C}$.

Example 1.2.5. Given two maps $\pi : X \to Z$ and $\psi : Y \to Z$ in an ∞-category $\mathcal{C}$, the *fiber product* $X \times_Z Y$ (if it exists) is characterized by the following universal property: to specify a map $W \to X \times_Z Y$ is to specify maps $W \to X$, $W \to Y$, *together with a homotopy between the induced composite maps* $W \to Z$. In particular, there is a map from $X \times_Z Y$ to the analogous fiber product in $h\,\mathcal{C}$ (if this fiber product exists), which need not be an isomorphism.

We want to emphasize the point of view that ∞-categories (with n-morphisms assumed invertible for $n > 1$) are in many ways a less drastic generalization of categories than ordinary 2-categories (in which there can exist non-invertible 2-morphisms). Virtually all definitions and constructions which make sense for ordinary categories admit straightforward generalizations to ∞-categories, while for 2-categories even the most basic notions need to be rethought (for example, one has a distinction between strict and lax pullbacks).

We now summarize some of the basic points to keep in mind when dealing with ∞-categories.

- Any ordinary category may be considered as an ∞-category: one takes each of the morphism spaces $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ to be discrete.
- By a *morphism* from X to Y in $\mathcal{C}$, we mean a point of the space $\mathrm{Hom}_{\mathcal{C}}(X, Y)$. Two morphisms are *equivalent*, or *homotopic*, if they lie in the same path component of $\mathrm{Hom}_{\mathcal{C}}(X, Y)$.
- If $f : C \to C'$ is a morphism in an ∞-category $\mathcal{C}'$, then we say that f is an *equivalence* if it becomes an isomorphism in the homotopy category $h\,\mathcal{C}$. In other words, f is an equivalence if it has a homotopy inverse.
- Any ∞-category $\mathcal{C}$ has an *opposite ∞-category* $\mathcal{C}^{op}$, having the same objects but with $\mathrm{Hom}_{\mathcal{C}^{op}}(X, Y) = \mathrm{Hom}_{\mathcal{C}}(Y, X)$.
- Given an ∞-category $\mathcal{C}$ and an object $X \in \mathcal{C}$, we can define a *slice ∞-category* $\mathcal{C}_{/X}$. The objects of $\mathcal{C}_{/X}$ are pairs (Y, f) with $Y \in \mathcal{C}$ and $f \in \mathrm{Hom}_{\mathcal{C}}(Y, X)$. The space of morphisms from $\mathrm{Hom}_{\mathcal{C}_{/X}}((Y, f), (Y', f'))$ is given by the homotopy fiber of the map $\mathrm{Hom}_{\mathcal{C}}(Y, Y') \stackrel{f'\circ}{\to} \mathrm{Hom}_{\mathcal{C}}(Y, X)$ over the point f.

  If the category $\mathcal{C}$ and the maps $f : Y \to X$, $f' : Y' \to X$ are clear from context, we will often write $\mathrm{Hom}_X(Y, Y')$ for $\mathrm{Hom}_{\mathcal{C}_{/X}}((Y, f), (Y', f'))$.

  Remark 1.2.6. Let $\mathcal{C}_0$ be a category which is enriched over topological spaces (in other words, the morphism sets in $\mathcal{C}$ are equipped with topologies such that the composition laws are all continuous), and let $\mathcal{C}$ denote the ∞-category which it represents. Given an object $X \in \mathcal{C}_0$, one may form the classical slice-category construction to obtain a category $\mathcal{C}_{0/X}$, which is again enriched over topological spaces. The associated ∞-category is not necessarily equivalent to the slice ∞-category $\mathcal{C}_{/X}$. The difference is that the morphism spaces in $\mathcal{C}_{/X}$ must be formed using homotopy fibers of maps $\phi : \mathrm{Hom}_{\mathcal{C}}(Y, Y') \to \mathrm{Hom}_{\mathcal{C}}(Y, X)$, while the morphism spaces in $\mathcal{C}_{0/X}$ are given by taking the ordinary fibers of the maps ϕ. If we want to construct an explicit model for $\mathcal{C}_{/X}$ as a topological category, then we need to choose an explicit construction for the homotopy fiber of a continuous map $\phi : E \to E'$. The standard construction for the homotopy fiber of ϕ over a point $e' \in E'$ is the space $\{(e, p) : e \in E, p : [0, 1] \to E', p(0) = \phi(e), p(1) = e'\}$. However, if one uses this construction to define the morphism spaces in $\mathcal{C}_{/X}$, then it is not possible to define a *strictly* associative composition law on morphisms. However, there exists a composition law which is associative up to *coherent* homotopy. This may then be replaced by a strictly associative composition law after slightly altering the morphism spaces.

  The same difficulties arise with virtually every construction we shall encounter. In order to perform some categorical construction on an ∞-category, it is not enough to choose a topological (or simplicial) category which models it and employ the same construction in the naive sense. However, this is merely a technical annoyance and not a serious problem: there are many strategies for dealing with this, one (at least) for each of the main definitions of ∞-categories. We will henceforth ignore these issues; the reader may refer to the literature for more detailed discussions and constructions.

- We call an ∞-category *small* if it has a set of objects, just as with ordinary categories. We call an ∞-category *essentially small* if the collection of equivalence classes of objects forms a set.
- Let $\mathcal{C}$ and $\mathcal{C}'$ be ∞-categories. One can define a notion of functor $F : \mathcal{C} \to \mathcal{C}'$. A functor F carries objects of $C \in \mathcal{C}$ to objects $FC \in \mathcal{C}'$, and induces maps between the morphism spaces $\mathrm{Hom}_{\mathcal{C}}(C, D) \to \mathrm{Hom}_{\mathcal{C}'}(FC, FD)$. These maps should be compatible with the composition operations, up to coherent homotopy. See [8] for an appropriate definition, in the context of simplicial categories. We remark that if we are given explicit models for $\mathcal{C}$ and $\mathcal{C}'$ as simplicial or topological categories, then we do not expect every functor F to arise from an ordinary functor (in other words, we do not expect to

be able to arrange the situation so that F is compatible with composition "on the nose"), although any ordinary functor which is compatible with a simplicial or topological enrichment does give rise to a functor between the associated ∞-categories.

- If $\mathcal{C}$ and $\mathcal{C}'$ are two small ∞-categories, then there is an ∞-category $\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$ of functors from $\mathcal{C}$ to $\mathcal{C}'$. We will sometimes write ${\mathcal{C}'}^{\mathcal{C}}$ for $\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$. We may summarize the situation by saying that there is an $(\infty, 2)$-category of small ∞-categories.
- We will also need to discuss functor categories $\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$ in cases when $\mathcal{C}$ and $\mathcal{C}'$ are not small. So long as $\mathcal{C}$ is small, this poses little difficulty: we must simply bear in mind that if $\mathcal{C}'$ is large, the functor category is also large. Care must be taken if $\mathcal{C}$ is also large: in this case $\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$ is ill-defined because its collection of objects could be *very large* (possibly not even a proper class), and the space of natural transformations between two functors might also be large. In practice, we will avoid this difficulty by only considering functors satisfying certain continuity conditions. After imposing these restrictions, $\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$ will be an honest ∞-category (though generally still a large ∞-category).
- Given an ∞-category $\mathcal{C}$ and some collection of objects S of $\mathcal{C}$, we can form an ∞-category $\mathcal{C}_0$ having the elements of S as objects, and the same morphism spaces as $\mathcal{C}$. We shall refer to those ∞-categories $\mathcal{C}_0$ which arise in this way *full subcategories* of $\mathcal{C}$. (Our use of "full subcategory" as opposed to "full sub-∞-category" is to avoid awkward language and is not intended to suggest that $\mathcal{C}_0$ is an ordinary category).
- A functor $F : \mathcal{C} \to \mathcal{C}'$ is said to be fully faithful if for any objects $C, C' \in \mathcal{C}$, the induced map $\mathrm{Hom}_{\mathcal{C}}(C, C') \to \mathrm{Hom}_{\mathcal{C}'}(FC, FC')$ is a homotopy equivalence. We say that F is *essentially surjective* if every object of $\mathcal{C}'$ is equivalent to an object of the form FC, $C \in \mathcal{C}$. We say that F is an *equivalence* if it is fully faithful and essentially surjective. In this case, one can construct a homotopy inverse for F if one assumes a sufficiently strong version of the axiom of choice.

  We shall understand that all relevant properties of ∞-categories are invariant under equivalence. For example, the property of being small is not invariant under equivalence, and is therefore not as natural as the property of being essentially small. We note that an ∞-category is essentially small if and only if it is equivalent to a small ∞-category.
- If a functor $F : \mathcal{C} \to \mathcal{C}'$ is fully faithful, then it establishes an equivalence of $\mathcal{C}$ with a full subcategory of $\mathcal{C}'$. In this situation, we may use F to identify $\mathcal{C}$ with its image in $\mathcal{C}'$.
- If X is an object in an ∞-category $\mathcal{C}$, then we sa that X is *initial* if $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ is contractible for all $Y \in \mathcal{C}$. The dual notion of a *final* object is defined in the evident way.
- If $\mathcal{C}$ is an ∞-category, then a *diagram* in $\mathcal{C}$ is a functor $F : \mathcal{I} \to \mathcal{C}$, where $\mathcal{I}$ is a small ∞-category. Just as with diagrams in ordinary categories, we may speak of limits and colimits of diagrams, which may or may not exist. These may be characterized in the following way (we restrict our attention to colimits; for limits, just dualize everything): a colimit of F is an object $C \in \mathcal{C}$ such that there exists a natural transformation $F \to F_C$ which induces equivalences $\mathrm{Hom}_{\mathcal{C}^{\mathcal{I}}}(F, F_D) \simeq \mathrm{Hom}_{\mathcal{C}}(C, D)$, where F_C and F_D denote the constant functors $\mathcal{I} \to \mathcal{C}$ having the values $C, D \in \mathcal{C}$. As we have explained, these are *homotopy-theoretic* limits and colimits, and do not necessarily enjoy any universal property in $h\mathcal{C}$.
- In the theory of ordinary categories (and in mathematics in general), the category of sets plays a pivotal role. In the ∞-categorical setting, the analogous role is filled by the ∞-category $\mathcal{S}$ of spaces. One can take the objects of $\mathcal{S}$ to be any suitable model for homotopy theory, such as CW complexes or fibrant simplicial sets: in either case, there are naturally associated "spaces" of morphisms, which may themselves be interpreted as objects of $\mathcal{S}$.
- An ∞-category is an ∞-*groupoid* if all of its morphisms are equivalences. A small ∞-groupoid is essentially the same thing as an object of $\mathcal{S}$.

Remark 1.2.7. This assertion has some content: for ∞-groupoids with a single object $*$, it reduces to Stasheff's theorem that the coherently associative composition on $X = \mathrm{Hom}(*, *)$ is precisely the data required to realize X as a loop space (see for example [25], where a version of this theorem is proved by establishing the equivalence of the homotopy theory of simplicial sets and simplicial groupoids).

- For any ∞-category $\mathcal{C}$, we define a *prestack* on $\mathcal{C}$ to be a functor $F : \mathcal{C}^{\mathrm{op}} \to \mathcal{S}$. We remark that this is not standard terminology: for us, prestacks and stacks will always be valued in ∞-groupoids, rather than in ordinary groupoids or in categories.
- Any object $C \in \mathcal{C}$ gives rise to a prestack $\mathcal{F}_C$ given by the formula $\mathcal{F}_C(C') = \mathrm{Hom}_{\mathcal{C}}(C', C)$. This induces a functor from $\mathcal{C}$ to the ∞-category $\mathcal{S}^{\mathcal{C}^{op}}$ of prestacks on $\mathcal{C}$, which generalizes the classical Yoneda embedding for ordinary categories. As with the classical Yoneda embedding, this functor is fully faithful; prestacks which lie in its essential image are called *representable*.
- The $(\infty, 2)$-category of ∞-categories has all $(\infty, 2)$-categorical limits and colimits. We do not want to formulate a precise statement here, but instead refer the reader to Appendix 5.1 for some discussion. Let us remark here that forming *limits* is relatively easy: one can work with them "componentwise," just as with ordinary categories.

  Example 1.2.8. Given any family $\{\mathcal{C}_\alpha\}$ of ∞-categories, their product $\mathcal{C} = \prod_\alpha \mathcal{C}_\alpha$ may be constructed in the following way: the objects of $\mathcal{C}$ consist of a choice of one object C_α from each $\mathcal{C}_\alpha$, and the morphism spaces $\mathrm{Hom}_{\mathcal{C}}(\{C_\alpha\}, \{C'_\alpha\}) = \prod_\alpha \mathrm{Hom}_{\mathcal{C}_\alpha}(C_\alpha, C'_\alpha)$.

  Example 1.2.9. Suppose that $F : \mathcal{C}' \to \mathcal{C}$ and $G : \mathcal{C}'' \to \mathcal{C}$ are functors between ∞-categories. One may form an ∞-category $\mathcal{C}' \times_{\mathcal{C}} \mathcal{C}''$, the *strict fiber product* of $\mathcal{C}'$ and $\mathcal{C}''$ over $\mathcal{C}$. An object of $\mathcal{C}' \times_{\mathcal{C}} \mathcal{C}''$ consists of a pair (C', C'', η) where $C' \in \mathcal{C}'$, $C'' \in \mathcal{C}''$, and $\eta : FC' \to FC''$ is an equivalence in $\mathcal{C}$. If one drops the requirement that η be an equivalence, then one obtains instead the *lax fiber product* of $\mathcal{C}'$ and $\mathcal{C}''$ over $\mathcal{C}$.

  Example 1.2.10. Let $\mathcal{C}$ and $\mathcal{I}$ be ∞-categories. Then the functor ∞-category $\mathcal{C}^{\mathcal{I}}$ may be regarded as a limit of copies of the ∞-category $\mathcal{C}$, where the limit is indexed by the ∞-category $\mathcal{I}$.

  Example 1.2.11. Let $\mathcal{C}$ be an ∞-category and $E \in \mathcal{C}$ an object. Then we may regard E as a determining a functor $* \to \mathcal{C}$, where $*$ is an ∞-category with a single object having a contractible space of endomorphisms. The slice category $\mathcal{C}_{/E}$ may be regarded as a lax fiber product of $\mathcal{C}$ and $*$ over $\mathcal{C}$ (via the functors E and the identity).

- One may describe ∞-categories by "generators and relations". In particular, it makes sense to speak of a *finitely presented* ∞-category. Such an ∞-category has finitely many objects and its morphism spaces are determined by specifying a finite number of generating morphisms, a finite number of relations among these generating morphisms, a finite number of relations among the relations, and so forth (a finite number of relations in all).

  Example 1.2.12. Let $\mathcal{C}$ be the free ∞-category generated by a single object X and a single morphism $f : X \to X$. Then $\mathcal{C}$ is a finitely presented ∞-category with a single object, and $\mathrm{Hom}_{\mathcal{C}}(X, X) = \{1, f, f^2, \ldots\}$ is infinite discrete. In particular, we note that the finite presentation of $\mathcal{C}$ does not guarantee finiteness properties of the morphism spaces.

  Example 1.2.13. If we identify ∞-groupoids with spaces, then writing down a presentation for an ∞-groupoid corresponds to giving a cell decomposition of the associated space. We therefore see that the finitely presented ∞-groupoids correspond precisely to the finite cell complexes.

  Example 1.2.14. Suppose that $\mathcal{C}$ is an ∞-category with only two objects X and Y, and that X and Y have contractible endomorphism spaces and that $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ is empty. Then $\mathcal{C}$ is completely determined by the morphism space $\mathrm{Hom}_{\mathcal{C}}(Y, X)$, which may be arbitrary. The ∞-category $\mathcal{C}$ is finitely presented if and only if $\mathrm{Hom}_{\mathcal{C}}(Y, X)$ is a finite cell complex (up to homotopy equivalence).

- Given two ∞-categories $\mathcal{C}$ and $\mathcal{C}'$ and functors $F : \mathcal{C} \to \mathcal{C}'$, $G : \mathcal{C}' \to \mathcal{C}$, we have two functors

$$\mathrm{Hom}_{\mathcal{C}}(X, GY), \mathrm{Hom}_{\mathcal{C}'}(FX, Y) : \mathcal{C}^{op} \times \mathcal{C}' \to \mathcal{S}$$

An equivalence between these two functors is called an *adjunction* between F and G: we say that F is *left adjoint* to G and that G is *right adjoint* to F. As with ordinary categories, adjoints are unique up to canonical equivalence when they exist.

- As with ordinary categories, many of the ∞-categories which arise in nature are large. Working with these objects creates foundational technicalities. We will generally ignore these technicalities, which can be resolved in many different ways.

1.3. **Accessibility.** Most of the categories which arise in nature (such as the category of groups) have a proper class of objects, even when taken up to isomorphism. However, there is a sense in which many of these large categories are determined by a bounded amount of information. For example, in the category of groups, every object may be represented as a filtered colimit of finitely presented groups, and the category of finitely presented groups is essentially small.

This phenomenon is common to many examples, and may be formalized in the notion of an *accessible category* (see for example [1] or [3]). In this section, we shall adapt the definition of accessible category and the basic facts concerning them to the ∞-categorical setting. We shall not assume that the reader is familiar with the classical notion of an admissible category.

Our reasons for studying accessibility are twofold: first, we would prefer to deal with ∞-categories as legitimate mathematical objects, without appealing to a hierarchy of universes (in the sense of Grothendieck) or some similar device. Second, we sometimes wish to make arguments which play off the difference between a “large” category and a “small” piece which determines it, in order to make sense of constructions which would otherwise be ill-defined. See, for example, the proof of Theorem 1.4.3.

Before we begin, we are going to need to introduce some size-related concepts. Recall that a cardinal κ is *regular* if κ cannot be represented as a sum of fewer than κ cardinals of size $< \kappa$.

Definition 1.3.1. Let κ be a regular cardinal. If $\kappa > \omega$, then an ∞-category $\mathcal{C}$ is called *κ-small* it has fewer than κ objects, and for any pair of objects X, Y, the space $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ may be presented using fewer than κ cells. If $\kappa = \omega$, then an ∞-category $\mathcal{C}$ is *κ-small* if it is finitely presented.

Remark 1.3.2. One may unify the two clauses of the definition as follows: an ∞-category $\mathcal{C}$ is κ-small if it may be presented by fewer than κ “generators and relations”. The reason that the above definition needs a special case for $\kappa = \omega$ is that a finitely presented category ∞-category can have morphism spaces which are not finite (see Example 1.2.12).

Definition 1.3.3. An ∞-category $\mathcal{C}$ is *κ-filtered* if for any diagram $F : \mathcal{I} \to \mathcal{C}$ where $\mathcal{I}$ is κ-small, there exists a natural transformation from F to a constant functor. We say that $\mathcal{C}$ is *filtered* if it is ω-filtered. A diagram $\mathcal{I} \to \mathcal{C}$ is called *κ-filtered* if $\mathcal{I}$ is κ-filtered, and *filtered* if $\mathcal{I}$ is ω-filtered.

As the definition suggests, the most important case occurs when $\kappa = \omega$.

Remark 1.3.4. Filtered ∞-categories are the natural generalization of filtered categories, which are in turn a mild generalization of directed partially ordered sets. Recall that a partially ordered set P is *directed* if every finite subset of P has some upper bound in P. One frequently encounters diagrams indexed by directed partially ordered sets, for example by the directed partially ordered set $\mathbb{Z}_{\geq 0} = \{0, 1, \ldots\}$ of natural numbers: colimits

$$\mathrm{colim}_{n \to \infty} X_n$$

of such diagrams are among the most familiar constructions in mathematics.

In classical category theory, it is convenient to consider not only diagrams indexed by directed partially ordered sets, but also more generally diagram indexed by filtered categories. A filtered category is defined to be a category $\mathcal{C}$ satisfying the following conditions:

- There exists at least one object in $\mathcal{C}$.
- Given any two objects $X, Y \in \mathcal{C}$, there exists a third object Z and morphisms $X \to Z$, $Y \to Z$.
- Given any two morphisms $f, g : X \to Y$ in $\mathcal{C}$, there exists a morphism $h : Y \to Z$ such that $h \circ f = h \circ g$.

The first two conditions are analogous to the requirement that any finite piece of $\mathcal{C}$ has an "upper bound", while the last condition guarantees that the upper bound is uniquely determined in some asymptotic sense.

One can give a similar definition of filtered ∞-categories, but one must assume a more general version of the third condition. A pair of maps $f, g : X \to Y$ may be regarded as a map $S^0 \to \mathrm{Hom}_{\mathcal{C}}(X, Y)$, where S^0 denotes the zero-sphere. The third condition asserts that there exists a map $Y \to Z$ such that the induced map $S^0 \to \mathrm{Hom}_{\mathcal{C}}(X, Z)$ extends to a map $D^1 \to \mathrm{Hom}_{\mathcal{C}}(X, Z)$, where D^1 is the 1-disk. To arrive at the definition of a filtered ∞-category, we must require the existence of $Y \to Z$ not only for maps $S^0 \to \mathrm{Hom}_{\mathcal{C}}(X, Y)$, but for any map $S^n \to \mathrm{Hom}_{\mathcal{C}}(X, Y)$ (and the requirement should be that the induced map extends over the disk D^{n+1}). With this revision, the above conditions are equivalent to our definition of ω-filtered ∞-categories.

The reader should not be intimidated by the apparent generality of the notion of a κ-filtered colimit. Given any κ-filtered diagram $\mathcal{I} \to \mathcal{C}$, one can always find a diagram $\mathcal{J} \to \mathcal{C}$ indexed by a partially ordered set $\mathcal{J}$ which is "equivalent" in the sense that they have the same colimit (provided that either colimit exists), and $\mathcal{J}$ is κ-filtered in the sense that every subset of $\mathcal{J}$ having cardinality $< \kappa$ has an upper bound in $\mathcal{J}$. We sketch a proof for this in Appendix 5.1.

Definition 1.3.5. We will call an ∞-category $\mathcal{C}$ *κ-closed* if every κ-filtered diagram has a colimit in $\mathcal{C}$. A functor between κ-closed ∞-categories will be called κ-continuous if it commutes with the formation of κ-filtered colimits. Finally, we will call an object C in a κ-closed ∞-category $\mathcal{C}$ *κ-compact* if the functor $\mathrm{Hom}_{\mathcal{C}}(C, \bullet)$ is a κ-continuous functor $\mathcal{C} \to \mathcal{S}$. When $\kappa = \omega$, we will abbreviate by simply referring to *continuous* functors and *compact* objects.

Definition 1.3.6. If $\mathcal{C}$ is any κ-closed ∞-category, we will let $\mathcal{C}_\kappa$ denote the full subcategory consisting of κ-compact objects.

Remark 1.3.7. Any κ-closed category (κ-continuous morphism, κ-compact object) is also κ'-closed (κ'-continuous, κ'-compact) for any $\kappa' \geq \kappa$.

Remark 1.3.8. There are a number of reasons to place emphasis on κ-filtered colimits (particularly in the case where $\kappa = \omega$). They exist more often and tend to be more readily computable than colimits in general. For example, the category of groups admits all colimits, but it is usually difficult to give an explicit description of a colimit (typically this involves solving a word problem). However, filtered colimits of groups are easy to describe, because the formation of filtered colimits is compatible with passage to the underlying set of a group. In other words, the forgetful functor F, which assigns to each group its underlying set, is continuous.

We now come to a key construction. Let $\mathcal{C}$ be a ∞-category and κ a regular cardinal. We will define a new ∞-category $\mathrm{Ind}_\kappa(\mathcal{C})$ as follows. The objects of $\mathrm{Ind}_\kappa(\mathcal{C})$ are small, κ-filtered diagrams $\mathcal{I} \to \mathcal{C}$. We imagine that an object of $\mathrm{Ind}_\kappa(\mathcal{C})$ is a *formal* κ-filtered colimit of objects in $\mathcal{C}$, and use the notation

$$\text{“}\operatorname{colim}_{D \in \mathcal{I}} F(D)\text{”}$$

to denote the object corresponding to a diagram $F : \mathcal{I} \to \mathcal{C}$. The morphism spaces are given by

$$\mathrm{Hom}_{\mathrm{Ind}_\kappa(\mathcal{C})}(\text{“}\operatorname{colim}_{D \in \mathcal{I}} F(D)\text{”}, \text{“}\operatorname{colim}_{D' \in \mathcal{I}'} F'(D')\text{”}) = \lim_{D \in \mathcal{I}} \operatorname{colim}_{D' \in \mathcal{I}'} \mathrm{Hom}_{\mathcal{C}}(FD, F'D')$$

When $\kappa = \omega$, we will sometimes write $\mathrm{Ind}(\mathcal{C})$ instead of $\mathrm{Ind}_\kappa(\mathcal{C})$. For a more detailed discussion of the homotopy theory of filtered diagrams, from the point of view of model categories, we refer the reader to [9].

Remark 1.3.9. If $\mathcal{C}$ is an ordinary category, then so is $\mathrm{Ind}_\kappa(\mathcal{C})$. Many categories have the form $\mathrm{Ind}(\mathcal{C})$. For example, the category of groups is equivalent to $\mathrm{Ind}(\mathcal{C})$, where $\mathcal{C}$ is the category of finitely presented groups. A similar remark applies to any other sort of (finitary) algebraic structure.

There is a functor $\mathcal{C} \to \mathrm{Ind}_\kappa(\mathcal{C})$, which carries an object $C \in \mathcal{C}$ to the constant diagram $* \to \mathcal{C}$ having value C, where $*$ denotes a category having a single object with a contractible space of endomorphisms. This functor is fully faithful, and identifies $\mathcal{C}$ with a full subcategory of $\mathrm{Ind}_\kappa(\mathcal{C})$. One can easily verify that

$\mathrm{Ind}_\kappa(\mathcal{C})$ is κ-closed, and that the essential image of $\mathcal{C}$ consists of κ-compact objects. In fact, $\mathrm{Ind}_\kappa(\mathcal{C})$ is universal with respect to these properties:

Proposition 1.3.10. *Let $\mathcal{C}$ and $\mathcal{C}'$ be ∞-categories, where $\mathcal{C}'$ is κ-closed. Then restriction induces an equivalence of ∞-categories*

$$\mathrm{Fun}^\kappa(\mathrm{Ind}_\kappa(\mathcal{C}), \mathcal{C}') \to \mathrm{Fun}(\mathcal{C}, \mathcal{C}')$$

where the left hand side denotes the ∞-category of κ-continuous functors. In other words, any functor $F : \mathcal{C} \to \mathcal{C}'$ extends uniquely to a κ-continuous functor $\widetilde{F} : \mathrm{Ind}_\kappa(\mathcal{C}) \to \mathcal{C}'$. If F is fully faithful and its image consists of κ-compact objects, then $\widetilde{F}$ is also fully faithful.

Proof. This follows more or less immediately from the construction. □

From this, we can easily deduce the following:

Proposition 1.3.11. *Let $\mathcal{C}$ be an ∞-category. The following conditions are equivalent:*

- *The ∞-category $\mathcal{C}$ is equivalent to $\mathrm{Ind}_\kappa \mathcal{C}_0$ for some small ∞-category $\mathcal{C}_0$.*
- *The ∞-category $\mathcal{C}$ is κ-closed, $\mathcal{C}_\kappa$ is essentially small, and the functor $\mathrm{Ind}_\kappa(\mathcal{C}_\kappa) \to \mathcal{C}$ is an equivalence of categories.*
- *The ∞-category $\mathcal{C}$ is κ-closed, and has a small subcategory $\mathcal{C}_0$ consisting of κ-compact objects such that every object of $\mathcal{C}$ can be obtained as a κ-filtered colimit of objects in $\mathcal{C}_0$.*

An ∞-category $\mathcal{C}$ satisfying the equivalent conditions of Proposition 1.3.11 will be called *κ-accessible*.

The only difficulty in proving Proposition 1.3.11 is in verifying that if $\mathcal{C}_0$ is small, then $\mathrm{Ind}_\kappa(\mathcal{C}_0)$ has only a set of κ-compact objects, up to equivalence. It is tempting to guess that any such object must be equivalent to an object of $\mathcal{C}_0$. The following example shows that this is not necessarily the case.

Example 1.3.12. Let R be a ring, and let $\mathcal{C}_0$ denote the (ordinary) category of finitely generated free R-modules. Then $\mathcal{C} = \mathrm{Ind}(\mathcal{C}_0)$ is equivalent to the category of flat R-modules (by Lazard's theorem; see for example the appendix of [10]). The compact objects of $\mathcal{C}$ are precisely the finitely generated projective R-modules, which need not be free.

However, a κ-compact object C in $\mathrm{Ind}_\kappa(\mathcal{C}_0)$ is not far from being an object of $\mathcal{C}_0$: it follows readily from the definition that C is a *retract* of an object of $\mathcal{C}_0$, meaning that there are maps

$$C \xrightarrow{f} D \xrightarrow{g} C$$

with $D \in \mathcal{C}_0$ and $g \circ f$ homotopic to the identity of C. In this case, we may recover C given the object D and appropriate additional data. Since $g \circ f$ is homotopic to the identity, we deduce easily that the morphism $r = f \circ g$ is idempotent in the homotopy category $h\mathcal{C}$. Moreover, in $h\mathcal{C}$, the object C is the equalizer of r and the identity morphism of D. This shows that there are only a bounded number of possibilities for C, and proves Proposition 1.3.11.

Definition 1.3.13. An ∞-category $\mathcal{C}$ is *accessible* if it is κ-accessible for some regular cardinal κ. A functor between accessible ∞-categories is *accessible* if it is κ-continuous for some κ (and therefore for all sufficiently large κ).

It is not necessarily true that a κ-accessible ∞-category $\mathcal{C}$ is κ' accessible for all $\kappa' > \kappa$. However, it is true that $\mathcal{C}$ is κ'-accessible for many other cardinals κ'. Let us write $\kappa' \gg \kappa$ if $\tau^\kappa < \kappa'$ for any $\tau < \kappa'$. Note that there exist arbitrarily large regular cardinals κ' with $\kappa' \gg \kappa$: for example, one may take κ' to be the successor of any cardinal having the form τ^κ.

Lemma 1.3.14. *If $\kappa' \gg \kappa$, then any κ'-filtered partially ordered set $\mathcal{I}$ may be written as a union of κ-filtered subsets having size $< \kappa'$. Moreover, the family of such subsets is κ'-filtered.*

Proof. It will suffice to show that every subset of $S \subseteq \mathcal{I}$ having cardinality $< \kappa'$ can be included in a larger subset S', such that $|S'| < \kappa'$, but S' is κ-filtered.

We define a transfinite sequence of subsets $S_\alpha \subseteq \mathcal{I}$ by induction. Let $S_0 = S$, and when λ is a limit ordinal we let $S_\lambda = \bigcup_{\alpha<\lambda} S_\alpha$. Finally, we let $S_{\alpha+1}$ denote a set which is obtained from S_α by adjoining an upper bound for every subset of S_α having size $< \kappa$ (which exists because $\mathcal{I}$ is κ'-filtered). It follows from the assumption $\kappa' \gg \kappa$ that if S_α has size $< \kappa'$, then so does $S_{\alpha+1}$. Since κ' is regular, we deduce easily by induction that $|S_\alpha| < \kappa'$ for all $\alpha < \kappa'$. It is easy to check that the set $S' = S_\kappa$ has the desired properties. □

We now show that accessible categories enjoy several pleasant boundedness properties:

Proposition 1.3.15. *Let $\mathcal{C}$ be an accessible category. Then every object of $\mathcal{C}$ is κ-compact for some regular cardinal κ. For any regular cardinal κ, the full subcategory of $\mathcal{C}$ consisting of κ-compact cardinals is essentially small.*

Proof. Suppose $\mathcal{C}$ is τ-accessible. Then every object of $\mathcal{C}$ may be written as a τ-filtered colimit of τ-compact objects. But it is easy to see that for $\kappa \geq \tau$, the colimit of a κ-small diagram of τ-compact object is κ-compact.

Now fix a regular cardinal $\kappa \gg \tau$, and let C be a κ-compact object of $\mathcal{C}$. Then C may be written as a τ-filtered colimit of a diagram $\mathcal{I} \to \mathcal{C}_\tau$. Without loss of generality, we may assume that $\mathcal{I}$ is actually the ∞-category associated to a partially ordered set. By Lemma 1.3.14, we may write $\mathcal{I}$ as a κ-filtered union of τ-filtered, κ-small subsets $\mathcal{I}_\alpha$. Let C_α denote the colimit of the diagram indexed by $\mathcal{I}_\alpha$; then $C = \operatorname{colim}_\alpha C_\alpha$ where *this* colimit is κ-filtered. Consequently, C is a retract of some C_α. Since there are only a bounded number of possibilites for $\mathcal{I}_\alpha$ and its functor to $\mathcal{C}_\tau$, we see that there are only a bounded number of possibilities for C, up to equivalence. □

Remark 1.3.16. It is not really necessary to reduce to the case where $\mathcal{I}$ is a partially ordered set. We could instead extend Lemma 1.3.14 to ∞-categories: the difficulties are primarily notational.

Proposition 1.3.17. *Let $\mathcal{C}$ be a κ-accessible ∞-category. Then $\mathcal{C}$ is κ'-accessible for any $\kappa' \gg \kappa$.*

Proof. It is clear that $\mathcal{C}$ is κ'-closed. Proposition 1.3.15 implies that $\mathcal{C}_{\kappa'}$ is essentially small. To complete the proof, it suffices to show that any object $C \in \mathcal{C}$ can be obtained as the colimit of a κ'-filtered diagram in $\mathcal{C}_{\kappa'}$. By assumption, C can be represented as the colimit of a κ-filtered diagram $F : \mathcal{I} \to \mathcal{C}_\kappa$. Without loss of generality, we may assume that $\mathcal{I}$ is a partially ordered set.

By Lemma 1.3.14, we may write $\mathcal{I}$ as a κ'-filtered union of κ'-small, κ-filtered partially ordered sets $\mathcal{I}_\alpha$. Let $F_\alpha = F|\mathcal{I}_\alpha$. Then we get

$$C = \operatorname{colim} F = \operatorname{colim}_\alpha(\operatorname{colim} F_\alpha)$$

Since each $\mathcal{I}_\alpha$ is κ'-small, the colimit of each F_α is κ'-compact. It follows that C is a κ'-filtered colimit of κ'-small objects, as desired. □

We now show that accessible functors between accessible ∞-categories are themselves determined by a bounded amount of data.

Proposition 1.3.18. *Let $F : \mathcal{C} \to \mathcal{C}'$ be an accessible functor between accessible ∞-categories. Then there exists a regular cardinal κ such that $\mathcal{C}$ and $\mathcal{C}'$ are κ-accessible, and F is the unique κ-continuous extension of a functor $F_\kappa : \mathcal{C}_\kappa \to \mathcal{C}'_\kappa$.*

Proof. Choose τ so large that F is τ-continuous. Enlarging τ if necessary, we may assume that $\mathcal{C}$ and $\mathcal{C}'$ are τ-accessible. Then $\mathcal{C}$ has only a bounded number of τ-compact objects. By Proposition 1.3.15, the images of these objects under F are all τ'-compact for some regular cardinal $\tau' \geq \tau$. We may now choose κ to be any regular cardinal such that $\kappa \gg \tau'$.

Any κ-compact object of $\mathcal{C}$ can be written as a κ-small, τ-filtered colimit of τ-compact objects. Since F preserves τ-filtered colimits, the image of a κ-compact of $\mathcal{C}$ is a κ-small, τ-filtered colimit of τ'-compact objects, hence κ-compact. Thus the restriction of F maps $\mathcal{C}_\kappa$ into $\mathcal{C}'_\kappa$. To complete the proof, it suffices to note that the τ-continuity of F implies κ-continuity since $\kappa > \tau$. □

The above proposition shows that there is an honest ∞-category of accessible functors from $\mathcal{C}$ to $\mathcal{C}'$, for any accessible ∞-categories $\mathcal{C}$ and $\mathcal{C}'$. Indeed, this ∞-category is just a (filtered) union of the essentially small ∞-categories $\operatorname{Hom}(\mathcal{C}_\kappa, \mathcal{C}'_\kappa)$, where κ ranges over regular cardinals such that $\mathcal{C}$ and $\mathcal{C}'$ are κ-accessible. We will denote this ∞-category by $\operatorname{Acc}(\mathcal{C}, \mathcal{C}')$. Any composition of accessible functors is accessible, so that we obtain a (large) $(\infty, 2)$-category Acc consisting of accessible ∞-categories and accessible functors.

Proposition 1.3.19. *The* $(\infty, 2)$*-category* Acc *has all* $(\infty, 2)$*-categorical limits, which are simply given by taking the limits of the underlying* ∞*-categories.*

In other words, a limit of accessible ∞-categories (and accessible functors) is accessible.

Proof. We may choose a regular cardinal κ such that all of the accessible categories are κ-accessible and all of the relevant functors between them are κ-continuous. Let $\mathcal{C}$ denote the limit in question. It is clear that $\mathcal{C}$ is κ-closed: one may compute κ-filtered colimits in $\mathcal{C}$ componentwise (since all of the relevant functors are κ-continuous). To complete the proof, we must show that (possibly after enlarging κ), the limit is generated by a set of κ-compact objects. The proof of this point involves tedious cardinality estimates, and will be omitted. The reader may consult [1] for a proof in the 1-categorical case, which contains all of the relevant ideas. □

Example 1.3.20. Let $\mathcal{C}$ be an accessible ∞-category, and let X be an object of $\mathcal{C}$. Then $\mathcal{C}_{/X}$ is an accessible ∞-category, since it is a lax pullback of accessible ∞-categories (see Example 1.2.11).

We conclude this section with the following useful observation:

Proposition 1.3.21. *Let* $G : \mathcal{C} \to \mathcal{C}'$ *be a functor between accessible* ∞*-categories. If* G *admits an adjoint, then* G *is accessible.*

Proof. If G is a left adjoint, then G commutes with all colimits and is therefore ω-continuous. Let us therefore assume that G is right adjoint to some functor F.

Choose κ so that $\mathcal{C}$ and $\mathcal{C}'$ are κ-accessible. Now choose $\kappa' \gg \kappa$ large enough that FC' is κ'-compact for all $C' \in \mathcal{C}_\kappa$. We claim that G is κ'-continuous.

Let $\pi : \mathcal{I} \to \mathcal{C}$ be any κ'-filtered diagram having colimit C. We must show that GC is the colimit of the diagram $G \circ \pi$.

We need to show that $\mathrm{Hom}_{\mathcal{C}'}(C', GC)$ is equivalent to $\mathrm{Hom}_{\mathcal{C}'}(C', \mathrm{colim}_{C_\alpha \in \mathcal{I}}\{GC_\alpha\})$ for all objects $C' \in \mathcal{C}'$. It suffices to check this for C' in $\mathcal{C}'_\kappa$, since every object in $\mathcal{C}'$ is a colimit of objects in $\mathcal{C}'_\kappa$. Then

$$\begin{aligned}
\mathrm{Hom}_{\mathcal{C}'}(C', \mathrm{colim}_{C_\alpha \in \mathcal{I}}\{GC_\alpha\}) &= \mathrm{colim}_{C_\alpha \in \mathcal{I}} \mathrm{Hom}_{\mathcal{C}'}(C', GC_\alpha) \\
&= \mathrm{colim}_{C_\alpha \in \mathcal{I}} \mathrm{Hom}_{\mathcal{C}}(FC', C_\alpha) \\
&= \mathrm{Hom}_{\mathcal{C}}(FC', C) \\
&= \mathrm{Hom}_{\mathcal{C}'}(C', GC),
\end{aligned}$$

as required. □

1.4. **Presentable ∞-Categories.**

In this section, we discuss the notion of a *presentable* ∞*-category*. Roughly speaking, this is an ∞-category with a set of small generators that admits arbitrary colimits. This includes, for example, any ordinary category with a set of small generators that admits arbitrary colimits: such categories are called locally presentable in [1]. The definition, in the ∞-categorical setting, was given by Carlos Simpson in [4], who calls them ∞-pretopoi.

Proposition 1.4.1. *Let* $\mathcal{C}$ *be an* ∞*-category. Then the following are equivalent:*

(1) *The* ∞*-category* $\mathcal{C}$ *is accessible and admits arbitrary colimits.*
(2) *For all sufficiently large regular cardinals* κ*, the* ∞*-category* $\mathcal{C}_\kappa$ *is essentially small, admits colimits for all* κ*-small diagrams, and the induced* $\mathrm{Ind}_\kappa(\mathcal{C}_\kappa) \to \mathcal{C}$ *is an equivalence.*
(3) *There exists a regular cardinal* κ *such that the* ∞*-category* $\mathcal{C}_\kappa$ *is essentially small, admits* κ*-small colimits, and the induced functor* $\mathrm{Ind}_\kappa(\mathcal{C}_\kappa) \to \mathcal{C}$ *is an equivalence.*
(4) *There is a small* ∞*-category* $\mathcal{C}'$ *which admits all* κ*-small colimits such that* $\mathcal{C}$ *is equivalent to* $\mathrm{Ind}_\kappa(\mathcal{C}')$*.*
(5) *The* ∞*-category* $\mathcal{C}$ *admits all colimits, and there exists a cardinal* κ *and a set* S *of* κ*-compact objects of* $\mathcal{C}$ *such that every object of* $\mathcal{C}$ *is a colimit of objects in* S*.*

Proof. Since κ-small colimits of κ-compact objects are κ-compact when they exist, it is clear that the (1) implies (2) for every cardinal κ such that $\mathcal{C}$ is κ-accessible. To complete the proof of this step, it suffices to show that under these hypotheses, the κ-accessibility of $\mathcal{C}$ implies κ'-accessibility for all $\kappa' > \kappa$. This follows from the following observation: if $C \in \mathcal{C}$ is the colimit of a κ-filtered diagram $\mathcal{I} \to \mathcal{C}$, then we may replace

$\mathcal{I}$ by a partially ordered set, which is the union of its subsets $\mathcal{I}_\alpha$ having size $< \kappa'$. Then C is the κ'-filtered colimit of objects C_α, where C_α denotes the colimit of the induced diagram $\mathcal{I} \to \mathcal{C}$ (which might not be filtered). Since C_α is a κ'-small colimit of κ-compact objects, it is κ'-compact.

The implications (2) $\Rightarrow$ (3) $\Rightarrow$ (4) $\Rightarrow$ (5) are obvious. We will complete the proof by showing that (5) $\Rightarrow$ (1). Assume that there exists a regular cardinal κ and a set S of κ-compact objects of $\mathcal{C}$ such that every object of $\mathcal{C}$ is a colimit of objects in S. We first claim that $\mathcal{C}_\kappa$ is essentially small: any κ-compact object $C \in \mathcal{C}$ is a colimit of a diagram involving only objects of S, hence a κ-filtered colimit of the colimits of κ-small diagrams in S. The κ-compactness of C implies that C is a retract of a κ-small colimit of objects in S, and thus there are only a bounded number of possibilities for C up to equivalence. To complete the proof, it suffices to show that every object of $\mathcal{C}$ is a κ-filtered colimit of objects of $\mathcal{C}_\kappa$. But this follows immediately, since the colimit of any κ-small diagram in S belongs to $\mathcal{C}_\kappa$. ☐

An ∞-category $\mathcal{C}$ satisfying the equivalent conditions of Proposition 1.4.1 is called *presentable*.

Remark 1.4.2. Let $\mathcal{C}$ be a presentable ∞-category. Then $\mathcal{C}$ is "tensored over $\mathcal{S}$" in the following sense: there exists a functor $\mathcal{C} \times \mathcal{S} \to \mathcal{C}$, which we shall denote by $\otimes$, having the property that

$$\mathrm{Hom}_{\mathcal{C}}(C \otimes S, C') = \mathrm{Hom}_{\mathcal{S}}(S, \mathrm{Hom}_{\mathcal{C}}(C, C'))$$

In order to prove this, we note that the $C \otimes S$ is canonically determined by its universal property and that its formation is compatible with colimits in S. Any space $S \in \mathcal{S}$ may be obtained as a colimit of points (if we regard S as an ∞-groupoid, then S is the colimit of the constant diagram $S \to \mathcal{S}$ having the value $*$). Thus, the existence of $C \otimes S$ can be deduced from the existence when S is a single point, in which case we may take $C \otimes S = C$.

The key fact about presentable ∞-categories which makes them very pleasant to work with is the following representability criterion:

Proposition 1.4.3. *Let $\mathcal{C}$ be a presentable ∞-category, and let $\mathcal{F}$ be a prestack on $\mathcal{C}$. Suppose that $\mathcal{F}$ carries colimits in $\mathcal{C}$ into limits in $\mathcal{S}$. Then $\mathcal{F}$ is representable.*

Proof. Choose a regular cardinal κ such that $\mathcal{C}$ is κ-accessible. Consider the restriction of $\mathcal{F}$ to $\mathcal{C}_\kappa$: we may view this as an ∞-category "fibered in spaces" over $\mathcal{C}_\kappa$. More precisely, we may form a new ∞-category $\widetilde{\mathcal{C}}$ whose objects consist of pairs $C \in \mathcal{C}_\kappa$, $\eta \in \mathcal{F}(C)$, with $\mathrm{Hom}_{\widetilde{\mathcal{C}}}((C, \eta), (C', \eta'))$ given by the fiber of $p : \mathrm{Hom}_{\mathcal{C}}(C, C') \to \mathcal{F}(C)$ over the point η, where p is induced by pulling back $\eta' \in \mathcal{F}(C')$. Note that $\widetilde{\mathcal{C}}$ admits all κ-small colimits, and in particular is κ-filtered.

There exists a functor $\pi : \widetilde{\mathcal{C}} \to \mathcal{C}_\kappa$. Regard π as a diagram in $\mathcal{C}$, and let F_κ denote its colimit. Then, by the compatibility of $\mathcal{F}$ with colimits, we obtain a natural "universal element" $\eta_\kappa \in \mathcal{F}(F_\kappa)$. We would like to show that this exhibits F_κ as representing the presheaf $\mathcal{F}$.

We first prove the F_κ is a good approximation to $\mathcal{F}$ in the following sense: given any morphism $f : X \to Y$ between κ-compact objects of $\mathcal{C}$, the natural map $\mathrm{Hom}_{\mathcal{C}}(Y, F_\kappa) \to \mathcal{F}(Y) \times_{\mathcal{F}(X)} \mathrm{Hom}_{\mathcal{C}}(X, F_\kappa)$ is surjective on π_0. Indeed, suppose we are given a compatible triple (X, η_X), (Y, η_Y), and $p : X \to F_\kappa$, where $\eta_Y \in \mathcal{F}(Y)$, $\eta_X = f^* \eta_Y \in \mathcal{F}(X)$, and we are given a homotopy between $p^*(\eta_\kappa)$ and η_X. Since X is κ-compact, p factors through some pair $(Z, \eta_Z) \in \widetilde{\mathcal{C}}$. Let $W = Y \coprod_X Z$. Since $\mathcal{F}$ is compatible with pushouts, the homotopy between $\eta_Z | X$ and $\eta_Y | X$ gives us $\eta_W \in \mathcal{F}(W)$ with $\eta_W | Z = \eta_Z$ and $\eta_W | Y = \eta_Y$. Then $(W, \eta_W) \in \widetilde{\mathcal{C}}$, and the natural map $Y \to W$ gives rise to a point of $\mathrm{Hom}_{\mathcal{C}}(Y, F_\kappa)$. One readily checks that this point has the desired properties.

Now let us show that F_κ represents $\mathcal{F}$. In other words, we want to show that the natural map

$$\mathrm{Hom}_{\mathcal{C}}(Y, F_\kappa) \to \mathcal{F}(Y)$$

is a homotopy equivalence for every Y. Since both sides are compatible with colimits in the variable Y, we may assume that Y is κ-compact. By Whitehead's theorem, it will suffice to show that the natural map

$$\mathrm{Hom}_{\mathcal{C}}(Y, F_\kappa) \to \mathcal{F}(Y) \times_{[S^n, \mathcal{F}(Y)]} [S^n, \mathrm{Hom}_{\mathcal{C}}(Y, F_\kappa)]$$

is surjective on π_0, where $[S^n, Z]$ denotes the "free loop space" consisting of maps from an n-sphere S^n into Z. This reduces to a special case of what we proved above if we take $X = Y \otimes S^n$. □

The representability criterion of Proposition 1.4.3 has many consequences, as we now demonstrate.

Corollary 1.4.4. *A presentable ∞-category admits arbitrary limits.*

Proof. Any limit of representable prestacks carries colimits into limits, and is therefore representable by Proposition 1.4.3. □

Remark 1.4.5. Now that we know that $\mathcal{C}$ has arbitrary limits, we can apply an argument dual to that of Remark 1.4.2 to show that $\mathcal{C}$ is *cotensored over* $\mathcal{S}$. In other words, for any $C \in \mathcal{C}$ and any $X \in \mathfrak{S}$, there exists an object $C^X \in \mathcal{C}$ and a natural equivalence $\mathrm{Hom}_{\mathcal{C}}(\bullet, C^X) = \mathrm{Hom}_{\mathfrak{S}}(X, \mathrm{Hom}_{\mathcal{C}}(\bullet, C)$.

From Proposition 1.4.3 we also get a version of the adjoint functor theorem:

Corollary 1.4.6. *Let $F : \mathcal{C} \to \mathcal{C}'$ be a functor between presentable ∞-categories. Then F has a right adjoint if and only if F preserves all colimits.*

Proof. Suppose we have a small diagram $\{X_i\}$ in $\mathcal{C}$ with colimit X. If F has a right adjoint G, then $\mathrm{Hom}_{\mathcal{C}'}(FX, Y) = \mathrm{Hom}_{\mathcal{C}}(X, GY) = \lim_i \mathrm{Hom}_{\mathcal{C}}(X_i, GY) = \lim_i \mathrm{Hom}_{\mathcal{C}'}(FX_i, Y)$, so that FX is a colimit of the induced diagram $\{FX_i\}$ in $\mathcal{C}'$. Conversely, if F preserves all colimits, then for each $Y \in \mathcal{C}'$, we may define a prestack $\mathcal{F}_Y$ on $\mathcal{C}$ by the equation $\mathcal{F}_Y(X) = \mathrm{Hom}_{\mathcal{C}'}(FX, Y)$. The hypothesis on F implies that $\mathcal{F}_Y$ carries colimits into limits, so $\mathcal{F}_Y$ is representable by an object $GY \in \mathcal{C}$. The universal property enjoyed by GY automatically ensures that G is a functor $\mathcal{C}' \to \mathcal{C}$, left adjoint to F. □

If $\mathcal{C}$ and $\mathcal{C}'$ are presentable, we let $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{C}, \mathcal{C}')$ denote the ∞-category of colimit-preserving functors from $\mathcal{C}$ to $\mathcal{C}'$. These give the appropriate notion of "morphism" between presentable ∞-categories.

The collection of colimit functors between presentable ∞-categories is stable under compositions. We may summarize the situation informally by saying that there is an $(\infty, 2)$-category $\mathrm{Fun}^{\mathrm{pres}}$ of presentable ∞-categories, with colimit-preserving functors as morphisms. This $(\infty, 2)$-category admits arbitrary $(\infty, 2)$-categorical limits, which are just $(\infty, 2)$-categorical limits of the underlying ∞-categories (such limits are accessible by Proposition 1.3.19, and colimits may be computed levelwise). This $(\infty, 2)$-category also admits arbitrary $(\infty, 2)$-categorical colimits, which may be constructed using generators and relations (see [1] for a discussion in the 1-categorical case), but we shall not need this.

Remark 1.4.7. If $\mathcal{C}$ is a presentable ∞-category, then any slice $\mathcal{C}_{/X}$ is also presentable. Indeed, $\mathcal{C}_{/X}$ is accessible (see Example 1.3.20) and has all colimits (which are also colimits in $\mathcal{C}$).

Remark 1.4.8. Let $\mathcal{C}$ be a small ∞-category, and let $\mathcal{P} = \mathcal{S}^{\mathcal{C}^{op}}$ denote the ∞-category of stacks on $\mathcal{C}$. Then $\mathcal{P}$ is presentable: in fact, it is a limit of copies of $\mathcal{S}$ indexed by $\mathcal{C}^{op}$ (one can also argue directly). But $\mathcal{P}$ also has a dual description: it is a *colimit* of copies of $\mathcal{S}$ indexed by $\mathcal{C}$. In other words, $\mathcal{P}$ is the *free presentable ∞-category generated by* $\mathcal{C}$, in the sense that for any ∞-pretopos $\mathcal{C}'$ we have an equivalence of ∞-categories

$$\pi : \mathrm{Fun}^{\mathrm{pres}}(\mathcal{P}, \mathcal{C}') \to \mathrm{Fun}(\mathcal{C}, \mathcal{C}')$$

where the left hand side denotes the ∞-category of right-exact continuous functors from $\mathcal{P}$ to $\mathcal{C}'$. The functor π is given by restriction, and it has a homotopy inverse which carries a functor $F : \mathcal{C} \to \mathcal{C}'$ to $\widetilde{F} : \mathcal{P} \to \mathcal{C}'$, where $\widetilde{F}(\mathcal{F})$ is the homotopy colimit of the diagram $F \circ p : \mathcal{I} \to \mathcal{C}'$, where $\mathcal{I}$ is the ∞-category fibered over $\mathcal{C}$ in groupoids defined by $\mathcal{F}$ and $p : \mathcal{I} \to \mathcal{C}$ is the canonical projection.

We conclude by showing that the $(\infty, 2)$-category of presentable ∞-categories admits an internal Hom-functor. The proof uses material from the next section.

Proposition 1.4.9. *Let $\mathcal{C}$ and $\mathcal{C}'$ be presentable ∞-categories. The ∞-category $\mathrm{Fun}^{pres}(\mathcal{C}, \mathcal{C}')$ of colimit-preserving functors from $\mathcal{C}$ to $\mathcal{C}'$ is presentable.*

Proof. We use Proposition 1.5.3, to be proved in the next section, which asserts that $\mathcal{C}$ is a localization of $\mathcal{S}^{\mathcal{C}_0^{op}}$ for some small ∞-category $\mathcal{C}_0$. This immediately implies that $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{C},\mathcal{C}')$ is a localization of $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{S}^{\mathcal{C}_0^{op}},\mathcal{C}')$. Using Proposition 1.5.3 again, it suffices to consider the case where $\mathcal{C} = \mathcal{S}^{\mathcal{C}_0^{op}}$. In this case, we deduce that $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{C},\mathcal{C}') = \mathrm{Fun}(\mathcal{C}_0,\mathcal{C}')$, which is a limit (indexed by $\mathcal{C}_0$) of ∞-categories having the form $\mathrm{Hom}(*,\mathcal{C}') = \mathcal{C}'$. Since $\mathcal{C}'$ is an presentable, the limit is also presentable. □

1.5. **Localization.** In this section, we discuss *localizations* of a presentable ∞-category $\mathcal{C}$. These ideas are due to Bousfield (see for example [11]). Suppose we have a collection S of morphisms in $\mathcal{C}$ which we would like to invert. In other words, we seek an presentable ∞-category $S^{-1}\mathcal{C}$ equipped with a colimit-preserving functor $\pi : \mathcal{C} \to S^{-1}\mathcal{C}$ which is "universal" with respect to the fact that every morphism in S becomes an equivalence in $S^{-1}\mathcal{C}$. It is an observation of Bousfield that we can often find $S^{-1}\mathcal{C}$ *inside of* $\mathcal{C}$, as the set of "local" objects of $\mathcal{C}$.

Example 1.5.1. Let $\mathcal{C}$ be the (ordinary) category of abelian groups, p a prime number, and let S denote the collection of morphisms f whose kernel and cokernel consist entirely of p-power torsion. A morphism f lies in S if and only if it induces an isomorphism after inverting the prime number p. In this case, we may identify $S^{-1}\mathcal{C}$ with the category of abelian groups which are *uniquely p-divisible*.

In Example 1.5.1, the natural functor $\mathcal{C} \to S^{-1}\mathcal{C}$ is actually left adjoint to an inclusion functor. Let us therefore begin by examining functors which are left adjoint to inclusions.

If $\mathcal{C}_0 \subseteq \mathcal{C}$ is a full subcategory of an ∞-category, and L is a left adjoint to the inclusion $\mathcal{C}_0 \to \mathcal{C}$, then we may also regard L as a functor from $\mathcal{C}$ to itself. The adjunction gives us a natural transformation $\alpha : \mathrm{id}_{\mathcal{C}} \to L$. From the pair (L,α) we can recover $\mathcal{C}_0$ as the full subcategory of $\mathcal{C}$ consisting of those objects $C \in \mathcal{C}$ such that $\alpha(C) : C \to LC$ is an equivalence. Conversely, if we begin with a functor $L : \mathcal{C} \to \mathcal{C}$ and a natural transformation $\alpha : \mathrm{id}_{\mathcal{C}} \to L$, and *define* $\mathcal{C}_0$ as above, then L is left adjoint to the inclusion $\mathcal{C}_0 \to \mathcal{C}$ if and only if (L,α) is *idempotent* in the following sense:

- For any $C \in \mathcal{C}$, the morphisms
$$L(\alpha(C)), \alpha(LC) : LC \to LLC$$
are equivalences which are homotopic to one another.

In this situation, we shall say that L is an *idempotent* functor $\mathcal{C} \to \mathcal{C}$ (the transformation α being understood).

Proposition 1.5.2. *Suppose that $\mathcal{C}$ is an accessible category and that $L : \mathcal{C} \to \mathcal{C}$ is idempotent. Then the induced functor $L_0 : \mathcal{C} \to L\mathcal{C}$ is accessible if and only if $L\mathcal{C}$ is an accessible category.*

Proof. If $L\mathcal{C}$ is accessible, then L_0 can be described as the left adjoint of the inclusion $L\mathcal{C} \subseteq \mathcal{C}$ which is accessible by Proposition 1.3.21. For the converse, suppose that L_0 is κ-continuous for some regular cardinal κ. Without loss of generality we may enlarge κ so that $\mathcal{C}$ is κ-accessible. It follows immediately that $L\mathcal{C}$ is stable under the formation of κ-filtered colimits. Since L_0 is a left adjoint, it preserves all colimits. Every object of $\mathcal{C}$ is a κ-filtered colimit of κ-compact objects of $\mathcal{C}$; hence every object of $L\mathcal{C}$ is a κ-filtered colimit of objects of the form LC, $C \in \mathcal{C}_\kappa$. To complete the proof, it suffices to note that every such LC is κ-compact in $L\mathcal{C}$. □

An accessible idempotent functor in an accessible category will be referred to as a *localization functor*. We will say that an ∞-category $\mathcal{C}_0$ is a *localization* of $\mathcal{C}$ if $\mathcal{C}_0$ is equivalent to the essential image of some localization functor $L : \mathcal{C} \to \mathcal{C}$.

We now give Simpson's characterization of presentable ∞-categories. The proof can also be found in [4], but we include it here for completeness.

Proposition 1.5.3. *Let $\mathcal{C}$ be an ∞-category. The following conditions are equivalent:*

(1) *There exists a regular cardinal κ such that the ∞-category $\mathcal{C}$ is κ-accessible and the Yoneda embedding $\mathcal{C} \to \mathcal{P} = \mathcal{S}^{\mathcal{C}_\kappa^{op}}$ has a left adjoint.*
(2) *The ∞-category $\mathcal{C}$ is a localization of an ∞-category of prestacks.*

(3) *The ∞-category $\mathcal{C}$ is a localization of a presentable ∞-category.*
(4) *The ∞-category $\mathcal{C}$ is presentable.*

Proof. If $\mathcal{C}$ satisfies (1), then we may regard it as a full subcategory of $\mathcal{P}$ via the Yoneda embedding. The left adjoint to the Yoneda embedding is then viewed as an localization functor $\mathcal{P} \to \mathcal{P}$ having $\mathcal{C}$ as its essential image, which proves (2).

It is obvious that (2) implies (3). To see that (3) implies (4), note that if $\mathcal{C}$ is the essential image of some localization functor defined on a larger presentable ∞-category $\mathcal{C}'$, then $\mathcal{C}$ is accessible by Proposition 1.3.21 and arbitrary colimits may be formed in $\mathcal{C}$ by forming those colimits in $\mathcal{C}'$ and then applying the functor L.

Finally, suppose that $\mathcal{C}$ is presentable, and choose a regular cardinal κ such that $\mathcal{C}$ is κ-accessible. Let $\mathcal{P} = \mathcal{S}^{\mathcal{C}_\kappa^{op}}$ denote the ∞-category of prestacks on $\mathcal{C}_\kappa$. The functor $\mathcal{C}_\kappa \to \mathcal{C}$ extends uniquely to a right-exact continuous functor $\mathcal{P} \to \mathcal{C}$. One can easily check that this functor is left adjoint to the Yoneda embedding, which proves (1). □

Remark 1.5.4. Carlos Simpson has also shown that the theory of presentable ∞-categories is equivalent to that of *cofibrantly generated model categories* (see [4]). Since most of the ∞-categories we shall meet are presentable, the subject could also be phrased in the language of model categories. But we shall not do this, since the restriction to presentable ∞-categories seems unnatural and is often technically inconvenient.

We now proceed to discuss localizations of presentable ∞-categories in more detail.

Definition 1.5.5. Let $\mathcal{C}$ be an ∞-category and S a collection of morphisms of $\mathcal{C}$. An object $C \in \mathcal{C}$ is *S-local* if, for any morphism $s : X \to Y$ in S, the induced morphism $\mathrm{Hom}_{\mathcal{C}}(Y, C) \to \mathrm{Hom}_{\mathcal{C}}(X, C)$ is an equivalence.

A morphism $f : X \to Y$ in $\mathcal{C}$ is an *S-equivalence* if the natural map $\mathrm{Hom}_{\mathcal{C}}(Y, C) \to \mathrm{Hom}_{\mathcal{C}}(X, C)$ is a homotopy equivalence whenever C is S-local.

Let $\to$ denote the (ordinary) category consisting of two objects and a single morphism between them. For any ∞-category $\mathcal{C}$, we may form a functor ∞-category $\mathcal{C}^{\to}$ whose objects are the morphisms of $\mathcal{C}$. We note that $\mathcal{C}^{\to}$ is presentable whenever $\mathcal{C}$ is presentable.

Definition 1.5.6. Let $\mathcal{C}$ be a presentable ∞-category. A collection S of morphisms of $\mathcal{C}$ will be called *saturated* if it satisfies the following conditions:

- Every equivalence belongs to S.
- If f and g are homotopic, then $f \in S$ if and only if $g \in S$.
- Suppose that $f : X \to Y$ and $g : Y \to Z$ are a composable pair of morphisms in $\mathcal{C}$. If any two of f, g, and $g \circ f$ belong to S, then so does the third.
- If $f : X \to Y$ and $g : X \to Z$ are morphisms of S and f belongs to S, then the induced morphism $f' : Z \to Z \coprod_X Y$ belongs to S.
- The full subcategory of $\mathcal{C}^{\to}$ consisting of morphisms which lie in S is stable under the formation of colimits.

Remark 1.5.7. Let S be a collection of morphisms in $\mathcal{C}$. It is clear that every morphism in S is an S-equivalence. If $\mathcal{C}$ is presentable, then the collection of S-equivalences is saturated.

Remark 1.5.8. If S is saturated, and $f : X \to Y$ lies in S, then the induced map $X \otimes S^n \to Y \otimes S^n$ lies in S.

Remark 1.5.9. If $\mathcal{C}$ is any presentable ∞-category, then the collection of all equivalences in $\mathcal{C}$ is saturated.

Remark 1.5.10. If $F : \mathcal{C} \to \mathcal{C}'$ is a colimit preserving functor and S' is a saturated collection of morphisms of $\mathcal{C}'$, then $S = \{f \in \mathrm{Hom}_{\mathcal{C}}(X, Y) : Ff \in S'\}$ is saturated. In particular, the collection of morphisms f such that Ff is an equivalence is saturated.

Remark 1.5.11. Let $\mathcal{C}$ be a presentable ∞-category. Then any intersection of saturated collections of morphisms in $\mathcal{C}$ is saturated. Consequently, for any collection S of morphisms of $\mathcal{C}$ there is a smallest saturated collection $\overline{S}$ containing S, which we call the *saturation* of S. It follows immediately that every morphism in $\overline{S}$ is an S-equivalence.

We shall say that a saturated collection S of morphisms is *setwise generated* if $S = \overline{S_0}$ for some *set* $S_0 \subseteq S$.

Proposition 1.5.12. *Let $\mathcal{C}$ be a presentable ∞-category, and let S be a set of morphisms in $\mathcal{C}$. Let $\mathcal{C}_S$ denote the full subcategory of $\mathcal{C}$ consisting of S-local objects. Then the inclusion $\mathcal{C}_S \to \mathcal{C}$ has an accessible left adjoint L. Furthermore, the following are equivalent for a morphism $f : X \to Y$ in $\mathcal{C}$:*

1. *The morphism f lies in the saturation $\overline{S}$.*
2. *The morphism f is an S-equivalence.*
3. *The morphism Lf is an equivalence.*

Proof. We may enlarge S by adjoining elements of $\overline{S}$ without changing the collection of S-local objects. Thus, we may assume that for any morphism $s : X \to Y$ in S, the morphism $X \coprod_{X \otimes S^n} (Y \otimes S^n) \to Y$ also belongs to S. This enlargement allows us to simplify the definition of an S-local object: by Whitehead's theorem, an object C is S-local if and only if

$$\mathrm{Hom}_{\mathcal{C}}(Y, C) \to \mathrm{Hom}_{\mathcal{C}}(X, C)$$

is surjective on π_0 for every morphism $s : X \to Y$ in S.

Choose a regular cardinal κ such that $\mathcal{C}$ is κ-accessible, and such that the source and target of every morphism in S is κ-compact. We will show that for every object $C \in \mathcal{C}_\kappa$, there exists a morphism $f : C \to C_S$ such that $f \in \overline{S}$ and C_S is S-local. This will prove the existence of the left adjoint L (carrying C to C_S), at least on the subcategory $\mathcal{C}_\kappa \subseteq \mathcal{C}$. This functor L has a unique κ-continuous extension to $\mathcal{C}$. Since any κ-filtered colimit of S-local objects is S-local, this κ-continuous extension will be a left adjoint defined on all of $\mathcal{C}$. The closure of $\overline{S}$ under κ-filtered colimits will also show that the adjunction morphism $C \to LC$ lies in $\overline{S}$, for any $C \in \mathcal{C}$.

Let us therefore assume that $C \in \mathcal{C}_\kappa$. Let $\mathcal{J}$ denote the ∞-category whose objects are morphisms $C \to C'$ in $\mathcal{C}_\kappa$, which lie in $\overline{S}$. One readily verifies (using the fact that $\overline{S}$ is saturated) that $\mathcal{J}$ is a κ-filtered category, equipped with a functor $\pi : \mathcal{J} \to \mathcal{C}_\kappa$ carrying the diagram $C \to C'$ to the object $C' \in \mathcal{C}_\kappa$. We may therefore regard π as a κ-filtered diagram in $\mathcal{C}_\kappa$, which has a colimit C_S in $\mathcal{C}$. Since π is a diagram of morphisms in $\overline{S}$, the natural map $C \to C_S$ lies in $\overline{S}$. We must now show that C_S is S-local.

Suppose $s : X \to Y$ is a morphism in S. We need to show that $\pi_0 \mathrm{Hom}_{\mathcal{C}}(Y, C_S) \to \pi_0 \mathrm{Hom}_{\mathcal{C}}(X, C_S)$ is surjective. To this end, choose any morphism $f : X \to C_S$. Since X is κ-compact, the morphism f factors through some morphism $X \to C'$, where $C \to C'$ appears in the κ-filtered diagram π. Then the pushout morphism $C \to C' \coprod_X Y$ lies in $\overline{S}$, and so also occurs in the diagram π. Thus we obtain a map $f' : Y \to (C' \coprod_X Y) \to C_S$. It is clear from the construction that f factors through f'.

It remains to verify the equivalence of the three conditions listed above. We have already seen that (1) $\implies$ (2), and the equivalence of (2) and (3) is formal. To complete the proof, it suffices to show that (3) implies (1). If $f : X \to Y$ induces an equivalence after applying L, then the morphisms $X \to Y_S$ and $Y \to Y_S$ both lie in $\overline{S}$, so that f lies in $\overline{S}$ since $\overline{S}$ is saturated. □

We note a universal property enjoyed by the localization $\mathcal{C}_S$:

Proposition 1.5.13. *Let $\mathcal{C}$ and $\mathcal{C}'$ be presentable ∞-categories, let S be a set of morphisms of $\mathcal{C}$, and let $L : \mathcal{C} \to \mathcal{C}_S$ be left adjoint to the inclusion. Composition with L induces a fully faithful functor $\pi : \mathrm{Fun}^{pres}(\mathcal{C}_S, \mathcal{C}') \to \mathrm{Fun}^{pres}(\mathcal{C}, \mathcal{C}')$ whose essential image consists of those functors $F : \mathcal{C} \to \mathcal{C}'$ for which Fs is an equivalence for all $s \in S$.*

Proof. Restriction induces a functor $\psi : \mathrm{Fun}^{\mathrm{pres}}(\mathcal{C}, \mathcal{C}') \to \mathrm{Fun}^{\mathrm{pres}}(\mathcal{C}_S, \mathcal{C}')$ which is right adjoint to π. It is clear that the adjunction $\psi \circ \pi \to 1$ is an equivalence, which shows that π is fully faithful. By construction, the essential image consists of all functors for which F induces equivalences $FC \simeq FLC$ for all $C \in \mathcal{C}$. Since every morphism $C \to LC$ lies in $\overline{S}$, it will suffice to show that if F inverts every morphism in S, then it inverts every morphism in $\overline{S}$. This follows immediately from the observation that the set of morphisms which become invertible after application of F is saturated. □

We conclude this section by showing that *every* localization of an presentable ∞-category $\mathcal{C}$ has the form $\mathcal{C}_S$, where S is some set of morphisms of $\mathcal{C}$.

We begin with the following observation: let $F : \mathcal{C} \to \mathcal{C}'$ be an accessible functor between accessible ∞-categories, and let S denote the full subcategory of $\mathcal{C}^{\to}$ consisting of morphisms s such that Fs is an equivalence. Then S may be constructed as a limit of accessible ∞-categories, so that S is accessible: in particular, there exists a *set* of objects $S_0 \subseteq S$ such that every object in S is a filtered colimit of objects in S_0. It follows that each element of S is an S_0-equivalence, so that if $\mathcal{C}$ is presentable then $S \subseteq \overline{S_0}$. If in addition F is a colimit-preserving functor between presentable ∞-categories, then S is saturated so we obtain $S = \overline{S_0}$. Consequently, we obtain the following:

Theorem 1.5.14. *Let $\mathcal{C}$ be a presentable ∞-category. Let S be a collection of morphisms in $\mathcal{C}$. The following conditions are equivalent:*

1. *The collection S is the saturation of some* set *of morphisms in $\mathcal{C}$.*
2. *There exists a colimit preserving functor $F : \mathcal{C} \to \mathcal{C}'$ such that S consists of all morphisms s for which Fs is an equivalence.*
3. *There exists a localization functor $L : \mathcal{C} \to \mathcal{C}$ such that S consists of all morphisms s for which Ls is an equivalence.*

Proof. The above discussion shows that (2) implies (1), and it is clear that (3) implies (2). Proposition 1.5.12 shows that (1) implies (3). ☐

It follows that *any* localization functor defined on a presentable ∞-category $\mathcal{C}$ may be obtained as a localization $\mathcal{C} \to \mathcal{C}_S$ for some *set* of morphisms S of $\mathcal{C}$.

2. ∞-Topoi

In this section, we will define the notion of an ∞-topos. These were defined by Simpson in [4] as ∞-categories which occur as left-exact localizations of ∞-categories of prestacks. In §2.2, we give our own definition, which we will show to be equivalent to Simpson's definition in §2.4. We will then explain how to extract topoi from ∞-topoi and vice versa. Finally, in §2.9, we explain how to obtain the Joyal-Jardine ∞-topos of hyperstacks on topos X as a localization of our ∞-topos of stacks on X.

There are several papers on higher topoi in the literature. The papers [15] and [16] both discuss a notion of 2-topos (the second from an elementary point of view). However, the basic model for these 2-topoi was the 2-category of (small) categories, rather than the 2-category of (small) groupoids. Jardine ([14]) has exhibited a model structure on the simplicial presheaves on a Grothendieck site, and the ∞-category associated to this model category is an ∞-topos in our sense. This construction was generalized from ordinary categories with a Grothendieck topology to simplicial categories with a Grothendieck topology (appropriately defined) in [17]; this also produces ∞-topoi. However, not every ∞-topos arises in this way: one can construct only ∞-topoi which are t-complete in the sense of [17]; we will summarize the situation in Section 2.9. We should remark here that our notion of ∞-topos is equivalent to the notion of a *Segal topos* defined in [17].

In our terminology, "topos" shall always mean "Grothendieck topos": we will discuss an "elementary" version of these ideas in a sequel to this paper.

2.1. Equivalence Relations and Groupoid Objects.

There are several equivalent ways to define a (Grothendieck) topos. The following is proved in [5]:

Proposition 2.1.1. *Let $\mathcal{C}$ be a category. The following are equivalent:*

1. *The category $\mathcal{C}$ is (equivalent to) the category of sheaves on some Grothendieck site.*
2. *The category $\mathcal{C}$ is (equivalent to) a left-exact localization of the category of presheaves on some small category.*
3. *The category $\mathcal{C}$ satisfies Giraud's axioms:*
   - *$\mathcal{C}$ has a set of small generators.*
   - *$\mathcal{C}$ admits arbitrary colimits.*
   - *Colimits in $\mathcal{C}$ are universal (that is, the formation of colimits commutes with base change).*
   - *Sums in $\mathcal{C}$ are disjoint.*
   - *All equivalence relations in $\mathcal{C}$ are effective.*

We will prove an analogue of this result (Theorem 2.4.1) in the ∞-categorical context. However, our result is quite so satisfying: it establishes the equivalence of analogues of conditions (2) and (3), but does not give an explicit method for constructing all of the left-exact localizations analogous to (1).

In the next section, we will define an ∞-topos to be an ∞-category which satisfies ∞-categorical versions of Giraud's axioms. The first four of these axioms generalize in a very straightforward way to ∞-categories, but the last (and most interesting axiom) is more subtle.

Recall that if X is an object in an (ordinary) category $\mathcal{C}$, then an *equivalence relation* R on X is an object of $\mathcal{C}$ equipped with a map $p : R \to X \times X$ such that for any S, the induced map $\mathrm{Hom}_{\mathcal{C}}(S, R) \to \mathrm{Hom}_{\mathcal{C}}(S, X) \times \mathrm{Hom}_{\mathcal{C}}(S, X)$ exhibits $\mathrm{Hom}_{\mathcal{C}}(S, R)$ as an equivalence relation on $\mathrm{Hom}_{\mathcal{C}}(S, X)$.

If $\mathcal{C}$ admits finite limits, then it is easy to construct equivalence relations in $\mathcal{C}$: given any map $X \to Y$ in $\mathcal{C}$, the induced map $X \times_Y X \to X \times X$ is an equivalence relation on X. If the category $\mathcal{C}$ admits finite colimits, then one can attempt to invert this process: given an equivalence relation R on X, one can form the coequalizer of the two projections $R \to X$ to obtain an object which we will denote by X/R. In the category of sets, one can recover R as the fiber product $X \times_{X/R} X$. In general, this need not occur: one always has $R \subseteq X \times_{X/R} X$, but the inclusion may be strict (as subobjects of $X \times X$). If equality holds, then R is said to be an *effective equivalence relation*, and the map $X \to X/R$ is said to be an *effective epimorphism*. In this terminology, we have the following:

Fact 2.1.2. *In the category $\mathcal{C}$ of sets, every equivalence relation is effective and the effective epimorphisms are precisely the surjective maps.*

The first assertion of Fact 2.1.2 remains valid in any topos, and according to the axiomatic point of view it is one of the defining features of topoi.

If $\mathcal{C}$ is a category with finite limits and colimits in which all equivalence relations are effective, then we obtain a one-to-one correspondence between equivalence relations on an object X and "quotients" of X (that is, isomorphism classes of effective epimorphisms $X \to Y$). This correspondence is extremely useful because it allows us to make elementary descent arguments: one can deduce statements about quotients of X from statements about X and about equivalence relations on X (which live over X). We would like to have a similar correspondence in certain ∞-categories.

In order to formulate the right notions, let us begin by considering the ∞-category $\mathcal{S}$ of spaces. The correct notion of surjection of spaces $X \to Y$ is a map which induces a surjection on components $\pi_0 X \to \pi_0 Y$. However, in this case, the fiber product $R = X \times_Y X$ is does not give an equivalence relation on X, because the map $R \to X \times X$ is not necessarily "injective" in any reasonable sense. However, it does retain some of the pleasant features of an equivalence relation: instead of transitivity, we have a *coherently associative* composition map $R \times_X R \to R$. In order to formalize this observation, it is convenient to introduce some simplicial terminology.

Let Δ denote the (ordinary) category of finite, nonempty, linearly ordered sets. If $\mathcal{C}$ is an ∞-category, then a *simplicial object* of $\mathcal{C}$ is a functor $C_\bullet : \Delta^{op} \to \mathcal{C}$. Since every object of Δ is isomorphic to $\{0, \ldots, n\}$ for some $n \geq 0$, we may think of $C_\bullet$ as being given by objects $\{C_n\}_{n \geq 0}$ in $\mathcal{C}$, together with various maps between the C_n which are compatible up to coherent homotopy. In particular, we have two maps $\pi_0, \pi_1 : C_1 \to C_0$, and we have natural maps $p_n : C_n \to C_1 \times_{C_0} \ldots \times_{C_0} C_1$, where there are n factors in the product and for each fiber product, the map from the left copy of C_1 to C_0 is given by π_1 and the map from the right copy of C_1 to C_0 is given by π_0. If p_n is an equivalence for each $n \geq 0$, then we shall say that C_n is a *category object of* $\mathcal{C}$. (Note that this definition does not really require that $\mathcal{C}$ has fiber products: a priori, the target of p_n may be viewed as a prestack on $\mathcal{C}$.)

Example 2.1.3. Suppose that $\mathcal{C}$ is the category of sets. Then a category object $C_\bullet$ of $\mathcal{C}$ consists of a set C_0 of *objects*, a set C_1 of *morphisms*, together with source and target maps $\pi_0, \pi_1 : C_1 \to C_0$ and various other maps between iterated fiber powers of C_1 over C_0. In particular, one has a map $C_1 \times_{C_0} C_1 \to C_1$ which gives rise to an multiplication on "composable" pairs of morphisms. One can show that this recovers the usual definition of a category.

Example 2.1.4. More generally, suppose that $\mathcal{C}$ is an ordinary category. Then what we have called a category object of $\mathcal{C}$ is equivalent to the usual notion of a category object of $\mathcal{C}$: namely, data which represents a "category-valued functor" on $\mathcal{C}$.

In the general case, the requirement that p_n be an equivalence shows that C_n is determined by C_1 and C_0. Thus, a category object in $\mathcal{C}$ may be viewed as a pair of morphisms $\pi_0, \pi_1 : C_1 \to C_0$, together with some additional data which formalizes the idea that there should be a coherently associative composition law on composable elements of C_1.

If an ∞-category $\mathcal{C}$ has fiber products, then any morphism $U \to X$ gives rise to a category object $U_\bullet$ of $\mathcal{C}$, with U_n given by the $(n+1)$-fold fiber power of U over X. However, this category object has a further important property: it is a *groupoid object*. This can be formalized in many ways: for example, one can assert the existence of an "inverse map" $i : U_1 \to U_1$ such that $\pi_0 \circ i \simeq \pi_1$, $\pi_1 \circ i \simeq \pi_0$, and various compositions involving the inverse map are homotopic to the identity. It is important to note that the inverse map is canonically determined by the category structure of $\{U_\bullet\}$ *if it exists*, and therefore we will not need to axiomatize its properties. In other words, a groupoid object in $\mathcal{C}$ is a category object of $\mathcal{C}$ which *satisfies conditions*, rather than a category object with *extra structure*.

If an ∞-category $\mathcal{C}$ admits all limits and colimits, then we may view any simplicial object $C_\bullet : \Delta^{op} \to \mathcal{C}$ as a diagram in $\mathcal{C}$ and form its colimit $|C_\bullet|$, which is also called the *geometric realization* of $C_\bullet$. There is a natural map $C_0 \to |C_\bullet|$, which induces a map of simplicial objects $\psi : C_\bullet \to U_\bullet$, where U_n denotes the $(n+1)$-fold fiber power of C_0 over $|C_\bullet|$. If the map ψ is an equivalence, then we shall say that $C_\bullet$ is an *effective groupoid*, and that $C_0 \to |C_\bullet|$ is an *effective epimorphism*.

Remark 2.1.5. Since $C_0 \simeq U_0$ and the objects C_n (U_n) are determined by taking fiber products of C_1 over C_0 (U_1 over U_0), we see that ψ is an equivalence if and only if it induces an equivalence $C_1 \simeq U_1$. In other words, $C_\bullet$ is effective if and only if $C_1 \simeq C_0 \times_{|C_\bullet|} C_0$.

The notion of a groupoid object in $\mathcal{C}$ will be our ∞-categorical replacement for the 1-categorical notion of an object with an equivalence relation. This new notion is useful thanks to the following analogue of Fact 2.1.2:

Proposition 2.1.6. *In the ∞-category $\mathcal{S}$ of spaces, all groupoid objects are effective and a map $X \to Y$ is an effective epimorphism if and only if $\pi_0 X \to \pi_0 Y$ is surjective.*

In contrast to Fact 2.1.2, Proposition 2.1.6 is nontrivial. For example, a category object $U_\bullet$ in $\mathcal{S}$ with $U_0 = *$ is more-or-less the same thing as an A_∞-space in the terminology of [13] (in other words, the space U_1 has a coherently associative multiplication operation). The first part of Proposition 2.1.6 asserts, in this case, that an A_∞-space whose set of connected components forms a group can be realized as a loop space. We refer the reader to our earlier discussion in Remark 1.2.7.

Remark 2.1.7. In some sense, the notion of a groupoid object is *simpler than* the notion of an equivalence relation. For example, let $\mathcal{C}$ be the category of sets. Then groupoid objects are simply groupoids in the usual sense, while equivalence relations correspond to groupoids satisfying an extra discreteness condition (no objects can have nontrivial automorphisms). The removal of this discreteness condition is what distinguishes the theory of ∞-topoi from the theory of ordinary topoi. It permits us to form quotients of objects using more general kinds of "gluing data". In geometric contexts, this extra flexibility allows the construction of useful objects such as orbifolds and algebraic stacks, which are useful in a variety of mathematical situations.

One can imagine weakening the gluing conditions even further, and considering axioms having the form "every category object is effective". This seems to be a very natural approach to a theory of topos-like (∞, ∞)-categories. We will expound further on this idea in a sequel to this paper.

2.2. **The Definition of an ∞-Topos.**

Just as a topos is a category which resembles the category of sets, we shall define an ∞-topos to be an ∞-category which resembles the ∞-category $\mathcal{S}$ of spaces. We will begin by extracting out certain properties enjoyed by $\mathcal{S}$. We will then take these properties as the axioms for an ∞-topos. They are the analogues of Giraud's conditions which characterize topoi in the case of ordinary categories.

Fact 2.2.1. *The ∞-category $\mathcal{S}$ is presentable.*

In particular, $\mathcal{S}$ admits all limits and colimits, so that fiber products exist in $\mathcal{S}$. However, fiber products have special properties in the ∞-category $\mathcal{S}$:

Fact 2.2.2. *In $\mathcal{S}$, the formation of colimits commutes with pullback.*

It follows from Fact 2.2.2 that the initial object $\emptyset$ of $\mathcal{S}$ really behaves as if it were "empty". Since $\emptyset$ is the colimit of the empty diagram, we see that for any map $X \to \emptyset$, X is the colimit of the pullback of the empty diagram (which is also empty). Thus $X \simeq \emptyset$.

Fact 2.2.3. *Sums in $\mathcal{S}$ are disjoint. That is, given any two spaces X and Y, the fiber product $X \times_{X \coprod Y} Y$ is an initial object.*

Finally, we recall Proposition 2.1.6, which implies in particular:

Fact 2.2.4. *Any groupoid object in $\mathcal{S}$ is effective.*

We are now in a position to define the notion of an ∞-topos.

Definition 2.2.5. An ∞-topos is an ∞-category $\mathcal{X}$ with the following properties:

- $\mathcal{X}$ is presentable.
- The formation of colimits in $\mathcal{X}$ commutes with pullback.
- Sums in $\mathcal{X}$ are disjoint.
- Every $\mathcal{X}$-groupoid is effective.

Then it follows from the discussion up to this point that:

Proposition 2.2.6. *The ∞-category $\mathcal{S}$ is an ∞-topos.*

We also note that any slice of an ∞-topos is an ∞-topos.

Proposition 2.2.7. *Let $\mathcal{X}$ be an ∞-topos, and let $E \in \mathcal{X}$ be an object. Then the slice ∞-category $\mathcal{X}_{/E}$ is an ∞-topos.*

Proof. Proposition 1.4.7 implies that $\mathcal{X}_{/E}$ is presentable. Since colimits and fiber products in $\mathcal{X}_{/E}$ agree with colimits and fiber products in $\mathcal{X}$, the other axioms follow immediately. □

Remark 2.2.8. Let $\mathcal{X}$ be a presentable ∞-category in which pullbacks commute with colimits. For any morphism $p : X \to E$ in $\mathcal{X}$, we define a presheaf $\mathcal{F}$ on $\mathcal{X}$ in the following way: $\mathcal{F}(Y) = \mathrm{Hom}_E(Y \times E, X)$. Since forming the product with E preserves colimits, we deduce from Proposition 1.4.3 that $\mathcal{F}$ is representable. In other words, there exists an object of $\mathcal{X}$ which represents the functor "sections of p".

2.3. **Geometric Morphisms.** The proper notion of morphism between ∞-topoi is that of a *geometric morphism*, which we now introduce. Recall that a functor $F : \mathcal{C} \to \mathcal{C}'$ between ∞-categories is said to be *left exact* if F preserves finite limits.

Definition 2.3.1. Let $\mathcal{X}$ and $\mathcal{Y}$ be ∞-topoi. A *geometric morphism* f from $\mathcal{X}$ to $\mathcal{Y}$ is a left exact functor $f^* : \mathcal{Y} \to \mathcal{X}$ which preserves all colimits.

It follows from Corollary 1.4.6 that f^* has a right adjoint, which we shall denote by f_*.

It is clear that the class of geometric morphisms is stable under composition. We may summarize the situation by saying that there is an $(\infty, 2)$-category Top^∞ of ∞-topoi and geometric morphisms. If $\mathcal{X}$ and $\mathcal{Y}$ are ∞-topoi, we will write $\mathrm{Top}^\infty(\mathcal{X}, \mathcal{Y})$ for the ∞-category of geometric morphisms from $\mathcal{X}$ to $\mathcal{Y}$.

Example 2.3.2. Let $\mathcal{X}$ be an ∞-topos, and let $E \in \mathcal{X}$ be an object. Then there is a geometric morphism $f_E : \mathcal{X}_{/E} \to \mathcal{X}$, defined by $f_E^*(E') = E' \times E$. The geometric morphisms which arise in this way are important enough to deserve a name: a geometric morphism $f : \mathcal{Y} \to \mathcal{X}$ is said to be *étale* if it factors as a composite $f_E \circ f'$, where $f' : \mathcal{Y} \to \mathcal{X}_{/E}$ is an equivalence.

Consider an étale geometric morphism $f : \mathcal{X}_{/E} \to \mathcal{X}$. Then the forgetful inclusion of $\mathcal{X}_{/E}$ into $\mathcal{X}$ is a *left* adjoint to f^*, which we shall denote by $f_!$. It follows that f^* commutes with *all* limits. The right adjoint of f^* is slightly more difficult to describe: it is given by $f_* X = Y$, where Y is the object of X representing sections of $X \to E$ (see Remark 2.2.8).

Proposition 2.3.3. *The* $(\infty, 2)$*-category* Top^{∞} *admits arbitrary* $(\infty, 2)$*-categorical colimits. These are formed by taking the limits of the underlying* ∞*-categories.*

Proof. We have seen that any limit of ∞-pretopoi is an ∞-pretopos. In the case where the ∞-categories involved are ∞-topoi and the functors between them are left-exact, one shows that the limit is actually an ∞-topos by checking the axioms directly. □

Corollary 2.3.4. *Let* $\mathcal{C}_0$ *be a small* ∞*-category. Then the category of presheaves* $\mathcal{P} = \mathcal{S}^{\mathcal{C}_0^{op}}$ *is an* ∞*-topos.*

Proof. The ∞-category $\mathcal{P}$ is a limit of copies of $\mathcal{S}$, indexed by $\mathcal{C}_0^{op}$. Hence it is a colimit of copies of $\mathcal{S}$ in Top^{∞}. □

In fact, Top^{∞} has all $(\infty, 2)$-categorical limits as well, but this is more difficult to prove and we will postpone a discussion until the sequel to this paper.

To conclude this section, we shall show that the ∞-category of geometric morphisms between two ∞-topoi is reasonably well-behaved. First, we need to introduce a technical notion and a lemma (which will also be needed later).

Definition 2.3.5. Let κ be a regular uncountable cardinal and $\mathcal{X}$ an ∞-topos. We shall say that $\mathcal{X}$ is *κ-coherent* if the following conditions are satisfied:

- $\mathcal{X}$ is κ-accessible.
- $\mathcal{X}_\kappa$ is stable under the formation of finite limits (in $\mathcal{X}$).
- The formation of finite limits in $\mathcal{X}$ commutes with κ-filtered colimits.

Remark 2.3.6. If $\mathcal{X}$ is κ-coherent, then $\mathcal{X}$ is κ'-coherent for any regular $\kappa' \geq \kappa$.

Remark 2.3.7. This does not seem to be the right notion when $\kappa = \omega$. In $\mathcal{S}$, the compact objects (retracts of finite complexes) are not stable under finite limits. For example, the fiber product $* \times_{S^1} *$ is the infinite discrete set $\pi_1(S^1) = \mathbb{Z}$, which is not compact. Compare with Remark 1.3.2.

To conclude this section, we will show that the ∞-category of geometric morphisms between two ∞-topoi is always an accessible category. We begin with the following Lemma, which will also be needed in the proof of Theorem 2.4.1.

Lemma 2.3.8. *Let* $\mathcal{X}$ *be an* ∞*-topos. Then there exists a regular cardinal* $\kappa > \omega$ *such that* $\mathcal{X}$ *is* κ*-coherent.*

Proof. Suppose that there exists a regular cardinal κ_0 such that the formation of finite limits in $\mathcal{X}$ commutes with κ_0-filtered colimits. Enlarging $\mathcal{X}$ if necessary, we may assume that $\mathcal{X}$ is κ_0-accessible. It then suffices to choose any uncountable $\kappa \geq \kappa_0$ such that all finite limits of diagrams in $\mathcal{X}_{\kappa_0}$ lie in $\mathcal{X}_\kappa$.

It remains to prove the existence of κ_0. If $\mathcal{X} = \mathcal{S}$, then we may take $\kappa_0 = \omega$. Working componentwise, we see that the $\kappa_0 = \omega$ works whenever $\mathcal{X}$ is an ∞-category of presheaves.

By Theorem 2.4.1, we may assume that $\mathcal{X}$ is the essential image of some left exact localization functor $L : \mathcal{P} \to \mathcal{P}$, where $\mathcal{P}$ is an ∞-category of presheaves. Then L commutes with fiber products, and with κ_0-filtered colimits for κ_0 sufficiently large. Thus, the assertion for $\mathcal{X}$ follows from the assertion for $\mathcal{P}$. □

The next result will not be used in the rest of the paper, and may be safely omitted by the reader. The proof involves routine, but tedious, cardinal estimates.

Proposition 2.3.9. *Let* $\mathcal{X}$ *and* $\mathcal{Y}$ *be* ∞*-topoi. Then the* ∞*-category* $\mathrm{Top}(\mathcal{X}, \mathcal{Y})$ *geometric morphisms from* $\mathcal{X}$ *to* $\mathcal{Y}$ *is accessible.*

Proof. The ∞-category $\mathrm{Top}(\mathcal{X}, \mathcal{Y})$ is a full subcategory of $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{Y}, \mathcal{X})$, which is an ∞-pretopos in which colimits may be computed pointwise. Let κ be a regular cardinal which satisfies the conclusion of Lemma 2.3.8 for both $\mathcal{X}$ and $\mathcal{Y}$. Without loss of generality, we may enlarge κ so that $\mathcal{X}$, $\mathcal{Y}$, and $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{Y}, \mathcal{X})$ are κ-accessible.

We next note that the condition that some functor $F \in \mathrm{Fun}^{\mathrm{pres}}(\mathcal{Y}, \mathcal{X})$ be left-exact can be checked by examining a bounded amount of data. Namely, we claim that F is left exact if and only if for any objects $Y' \rightarrow Y \leftarrow Y''$ in $\mathcal{Y}_\kappa$ and any $X \in \mathcal{X}_\kappa$, the natural map $\pi^F_{X,Y} : \mathrm{Hom}_{\mathcal{X}}(X, F(Y' \times_Y Y'')) \rightarrow \mathrm{Hom}_{\mathcal{X}}(X, FY') \times_{\mathrm{Hom}_{\mathcal{X}}(X,FY)} \mathrm{Hom}_{\mathcal{X}}(X, FY'')$ is an equivalence. This condition is obviously necessary. Assume that the condition holds. *Any* diagram $Y' \rightarrow Y \leftarrow Y''$ in $\mathcal{Y}$ may be obtained as a κ-filtered colimit of diagrams $Y'_\alpha \rightarrow Y_\alpha \leftarrow Y''_\alpha$ in $\mathcal{Y}_\kappa$. Since $\mathcal{Y}$ is κ-coherent, we deduce that $Y' \times_Y Y''$ is the colimit of the diagram $\{Y'_\alpha \times_{Y_\alpha} Y''_\alpha\}$. Using the continuity of F, we see that $\pi^F_{X,Y}$ is an equivalence for every κ-compact X (with no compactness assumptions on the triple (Y, Y', Y''). Since every object of $\mathcal{X}$ can be written as a colimit of κ-compact objects, we deduce that $\pi^F_{X,Y}$ is an equivalence in general.

Let κ' be a regular cardinal $\geq \kappa$, such that $\mathcal{X}_\kappa$ and $\mathcal{Y}_\kappa$ have cardinality $< \kappa'$, and the space $\mathrm{Hom}_{\mathcal{X}}(A, B)$ is κ'-presented for all $A, B \in \mathcal{X}_\kappa$. To complete the proof, we will show that if $F \in \mathrm{Fun}^{\mathrm{pres}}(\mathcal{Y}, \mathcal{X})$ is left exact, then F can be written as a κ'-filtered colimit of left-exact objects in $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{Y}, \mathcal{X})_{\kappa'}$. Since $\mathrm{Fun}^{\mathrm{pres}}(\mathcal{Y}, \mathcal{X})$ is κ'-accessible, we may write F as a κ'-filtered colimit of some diagram $\pi : \mathcal{I} \rightarrow \mathrm{Fun}^{\mathrm{pres}}(\mathcal{Y}, \mathcal{X})_{\kappa'}$. To simplify the discussion, let us assume that $\mathcal{I}$ is a κ'-directed partially ordered set.

Let P denote the set of all directed, κ'-small subsets of $\mathcal{I}$. For each $S \in P$, let F_S denote the colimit of the system $\{\pi(s)\}_{s \in S}$. Then F may also be written as a colimit of the functors $\{F_S\}_{S \in P}$. To complete the proof, it will suffice that $P_0 = \{S \in P : F_S \in \mathrm{Top}(\mathcal{X}, \mathcal{Y})\}$ is cofinal in P.

Choose $S \in P$. We must show that there exists $S' \in P_0$ which contains S. We will construct a transfinite sequence of subsets $S_\alpha \in P$, with $S_0 = S$ and $S_\lambda = \bigcup_{\alpha < \lambda} S_\alpha$ for limit ordinals λ. For successor ordinals, our construction must be a little more complicated. Suppose that S_α has been defined, and let $F_\alpha = F_{S_\alpha}$. Consider all instances of the following data:

- $X \in \mathcal{X}_\kappa$
- $Y' \rightarrow Y \leftarrow Y'' \in \mathcal{Y}_\kappa$
- p is a point of $\mathrm{Hom}_{\mathcal{X}}(X, F_\alpha Y' \times_{F_\alpha Y} F_\alpha Y'')$
- ϕ is a map from S^n to the fiber of $\pi^{F_\alpha}_{X,Y}$, for some $n \geq 0$.

Let p' denote the image of p in $\mathrm{Hom}_{\mathcal{X}}(X, FY' \times_{FY} FY'')$. Since F is left-exact, we deduce that the induced map ϕ' from S^n to the fiber of $\pi^F_{X,Y}$ over p' is homotopic to a constant. Since the formation of this fiber is κ'-continuous, the same holds if we replace F by $F_{S_\alpha \cup T(x)}$, where $T(x) \subseteq \mathcal{I}$ is κ'-small, and x denotes the complicated data described above.

The hypotheses guarantee that there are fewer than κ' possibilities for the data x. Consequently, we may choose a filtered subset $S_{\alpha+1} \subseteq \mathcal{I}$ containing S_α and each $T(x)$. This completes the construction. We note that S_α is κ'-small for each $\alpha < \kappa'$. Set $S' = S_\kappa$.

To complete the proof, it suffices to show that F_κ is left exact. In other words, we must show that for any $X \in \mathcal{X}_\kappa$, $Y' \rightarrow Y \leftarrow Y''$ in $\mathcal{Y}_\kappa$, the map $\pi^{F_\kappa}_{X,Y}$ is an equivalence. By Whitehead's theorem, it suffices to show that for any data x of the sort described above for $\alpha = \kappa$, ϕ is already homotopic to a constant map. This follows from the construction, once we note that the formation of all relevant objects commutes with κ-filtered colimits. □

2.4. **Giraud's Theorem.** We now come to the first non-trivial result of this paper, which asserts that any ∞-topos can be obtained as a left-exact localization of an ∞-category of prestacks. This should be considered a version of Giraud's theorem for ordinary topoi, which shows that any abstract category satisfying a set of axioms actually arises as a category of sheaves on some site. Our result is not quite as specific because we do not have an explicit understanding of the localization functor as a "sheafification process".

Theorem 2.4.1. *Let $\mathcal{X}$ be an ∞-category. The following are equivalent:*

- *The ∞-category $\mathcal{X}$ is an ∞-topos.*
- *There exists a cardinal κ such that $\mathcal{X}$ is κ-accessible and the Yoneda embedding $\mathcal{X} \rightarrow \mathcal{S}^{\mathcal{X}_\kappa^{op}}$ has a left exact left adjoint.*
- *There exists an ∞-topos $\mathcal{P}$ such that $\mathcal{X}$ is a left exact localization of $\mathcal{P}$.*

In order to prove Theorem 2.4.1, we will need a rather technical lemma regarding the structure of "free groupoids", which we will now formulate. Let $\mathcal{C}$ be an ∞-category. If U is an object of $\mathcal{C}$, then we shall call a pair of morphisms $\pi_0, \pi_1 : R \to U$ a *coequalizer diagram over* U. For fixed U, the coequalizer diagrams over Y form an ∞-category $\mathfrak{CE}_U$. A *groupoid over* U is a $\mathcal{C}$-groupoid $U_\bullet$ together with an identification $U_0 \simeq U$. These groupoids also form an ∞-category which we shall denote by $\mathfrak{GPD}_U$.

Any groupoid $U_\bullet$ over U determines a coequalizer diagram over U, using the two natural maps $\pi_0, \pi_1 : U_1 \to U$. This determines a functor $G^{\mathcal{C}} : \mathfrak{GPD}_U \to \mathfrak{CE}_U$. We are interested in constructing a left-adjoint to $G^{\mathcal{C}}$. In other words, we would like to construct the "free groupoid over Y generated by a coequalizer diagram over Y". Moreover, we will need to know that the construction of this free groupoid is compatible with various functors.

Lemma 2.4.2. *Let $\mathcal{X}$ be an ∞-topos, and let U be an object of $\mathcal{X}$. Then the following assertions hold:*

- *The functor $G^{\mathcal{X}} : \mathfrak{GPD}_U \to \mathfrak{CE}_U$ has a left adjoint $F^{\mathcal{X}} : \mathfrak{CE}_U \to \mathfrak{GPD}_U$.*
- *If $T : \mathcal{X} \to \mathcal{X}'$ is a functor between ∞-topoi, then the identification $G^{\mathcal{C}'} \circ T \simeq T \circ G^{\mathcal{C}}$ induces a natural transformation $F^{\mathcal{C}'} \circ T \to T \circ F^{\mathcal{C}}$. This natural transformation is an equivalence if T preserves colimits and fiber products over U.*

Proof. We sketch the proof, which is based on Proposition 5.4.1 from the appendix, and other ideas introduced there. We begin by noting that the ∞-category $\mathfrak{CE}_U$ has a monoidal structure. The identity object is given by the diagram $1 = \{\mathrm{id}_U, \mathrm{id}_U : U \to U$. Given two coequalizer diagrams $A = \{\pi_0, \pi_1 : R \to U\}$ and $A' = \{\pi'_0, \pi'_1 : R' \to U\}$, we may define $A \otimes A' = \{\pi_0, \pi'_1 : R \otimes_U R' \to U\}$, where the fiber product is formed using π_1 and π'_0.

Since $\mathcal{X}$ is an ∞-topos, we deduce that $\mathfrak{CE}_U$ is presentable and that $\otimes$ is colimit-preserving (since colimits commute with pullback in $\mathcal{X}$). We may therefore apply Proposition 5.4.1 to produce monoid objects in $\mathfrak{CE}_U$. We note that a monoid object in $\mathfrak{CE}_U$ is simply a *category object* $U_\bullet$ of $\mathcal{X}$, together with an identification $U \simeq U_0$.

To prove the first part of the lemma, we need to show that for any object $N \in \mathfrak{CE}_U$, the "free groupoid generated by N" exists. To see this, we begin with the object $1 \in \mathfrak{CE}_U$ and note that it comes equipped with a natural monoid structure (this corresponds to the constant simplicial object $U_\bullet$ having value U). Applying Proposition 5.4.1 to the morphism $1 \to 1 \coprod N$ in $\mathfrak{CE}_U$, we deduce the existence of a monoid object $M^0 \in \mathfrak{CE}_U$. We may identify this monoid object with a simplicial object $M^0_\bullet$ in $\mathcal{X}$ which is a category object equipped with an identification $M^0_0 \simeq U$. However, we are not yet done, because $M^0_\bullet$ is not necessarily a groupoid.

Our next goal is to promote M^0 to a groupoid. To do this, we will define a transfinite sequence of monoid objects $M^\alpha \in \mathfrak{CE}_U$, equipped with a coherent system of maps $M^\beta \to M^\alpha$ for $\beta < \alpha$. We have already defined M^α for $\alpha = 0$, and when α is a limit ordinal we will simply define M^α to be the appropriate colimit. We are thereby reduced to handling the successor stages.

Assume that M^α has been defined. We will construct $M^{\alpha+1}$ by freely adjoining inverses for all of the morphisms in M^α. This can be achieved using four applications of Proposition 5.4.1. The idea is straightforward but notationally difficult to describe, so we just sketch the idea in a simplified example. Suppose that $\mathcal{C}$ is an ordinary category, and that $f : C \to C'$ is a morphism in $\mathcal{C}$ which we would like to formally adjoin an inverse for. We can obtain this inverse in three steps:

- Adjoin a new morphism $g : C' \to C$.
- Impose the relation $f \circ g = \mathrm{id}_{C'}$.
- Impose the relation $g \circ f = \mathrm{id}_C$.

If $\mathcal{C}$ is an ∞-category, then we need to be a bit more careful, since we have now identified g with the inverse of f in two different ways. In this case, we need four steps:

- Adjoin a new morphism $g : C' \to C$.
- Adjoin a path from $f \circ g$ to $\mathrm{id}_{C'}$.
- Adjoin a path from $g \circ f$ to id_C.
- Adjoin a homotopy between the two paths that we now have from $g \circ f \circ g$ to g.

The construction described above can be carried out in a relative situation, using Proposition 5.4.1. The result is that we obtain a new monoid $M^{\alpha+1} \in \mathfrak{CE}_U$. This monoid is equipped with a monoid map

$M^\alpha \rightarrow M^{\alpha+1}$. If we think of $M^{\alpha+1}$ as a category object of $\mathcal{X}$, then it is obtained from M^α by forcing all of the morphisms in M^α to become invertible.

By induction, one can easily check that for any groupoid object $V_\bullet$ of $\mathcal{C}$, the map $\mathrm{Hom}(M^\alpha_\bullet, V_\bullet) \rightarrow \mathrm{Hom}(M^\beta_\bullet, V_\bullet)$ is an equivalence. Moreover, using the fact that $\mathcal{X}$ is presentable and some cardinality estimates, one can show that M^α is itself a groupoid for sufficiently large α (it seems plausible that we can even take $\alpha = 1$, but this is not important). Thus for large α, M^α is the desired free groupoid.

The second assertion of the Lemma follows because the free groupoid was constructed using colimits and iterated application of Proposition 5.4.1, and the second assertion of Proposition 5.4.1 guarantees that the free constructions that we use are compatible with colimit preserving, monoidal functors. □

If $\mathcal{X}$ is an ∞-topos, there is an easier way to see that T has a left adjoint. Indeed, all groupoids in $\mathcal{C}$ are effective, so a groupoid $Y_\bullet$ over Y is uniquely determined by the surjective map $Y \rightarrow |Y_\bullet|$, and conversely. It is then easy to see that the left adjoint to T should assign to any coequalizer diagram $X \rightrightarrows Y$ the groupoid $Y_\bullet$ associated to the map $Y \rightarrow Z$, where Z is the coequalizer of $\pi_0, \pi_1 : X \rightarrow Y$. In particular, we have $Y_1 = Y \times_Z Y$. In particular, this proves the following:

Lemma 2.4.3. *Let $\pi_0, \pi_1 : X \rightarrow Y$ be a coequalizer diagram in an ∞-topos $\mathcal{X}$, with coequalizer Z. Suppose $F : \mathcal{X} \rightarrow \mathcal{X}'$ is a functor between ∞-topoi which commutes with all colimits, and with fiber products over Y. Then the natural map*

$$F(Y \times_Z Y) \rightarrow FY \times_{FZ} FY$$

is an equivalence.

Proof. This follows immediately from the preceding discussion and Lemma 2.4.2. □

We now proceed to the main point.

Lemma 2.4.4. *Let $\mathcal{X}$ be an ∞-topos which is κ-coherent, let $\mathcal{P} = \mathcal{S}^{\mathcal{X}_\kappa^{op}}$ be the ∞-category of prestacks on $\mathcal{X}_\kappa$, and let $G : \mathcal{X} \rightarrow \mathcal{P}$ be the Yoneda embedding. Then the left adjoint F of G is left exact.*

Proof. We first note that since $\mathcal{X}$ is κ-accessible, the left adjoint F of G exists by the proof of Proposition 1.5.3. We must show that F commutes with all finite limits. It will suffice to show that F commutes with pullbacks and preserves the final object. The second point is easy: since the final object $1 \in \mathcal{X}$ is κ-compact, it follows that the final object in $\mathcal{P}$ is the presheaf represented by $1 \in \mathcal{X}_\kappa$, and F carries this presheaf into 1.

Let us now show that F commutes with pullbacks. In other words, we must show that for any diagram $X \rightarrow Y \leftarrow Z$ in $\mathcal{P}$, the natural morphism

$$\eta : F(X \times_Y Z) \rightarrow FX \times_{FY} FZ$$

is an equivalence. We proceed as follows: let S denote the collection of all objects $Y \in \mathcal{P}$ such that η is an isomorphism for *any* pair of objects X, Z over Y. We need to show that every object of $\mathcal{P}$ lies in S. Since every prestack on $\mathcal{X}_\kappa$ is a colimit of representable prestacks, it will suffice to show that every representable prestack lies in S and that S is closed under the formation of colimits.

First of all, we note that since both $\mathcal{X}$ and $\mathcal{P}$ are ∞-topoi, colimits commute with pullbacks. Since F commutes with all colimits, we see that both $F(X \times_Y Z)$ and $FX \times_{FY} FZ$ are compatible with arbitrary colimits in X and Z. Since every prestack (over Y) is a colimit of representable prestacks (over Y), in order to show that $Y \in S$ it suffices to show that η is an isomorphism whenever X and Z are representable. Suppose that X, Y, and Z are all representable by objects $x, y, z \in \mathcal{X}_\kappa$, and let $w = x \times_y z$ so that w represents the prestack $W = X \times_Y Z$. Then $L(X \times_Y Z) \simeq L(W) \simeq w \simeq x \times_y z \simeq LX \times_{LY} LZ$, so that $Y \in S$. Therefore S contains every representable prestack.

To complete the proof, we must show that S is stable under the formation of colimits. It will suffice to show that S is stable under the formation of sums and coequalizers (see Appendix 5.1). We consider these two cases separately.

Suppose $Y \in \mathcal{P}$ is a sum of some family of objects $\{Y_\alpha\}_{\alpha \in A} \subseteq S$. As above, we may assume that X and Z are represented by objects $x, z \in \mathcal{X}_\kappa$. The maps $X, Z \rightarrow Y$ correspond to points $p_x \in Y_\alpha(x)$, $p_z \in Y_\beta(z)$

for some indices α and β. Then $L(X \times_Y Z) = L(\emptyset) = \emptyset$ if $\alpha \neq \beta$, and

$$L(X \times_Y Z) = L(X \times_{Y_\alpha} Z) = LX \times_{LY_\alpha} LZ$$

if $\alpha = \beta$, where the last equality uses the fact that $Y_\alpha \in S$. On the other hand,

$$LX \times_{LY} LZ = x \times_{\coprod_\gamma LY_\gamma} z = x \times_{LY_\alpha} (LY_\alpha \times_{\coprod_\gamma LY_\gamma} LY_\beta) \times_{LY_\beta} z$$

. Since sums are disjoint in $\mathcal{X}$, this fiber product is empty if $\alpha \neq \beta$ and is equal to $x \times_{LY_\alpha} z$ otherwise.

We now come to the core of the argument, which is showing that S is stable under the formation of coequalizers. Fortunately, the hard work is already done. Suppose that Y is the coequalizer of a diagram $Y_1 \rightrightarrows Y_0$ in $\mathcal{P}$, where $Y_0, Y_1 \in S$. As above, we may assume that X and Z are representable by objects $x, z \in \mathcal{C}_\kappa$. Then any point of $\mathrm{Hom}_{\mathcal{P}}(X, Y) = Y(x)$ may be lifted to a point of $Y_0(x)$. Thus we may assume that the maps $X, Z \to Y$ both factor through Y_0. Since $X \times_Y Z = X \times_{Y_0} (Y_0 \times_Y Y_0) \times_{Y_0} Z$ and we already know that L commutes with fiber products over Y_0, we can reduce to the case where $X = Z = Y_0$ (no longer assuming X and Z to be representable). Now we simply apply Lemma 2.4.3, noting that L commutes with all colimits and with fiber products over Y_0. □

We can now give the proof of Theorem 2.4.1:

Proof. We have just seen that (1) ⇒ (2). Since prestack ∞-categories are ∞-topoi, it is obvious that (2) ⇒ (3). Let us prove that (3) ⇒ (1). By Proposition 1.5.3, $\mathcal{X}$ is presentable. It remains to check the other axioms. Let $L : \mathcal{P} \to \mathcal{X}$ denote the left exact localization functor.

Sums in $\mathcal{X}_0$ are disjoint: given two objects E and E' of $\mathcal{X}_0$, their sum in $\mathcal{X}_0$ is $L(E \coprod E')$, so that the fiber product $E \times_{L(E \coprod E')} E' \simeq LE \times_{L(E \coprod E')} LE' \simeq L(E \times_{E \coprod E'} E') \simeq L(\emptyset)$, which is the initial object of $\mathcal{X}_0$.

Finally, suppose that $X_\bullet$ is a groupoid in $\mathcal{X}_0$. Then for any $Y \in \mathcal{X}$, $\mathrm{Hom}(Y, X_\bullet) = \mathrm{Hom}(LY, X_\bullet)$, so that $X_\bullet$ is also a groupoid in $\mathcal{X}$. Thus $X_\bullet$ is effective in $\mathcal{X}$. Since L is left exact, it follows that $X_\bullet$ is effective in $\mathcal{X}_0$. □

Theorem 2.4.1 is extremely useful. It allows us to reduce the proofs of many statements about arbitrary ∞-topoi to the case of ∞-categories of prestacks. We can then often reduce to the case of the ∞-topos $\mathcal{S}$ by working componentwise. This places the full apparatus of classical homotopy theory at our disposal.

2.5. **Truncated Objects.**

In this section, we will show that for any ∞-topos $\mathcal{X}$, the full subcategory of $\mathcal{X}$ consisting of "discrete objects" is an (ordinary) topos. This observation gives us a functor from the $(\infty, 2)$-category of ∞-topoi to the 2-category of topoi.

We will begin with some generalities concerning truncated objects in an ∞-category.

Definition 2.5.1. Let $\mathcal{C}$ be an ∞-category and k an integer ≥ -1. An object $C \in \mathcal{C}$ is called *k-truncated* if for any morphism $\eta : D \to C$, one has $\pi_n(\mathrm{Hom}_{\mathcal{C}}(D, C), \eta) = *$ for $n > k$. By convention we shall say that an object is (-2)-truncated if it is a final object of $\mathcal{C}$. An object C is called *discrete* if it is 0-truncated.

In other words, an object $C \in \mathcal{C}$ is discrete if $\mathrm{Hom}_{\mathcal{C}}(D, C)$ is a discrete set for any $D \in \mathcal{C}$. We see immediately that the discrete objects of $\mathcal{C}$ form an ordinary category.

Remark 2.5.2. The k-truncated objects in an ∞-category $\mathcal{C}$ are stable under the formation of all limits which exist in $\mathcal{C}$.

Example 2.5.3. Suppose $\mathcal{C} = \mathcal{S}$. Then the (-2)-truncated objects are precisely the contractible spaces. The (-1)-truncated objects are spaces which are either empty or contractible. The 0-truncated objects are the spaces which are discrete (up to homotopy). More generally, the n-truncated objects are the spaces X all of whose homotopy groups $\pi_k(X, x)$ vanish for $k > n$.

Let us say that a morphism $f : C \to C'$ in an ∞-category $\mathcal{C}$ is *k-truncated* if it exhibits C as a k-truncated object in the slice category $\mathcal{C}_{/C'}$ of objects over C'. For example, the (-2)-truncated morphisms are precisely the equivalences in $\mathcal{C}$. The (-1)-truncated morphisms are called *monomorphisms.* A morphism $f : C \to C'$

is a monomorphism if the induced map $\mathrm{Hom}_{\mathcal{C}}(X, C) \to \mathrm{Hom}_{\mathcal{C}}(X, C')$ is equivalent to the inclusion of a summand, for any $X \in \mathcal{C}$.

The following easy lemma, whose proof is left to the reader, gives a recursive characterization of n-truncated morphisms.

Lemma 2.5.4. *Let $\mathcal{C}$ be an ∞-category with finite limits and let $k \geq -1$. A morphism $f : C \to C'$ is k-truncated if and only if the diagonal $C \to C \times_{C'} C$ is $(k-1)$-truncated.*

This immediately implies the following:

Proposition 2.5.5. *Let $F : \mathcal{C} \to \mathcal{C}'$ be a left-exact functor between ∞-categories which admit finite limits. Then F carries k-truncated objects into k-truncated objects and k-truncated morphisms into k-truncated morphisms.*

Proof. An object C is k-truncated if and only if the morphism $C \to 1$ to the final object is k-truncated. Since F preserves the final object, it suffices to prove the assertion concerning morphisms. Since F commutes with fiber products, Lemma 2.5.4 allows us to use induction on k, thereby reducing to the case where $k = -2$. But the (-2)-truncated morphisms are precisely the equivalences, and these are preserved by any functor. □

Proposition 2.5.5 applies in particular whenever $F = f_*$ or $F = f^*$, if f is a geometric morphism between ∞-topoi.

Proposition 2.5.6. *Let $\mathcal{C}$ be a presentable ∞-category, $k \geq -2$. Then there exists an accessible functor $\tau_k : \mathcal{C} \to \mathcal{C}$, together with a natural transformation $\mathrm{id}_{\mathcal{C}} \to \tau_k$, with the following property: for any $C \in \mathcal{C}$, the object $\tau_k C$ is k-truncated, and the natural map $\mathrm{Hom}_{\mathcal{C}}(\tau_k C, D) \to \mathrm{Hom}_{\mathcal{C}}(C, D)$ is an equivalence for any k-truncated object D.*

Remark 2.5.7. If the ∞-category $\mathcal{C}$ is potentially unclear in context, then we will write $\tau_k^{\mathcal{C}}$ for the truncation functor in $\mathcal{C}$.

Proof. Suppose first that $\mathcal{C} = \mathcal{S}$. In this case, the functor τ_k is the classical "Postnikov truncation" functor. It is characterized by the fact that there exists a surjective map $X \to \tau_k X$ which induces an isomorphism on π_n for $n \leq k$, and $\pi_n(\tau_k X) = 0$ for $n > k$, for any choice of basepoint. If we represent spaces by fibrant simplicial sets, then the functor τ_k can be implemented by replacing a simplicial set X by its k-coskeleton. In particular, we note that τ_k is ω-continuous.

Now suppose that $\mathcal{C}$ is an ∞-category of prestacks. The existence of τ_k in this case follows easily from the existence of τ_k when $\mathcal{C} = \mathcal{S}$: we simply work componentwise. Once again, the functor τ_k is actually ω-continuous.

We now handle the case of a general presentable ∞-category $\mathcal{C}$. By Proposition 1.5.3, we may view $\mathcal{C}$ as the image of a localization functor $L : \mathcal{C}' \to \mathcal{C}'$, where $\mathcal{C}'$ is an ∞-category of prestacks. We will suppose that a truncation functor $\tau_k^{\mathcal{C}'}$ has already been constructed for $\mathcal{C}'$.

We now construct a transfinite sequence of functors $\tau_k^{\alpha} : \mathcal{C} \to \mathcal{C}$, together with a coherent collection of natural transformations $\tau_k^{\alpha_0} \to \ldots \to \tau_k^{\alpha_n}$ for $\alpha_0 \leq \ldots \leq \alpha_n$. Let τ_k^0 be the identity functor, and for limit ordinals λ we let τ_k^{λ} be the (filtered) colimit of $\{\tau_k^{\alpha}\}_{\alpha < \lambda}$. Finally, let $\tau_k^{\alpha+1} = L \tau_k^{\mathcal{C}'} \tau_k^{\alpha}$.

One verifies easily by induction that for any k-truncated object $D \in \mathcal{C}$, the natural morphism

$$\mathrm{Hom}_{\mathcal{C}}(\tau_k^{\alpha} C, D) \to \mathrm{Hom}_{\mathcal{C}}(C, D)$$

is an equivalence. Note also that each of the functors τ_k^{α} is accessible. To complete the proof, we note that L is κ-continuous for some κ. Thus κ-filtered colimits in $\mathcal{C}$ agree with κ-filtered colimits in $\mathcal{C}'$, so that $\tau_k^{\kappa} C$ is a κ-filtered colimit of k-truncated objects in $\mathcal{C}'$, hence k-truncated in $\mathcal{C}'$ and thus also in $\mathcal{C}$. We may then take $\tau_k = \tau_k^{\kappa}$. □

Remark 2.5.8. The preceding proof can be simplified if $\mathcal{C}$ is an ∞-topos. In this case, the localization functor L can be chosen to be left exact, and therefore to preserve the property of being k-truncated, so that one may simply define $\tau_k = L \tau_k^{\mathcal{C}'}$.

Remark 2.5.9. It is immediate from the definition of τ_k that it is a localization functor. We deduce immediately that if $\mathcal{C}$ is a presentable ∞-category, then the subcategory $\tau_k \mathcal{C}$ of k-truncated objects of $\mathcal{C}$ is also presentable. In particular, the discrete objects of $\mathcal{C}$ form a presentable category.

Proposition 2.5.10. *Let $F : \mathcal{C} \to \mathcal{C}'$ be a functor between presentable ∞-categories, which is left-exact and colimit preserving. Then there are natural equivalences $F(\tau_k X) \simeq \tau_k F(X)$.*

Proof. Since F is left-exact, it preserves k-truncated objects. Thus the natural map $FX \to F(\tau_k X)$ factors uniquely through some map $\phi : \tau_k FX \to F(\tau_k X)$. To show that ϕ is an equivalence, we must show that $\mathrm{Hom}_{\mathcal{C}'}(FX, W) \simeq \mathrm{Hom}_{\mathcal{C}}(F(\tau_k X), W)$ is an equivalence for any k-truncated W in $\mathcal{C}'$.

Since F preserves colimits, it has a right adjoint G. Thus we can rewrite the desired conclusion as $\mathrm{Hom}_{\mathcal{C}}(X, GW) \simeq \mathrm{Hom}_{\mathcal{C}}(\tau_k X, GW)$. Since G preserves all limits, GW is k-truncated and we are done. □

Proposition 2.5.10 applies in particular whenever $F = f^*$, where f is a geometric morphism between ∞-topoi.

We will now need a few basic facts about truncated objects in an ∞-topos.

Proposition 2.5.11. *Let $\mathcal{X}$ be an ∞-topos containing a diagram $X \to Y \leftarrow Z$, where Y and Z are k-truncated. Then the natural map $f : \tau_k(X \times_Y Z) \to (\tau_k X) \times_Y Z$ is an equivalence.*

Proof. By Theorem 2.4.1, we may view $\mathcal{X}$ as a left-exact localization of an ∞-category $\mathcal{P}$ of prestacks. Let L denote the localization functor. Since $\tau_k^{\mathcal{X}} = L\tau_k^{\mathcal{P}}$, it will suffice to prove the proposition after $\mathcal{X}$ has been replaced by $\mathcal{P}$. Working componentwise, we can reduce to the case where $\mathcal{X} = \mathcal{S}$.

In order to show that f is a homotopy equivalence, it will suffice to show that f induces isomorphisms on homotopy sets for any choice of basepoint $p \in \tau_k(X \times_Y Z)$. We may lift p to a point $\widetilde{p} \in X \times_Y Z$. Since both the source and target of f are k-truncated, it will suffice to prove that $\pi_n(X \times_Y Z, \widetilde{p}) \simeq \pi_n(\tau_k X \times_Y Z, f(p))$ is bijective for $n \leq k$. If $n = 0$, this follows by inspection, since X and $\tau_k X$ have the same connected component structure when $k \geq 0$.

If $n > 0$, we use the fiber sequences $F \to X \times_Y Z \to X$ and $F \to (\tau_k X) \times_Y Z \to \tau_k X$ with the same fiber (containing the point $\widetilde{p}$). Using the long exact sequence associated to these fibrations, we deduce the bijectivity from the (slightly nonabelian) five lemma, using the fact that $\pi_j(X, x) \to \pi_j(\tau_k X, x)$ is bijective for $j \leq k$ and surjective for $j = k + 1$. □

We can now show how every ∞-topos $\mathcal{X}$ determines an ordinary topos. The proof requires Corollary 2.6.8, which is proven in the next section.

Theorem 2.5.12. *Let $\mathcal{X}$ be an ∞-topos. Then the full subcategory of discrete objects of $\mathcal{X}$ is a topos.*

Proof. We have seen that these objects form a presentable category $\tau_0 \mathcal{X}$. To complete the proof, it will suffice to verify the remainder of Giraud's axioms (see Proposition 2.1.1), namely:

- Sums in $\tau_0 \mathcal{X}$ are disjoint. It will suffice to show that fiber products in $\tau_0 \mathcal{X}$ agree with the (homotopy) fiber products in $\mathcal{X}$ and that sums in $\tau_0 \mathcal{X}$ agree with sums in $\mathcal{X}$. The first assertion follows from the definition of a discrete object (both fiber products enjoy the same universal property). To prove the second, it suffices to show that a sum of discrete objects in $\mathcal{X}$ is discrete, which follows easily from the disjointness of sums in $\mathcal{X}$.
- Colimits in $\tau_0 \mathcal{X}$ commute with pullback. Since $\tau_0 \mathcal{X}$ is a localization of $\mathcal{X}$, colimits in $\tau_0 \mathcal{X}$ may be computed by first forming the appropriate colimit in $\mathcal{X}$ and then applying τ_0. The required commutativity follows immediately from Proposition 2.5.11.
- All equivalence relations in $\mathcal{X}_1$ are effective. If $R \subseteq U \times U$ is an equivalence relation in $\tau_0 \mathcal{X}$, then we obtain a groupoid $X_\bullet$ in $\tau_0 \mathcal{X}$, with $X_n = R \times_U \ldots \times_U R$. Since $\mathcal{X}$ is an ∞-topos, this groupoid is effective. Let $X = |X_\bullet|$. To complete the proof, it will suffice to show that X is discrete. This follows from Corollary 2.6.8, since $U \times_X U \to U \times U$ is a monomorphism (by the definition of an equivalence relation).

□

The construction $\mathcal{X} \rightsquigarrow \tau_0\,\mathcal{X}$ is functorial on the $(\infty, 2)$-category of ∞-topoi. Indeed, if $f : \mathcal{X} \to \mathcal{X}'$ is a geometric morphism of ∞-topoi, then f^* and f_* are both left exact and therefore carry discrete objects into discrete objects. They therefore induce (left exact) adjoint functors between $\tau_0\,\mathcal{X}$ and $\tau_0\,\mathcal{X}'$, which is precisely the definition of a geometric morphism of topoi.

One should think of $\tau_0\,\mathcal{X}$ as a sort of "Postnikov truncation" of $\mathcal{X}$. The classical 1-truncation of a homotopy type X remembers only the fundamental groupoid of X. It therefore knows all about local systems of sets on X, but nothing about fibrations over X with non-discrete fibers. The relationship between $\mathcal{X}$ and $\tau_0\,\mathcal{X}$ is analogous: $\tau_0\,\mathcal{X}$ knows about the "sheaves of sets" on $\mathcal{X}$, but has forgotten about sheaves with nondiscrete spaces of sections.

Remark 2.5.13. In view of the above, the notation $\tau_0\,\mathcal{X}$ is unfortunate because the analogous notation for the 1-truncation of a homotopy type X is $\tau_1 X$. We caution the reader not to regard $\tau_0\,\mathcal{X}$ not as the result of applying an operation τ_0 to $\mathcal{X}$; it instead denotes the essential image of the functor $\tau_0 : \mathcal{X} \to \mathcal{X}$.

We would expect, by analogy with homotopy theory, that there in some sense a geometric morphism $\mathcal{X} \to \tau_0\,\mathcal{X}$. In order to properly formulate this, we need to show how to build an ∞-topos from an ordinary topos. We will return to this idea after a brief digression.

2.6. **Surjections.** Let $\mathcal{X}$ be an ∞-topos. The assumption that every groupoid in $\mathcal{X}$ is effective leads to a good theory of surjections (also called "effective epimorphisms") in $\mathcal{X}$, as we shall explain in this section.

Let U_0 be an object in an ∞-topos $\mathcal{X}$. We define a groupoid object $U_\bullet$ of $\mathcal{X}$ by letting U_n denote the $(n+1)$-fold product of U_0 with itself. Let $U = |U_\bullet|$. Then we have the following:

Proposition 2.6.1. *The morphism $U_0 \to U$ is equivalent to the natural morphism $U_0 \to \tau_{-1} U_0$.*

Proof. We first show that U is (-1)-truncated. It suffices to show that the diagonal map $U \to U \times U$ is an equivalence. Note that $U = U \times_U U$. Since colimits commute with fiber products and $U = |U_n|$, it suffices to show that $p_n : U_n \times_U U_n \to U_n \times U_n$ is an equivalence for all $n \geq 0$. Note that $p_n = U_n \times_{U_0} p_0 \times_{U_0} U_n$; thus it suffices to consider the case where $n = 0$. In this case, the fiber product $U_0 \times U_0 = U_1$ by definition, and $U_0 \times_U U_0 = U_1$ since $U_\bullet$ is effective.

To complete the proof, it suffices to show that $\mathrm{Hom}_{\mathcal{X}}(U, E) \to \mathrm{Hom}_{\mathcal{X}}(U_0, E)$ is an equivalence whenever E is (-1)-truncated. By definition, the condition on E means that both spaces are either empty or contractible; we must show that if the target $\mathrm{Hom}_{\mathcal{X}}(U_0, E)$ is nonempty, then so is $\mathrm{Hom}_{\mathcal{X}}(U, E)$. But this is clear from the description of U as a colimit of objects, each of which admits a map to E (automatically unique). $\square$

If $U \simeq 1$, then we shall say that $U_0 \to 1$ is a *surjection*. More generally, we shall say that a morphism $U_0 \to E$ is a *surjection* if it is a surjection in the ∞-topos $\mathcal{X}_{/E}$. In other words, $U_0 \to E$ is a surjection if $E \simeq |U_\bullet|$, where U_n denotes the $(n+1)$-fold fiber power of U_0 over E.

Remark 2.6.2.
(1) Any equivalence is surjective.
(2) Any morphism which is homotopic to a surjection is surjective.
(3) If $f : \mathcal{X} \to \mathcal{Y}$ is a geometric morphism and $Y \to Y'$ is a surjection in $\mathcal{Y}$, then $f^*Y \to f^*Y'$ is a surjection in $\mathcal{X}$. In particular, we see that $Y \times_{Y'} Y'' \to Y''$ is surjective for any $Y'' \to Y'$.

To show that surjections are closed under composition, it is convenient to recast the definition of a surjection in the following way. Given an object $E \in \mathcal{X}$, let $\mathcal{P}(E)$ denote the collection of all "subobjects" of E: namely, all equivalence classes (-1)-truncated morphisms $E_0 \to E$. We may partially order this collection so that $E_0 \leq E_1$ if there is a factorization (automatically unique) $E_0 \to E_1 \to E$. It is not difficult to verify that $\mathcal{P}(E)$ is actually a set (it is the set of subobjects of 1 in the topos $\tau_0\,\mathcal{X}_{/E}$), but we shall not need this. Note that given any morphism $f : E' \to E$, we obtain an order-preserving pullback map $f^* : \mathcal{P}(E) \to \mathcal{P}(E')$, given by $f^*(E_0') = E_0' \times_{E'} E$.

Proposition 2.6.3. *A morphism $f : E' \to E$ in $\mathcal{X}$ is surjective if and only if $f^* : \mathcal{P}(E) \to \mathcal{P}(E')$ is injective.*

Proof. Suppose first that f^* is injective. Let U_n denote the $(n+1)$-fold fiber power of E' over E. Then $|U_\bullet| \to E$ is (-1)-truncated. Thus $|U_\bullet|$ and E give two elements of $\mathcal{P}(E)$. It is easy to see that $f^*|U_\bullet| = f^*E \in \mathcal{P}(E')$. The injectivity of f^* then implies that $|U_\bullet| \to E$ is an equivalence, so that f is surjective.

Conversely, suppose that f is surjective, so that $|U_\bullet| \to E$ is an equivalence. Consider any pair of injections $E_0 \to E$, $E_1 \to E$. Suppose that $f^*E_0 = f^*E_1$ in $\mathcal{P}(E')$; we must show that $E_0 = E_1$ in $\mathcal{P}(E)$. It will suffice to show that $E_0 \times_E E_1$ maps isomorphically to both E_0 and E_1. By symmetry it will suffice to prove this for E_1. Replacing E_0 by $E_0 \times_E E_1$, we may suppose that there is a factorization $E_0 \stackrel{g}{\to} E_1 \to E$. We wish to show that g is an equivalence.

Since pullbacks commute with colimits, the morphism g is obtained as a colimit of morphisms $g_n : E_0 \times_E U_n \to E_1 \times_E U_n$. Thus it suffices to show that each g_n is an equivalence. Since $g_n = g_0 \times_{U_0} U_n$, it suffices to show that g_0 is an equivalence. But this is precisely what the assumption that $f^*E_0 = f^*E_1$ tells us. □

From this we immediately deduce some corollaries.

Corollary 2.6.4. *A composite of surjective morphisms is surjective.*

Corollary 2.6.5. *Let $f_\alpha : E'_\alpha \to E_\alpha$ be a collection of morphisms in an ∞-topos $\mathcal{X}$, and let $f : E' \to E$ be their sum. Then f is surjective if and only if each f_α is surjective.*

Proof. Note that $\mathcal{P}(E) = \Pi_\alpha \mathcal{P}(E_\alpha)$ and $\mathcal{P}(E') = \Pi_\alpha \mathcal{P}(E'_\alpha)$. If each f^*_α is injective, then $f^* = \Pi_\alpha f^*_\alpha$ is also injective. The converse also holds, since each of the factors $\mathcal{P}(E_\alpha)$ is nonempty. □

Corollary 2.6.6. *Suppose $X \stackrel{f}{\to} Y \stackrel{g}{\to} Z$ is a diagram with $g \circ f$ surjective. Then g is surjective.*

The notion of a surjective morphism gives a good mechanism for proving theorems by descent. The following gives a sampler:

Proposition 2.6.7. *Let $\mathcal{X}$ be an ∞-topos, let $f : E \to S$ be a morphism in $\mathcal{X}$, let $g : S' \to S$ be another morphism, and let $f' : E' \to S'$ be the base change of f by $S' \to S$. Then:*

- *If the morphism f is n-truncated, then so is f'. The converse holds if $S' \to S$ is surjective.*
- *If the morphism f is surjective, then so is f'. The converse holds if $S' \to S$ is surjective.*

Proof. We begin with the first claim. If f is n-truncated, then f' is clearly n-truncated. For the converse, we note that f is n-truncated if and only it induces an equivalence $E \to \tau_n^{\mathcal{X}/S} E$. Since base change by S' commutes with truncation, it will suffice to prove that a morphism which becomes an equivalence after base change to S' is an equivalence to begin with (in other words, it suffices to consider the case where $n = -2$).

Let $S'_\bullet$ denote the groupoid obtained by taking fiber powers of S' over S, and let $E'_\bullet$ be the groupoid obtained by taking iterated fiber powers of E' over E. Then if f' is an equivalence, we get an induced equivalence $E'_\bullet \to S'_\bullet$, hence an equivalence $|E'_\bullet| \to |S'_\bullet|$. Since g is surjective, the target is equivalent to S. The map $E' \to E$ is a base change of g, hence surjective, so the same argument shows that the source $|E'_\bullet|$ is equivalent to E, which completes the proof.

Now let us consider the second claim. We have already seen (and used) that the surjectivity of f implies the surjectivity of f'. For reverse direction, use Corollary 2.6.6. □

The following consequence was needed in the last section:

Corollary 2.6.8. *Let $\mathcal{C}$ be an ∞-topos, $k \geq -1$, and let $f : U \to X$ be a surjective morphism in $\mathcal{C}$. Suppose that $U \times_X U \to U \times U$ is $(k-1)$-truncated. Then X is k-truncated.*

Proof. The morphism $U \times_X U \to U \times U$ is obtained from the diagonal morphism $X \to X \times X$ by a surjective base change. It follows that the diagonal of X is $(k-1)$-truncated, so that X is k-truncated by Lemma 2.5.4. □

We saw in the last section that if $\mathcal{X}$ is an ∞-topos, then the full subcategory $\mathcal{X}_1$ consisting of discrete objects is an ordinary topos. By Giraud's theorem 2.4.1, $\mathcal{X}_1$ actually arises as the category of sheaves on some site equipped with a Grothendieck topology. We will now construct such a site.

Suppose that $\mathcal{X}$ is κ-coherent. Via the Yoneda embedding, every 1-truncated object determines a presheaf of *sets* on $\mathcal{X}_\kappa$. Note that a presheaf on an ∞-category $\mathcal{C}$ is the same thing as a presheaf on the homotopy category $h\,\mathcal{C}$. We will show that $\mathcal{X}_1$ is precisely the collection of sheaves of sets for a Grothendieck topology on $h\,\mathcal{X}_\kappa$ which we now describe. We will say that a collection of morphisms $\{E_\alpha \to E\}$ in $h\,\mathcal{X}_\kappa$ is a covering of E if the induced map (which is well-defined up to homotopy) $\coprod_\alpha E_\alpha \to E$ is surjective.

Proposition 2.6.9. *The covering families defined above determine a Grothendieck topology on $h\,\mathcal{X}_\kappa$. The presheaf represented by any discrete object of $\mathcal{X}$ is a sheaf of sets on $h\,\mathcal{X}_\kappa$. The Yoneda embedding induces an equivalence of categories between $\mathcal{X}_1$ and the category of sheaves of sets on $h\,\mathcal{X}_\kappa$.*

Proof. Let us first show that the covering families determine a Grothendieck topology on $h\,\mathcal{X}_\kappa$.

- We must show that the one element family $\{E \to E\}$ is covering for any E, and that any covering family remains a covering family when it is enlarged. Both of these properties are clear from the definitions.
- Suppose that some family of morphisms $\{E_\alpha \to E\}$ give a covering of E, and the for each α we have a covering $\{E_{\alpha\beta} \to E_\alpha\}$. We must show that the family of composites $\{E_{\alpha\beta} \to E\}$ is also a covering family. The induced morphism
$$\coprod_{\alpha,\beta} E_{\alpha\beta} \to \coprod_\alpha E_\alpha$$
is surjective by Corollary 2.6.5, hence the morphism
$$\coprod_{\alpha,\beta} E_{\alpha\beta} \to E$$
is surjective by Corollary 2.6.4.
- Let $\{E_\alpha \to E\}$ be a covering and $f : E' \to E$ any morphism. We must show that there exists a covering of $\{E'_\beta \to E'\}$ such that each composite map $E'_\beta \to E$ factors through some E_α. In fact, we may take the family of morphisms $\{E_\alpha \times_E E' \to E'\}$. (Note that $E_\alpha \times_E E'$ is *not necessarily* a fiber product in $h\,\mathcal{C}_\kappa$: these do not necessarily exist, and are not needed).

Now suppose that $X \in \mathcal{X}_1$ is a discrete object of $\mathcal{X}$. We claim that $\mathcal{F}(\bullet) = \mathrm{Hom}(\bullet, X)$ is a sheaf on $h\,\mathcal{C}_\kappa$. Suppose given a covering $\{E_\alpha \to E\}$ of an object E and elements $\eta_\alpha \in \mathcal{F}(E_\alpha)$. We must show that if the η_α agree on overlaps, then they glue to a unique section of $\mathcal{F}(E)$. Set $E' = \coprod E_\alpha$ and U_n denote the $(n+1)$-fold fiber power of E' over E. Then $U_\bullet$ is a groupoid with $|U_\bullet| = E$. Consequently, we have $\mathrm{Hom}_{\mathcal{X}}(E, X) = |\mathrm{Hom}_{\mathcal{X}}(U_\bullet, X)|$. Since X is discrete, the right hand side is simply the kernel of the pair of maps $p, q : \mathcal{F}(E') \to \mathcal{F}(E' \times_E E')$. The condition that the η_α agree on overlap is precisely the condition that they give rise to an element of this kernel.

Let $\mathcal{C}$ denote the topos of sheaves on $h\,\mathcal{X}_\kappa$. Since $\mathcal{X}$ was assumed to be κ-accessible, the Yoneda embedding $\pi : \mathcal{X}_1 \to \mathcal{C}$ is fully faithful. To complete the proof, it will suffice to show that π is essentially surjective. In other words, we must prove that any sheaf of sets $\mathcal{F}$ on $h\,\mathcal{C}_\kappa$ is representable by an object of $\mathcal{X}$ (automatically discrete). Regard $\mathcal{F}$ as a presheaf on $\mathcal{C}_\kappa$; we need to show that it is representable. If $\mathcal{F}$ carries all κ-small colimits into limits, then it has a unique κ-continuous extension to $\mathcal{X}$ which is representable by Proposition 1.4.3. In order to show that $\mathcal{F}$ commutes with all κ-small limits, it will suffice to show that $\mathcal{F}$ commutes with κ-small sums and with coequalizers (see Appendix 5.1).

If E is a sum of κ-compact objects E_α, then $\{E_\alpha \to E\}$ is a covering family so that $\mathcal{F}(E) \to \Pi_\alpha \mathcal{F}(E_\alpha)$ is injective. The surjectivity follows from the fact that sums are disjoint in $\mathcal{X}$, so that no compatibility conditions are required to glue elements of $\mathcal{F}(E_\alpha)$ to an element of $\mathcal{F}(E)$.

Finally, suppose E is presented as a homotopy coequalizer of a diagram $\pi_0, \pi_1 : F_1 \to F_0$ of κ-compact objects. We must show that the diagram $\mathcal{F}(E) \to \mathcal{F}(F_0) \rightrightarrows \mathcal{F}(F_1)$ is a (homotopy) equalizer. Since $\mathcal{F}$ is set-valued, this means only that $\mathcal{F}(E)$ is the kernel of the pair of maps $\mathcal{F}(F_0) \rightrightarrows \mathcal{F}(F_1)$. The injectivity of $\mathcal{F}(E) \to \mathcal{F}(F_0)$ follows from the fact that $\{F_0 \to E\}$ is a covering (and that $\mathcal{F}$ is a separated presheaf).

To complete the proof, we need to show that given any η in the kernel of $\pi_0, \pi_1 : \mathcal{F}(F_0) \to \mathcal{F}(F_1)$ arises from an object in $\mathcal{F}(E)$. Since $F_0 \to E$ is surjective, it suffices to show that the two restrictions of η to $\mathcal{F}(F_0 \times_E F_0)$ coincide. Let $\pi'_0, \pi'_1 : F_0 \times_E F_0 \to F_0$ denote the two projections.

To complete the proof, we need to recall the notation of Lemma 2.4.2: the ∞-category $\mathfrak{CE}_{F_0}$ has a coherently associative multiplication $\otimes$. Let $\mathcal{C}$ denote the full subcategory of $\mathfrak{CE}_{F_0}$ consisting of those triples (A, p, q) such that $p^*\eta = q^*\eta \in \mathcal{F}(A)$. It is clear from the definition that $\mathcal{C}$ is stable under $\otimes$ and under direct sum. Since $\mathcal{F}$ is a separated presheaf, any surjective morphism $\psi : A \to A'$ induces an injection $\psi^* : \mathcal{F}(A') \to \mathcal{F}(A)$. It follows that if $(A, p, q) \in \mathcal{C}$ and there is a morphism $(A, p, q) \to (A', p', q')$ in $\mathfrak{CE}_{F_0}$ which induces a surjection $A \to A'$, then $(A', p', q') \in \mathcal{C}$.

Let R_{ij} denote the object $(F_1, \pi_i, \pi_j) \in \mathfrak{CE}_Y$. The proof of Lemma 2.4.2 implies that we can construct $F_0 \times_E F_0$ as a colimit of $\otimes$-products of the objects R_{ij}. Consequently, there exists a morphism $(A, p, q) \to (F_0 \times_E F_0, \pi'_0, \pi'_1)$ which induces a surjection $A \to F_0 \times_E F_0$, where A is a disjoint union of $\otimes$-products of the objects R_{ij}. We need to show that $(F_0 \times_E F_0, \pi'_0, \pi'_1) \in \mathcal{C}$. By the preceding argument, it suffices to show that each R_{ij} lies in $\mathcal{C}$, which is clear. □

Now we come to a crucial fact which will be needed subsequently: a map is surjective if and only if it induces a surjective map on connected components. Note that a map between discrete objects of $\mathcal{X}$ is surjective if and only if it is a surjective map in the ordinary topos $\tau_0\,\mathcal{X}$, in the usual sense.

Proposition 2.6.10. *Let $\mathcal{X}$ be an ∞-topos, and let $\phi : E \to F$ be a morphism in $\mathcal{X}$. Then ϕ is surjective if and only if $\tau_0\phi : \tau_0 E \to \tau_0 F$ is surjective.*

Proof. Suppose that ϕ is surjective, so that $F \simeq |E_\bullet|$ where E_n denotes the $(n+1)$-fold fiber power of E over F. Since $\tau_0 : \mathcal{X} \to \tau_0\,\mathcal{X}$ commutes with colimits (it is a left adjoint), we deduce that $\tau_0 F$ is the geometric realization of $\tau_0 E_\bullet$ in the ordinary category $\mathcal{X}_1$. In other words, it is the coequalizer of a pair of morphisms $\tau_0 E_1 \to \tau_0 E_0$. This proves that $\tau_0 E \to \tau_0 F$ is surjective.

For the converse, let us suppose that $\tau_0 E \to \tau_0 F$ is surjective. Choose κ such that $\mathcal{X}$ is κ-coherent and $E, F \in \mathcal{X}_\kappa$. Then E represents a prestack on $\mathcal{X}_\kappa$; let hE denote the corresponding presheaf on $h\,\mathcal{X}_\kappa$, defined by $hE(X) = \pi_0 E(X) = \pi_0 \operatorname{Hom}_{\mathcal{X}}(X, E)$. Similarly let hF denote the presheaf corresponding to F. Then $\tau_0 E$ and $\tau_0 F$ are the sheafifications of these presheaves. The identity map $F \to F$ gives an element of $\eta \in hF(F)$. Since $\tau_0 E$ surjects onto $\tau_0 F$, there exists a covering $\{F_\alpha \to F\}$ and liftings $\eta_\alpha \in hE(F_\alpha)$. Enlarging κ if necessary, we may assume that the sum F' of the F_α is κ-compact, so that η lifts to $\eta' \in \pi_0 \operatorname{Hom}_{\mathcal{X}}(F', E)$ with $F' \to F$ surjective. By Corollary 2.6.6, ϕ is also surjective. □

2.7. **Stacks on a Topos.** In this section, we introduce a construction which passes from ordinary topoi to ∞-topoi, and show that it is right adjoint to the functor $\mathcal{X} \rightsquigarrow \tau_0\,\mathcal{X}$ constructed in §2.5. The construction is in some sense dual to what we did in the last section: instead of starting with an ∞-topos and producing a site, we will begin with a site and produce an ∞-topos of "sheaves of spaces" on that site.

Definition 2.7.1. Let $\mathfrak{X}$ be a topos, and $\mathcal{F}$ a prestack on $\mathfrak{X}$. We will say that $\mathcal{F}$ is a *stack* on $\mathfrak{X}$ if it satisfies the following descent conditions:

- Given any collection $\{C_i\}$ of objects of $\mathfrak{X}$ with (disjoint) sum C, the natural map $\mathcal{F}(C) \to \Pi_i\,\mathcal{F}(C_i)$ is an equivalence.
- Given any surjection $C \to D$ in $\mathfrak{X}$, let $C_\bullet$ denote the simplicial object of $\mathfrak{X}$ with C_n the $(n+1)$-fold fiber power of C over D. Then the natural map $\mathcal{F}(D) \to |\,\mathcal{F}(C_\bullet)|$ is an equivalence. (Here $|\,\mathcal{F}(C_\bullet)|$ denotes the geometric realization of the cosimplicial space).

Lemma 2.7.2. *Let $\mathfrak{X}$ be a topos, and let $\mathfrak{X}_0$ be a full subcategory of $\mathfrak{X}$ which is closed under finite limits and generates $\mathfrak{X}$. Let $f : \mathcal{F} \to \mathcal{G}$ be a map of sheaves on $\mathfrak{X}$. If f induces an equivalence $\mathcal{F}(C) \to \mathcal{G}(C)$ for all $C \in X_0$, then f is an equivalence.*

Proof. Without loss of generality, we may enlarge $\mathfrak{X}_0$ so that $\mathfrak{X}_0$ is closed under sums: $\mathfrak{X}_0$ remains stable under the formation of fiber products, and the value of a sheaf on a sum is determined by its values on the summands.

We must show that for all $C \in \mathfrak{X}$, $\mathcal{F}(C) \to \mathcal{G}(C)$ is an equivalence. By hypothesis, this is true when $C \in \mathfrak{X}_0$. Next, suppose that there is a monomorphism $C \to C'$, with $C' \in \mathfrak{X}_0$. Choose a surjection $U \to C$, where $U \in \mathfrak{X}_0$. Then we have $U \times_C U \times_C \ldots \times_C U = U \times_{C'} U \times_{C'} \ldots \times_{C'} U$, so that all fiber powers of U over C lie in $\mathfrak{X}_0$. Let U_n be the $(n+1)$-fold fiber power of U over C. Then $\mathcal{F}(C) \simeq |\,\mathcal{F}(U_\bullet)|$, and similarly for $\mathcal{G}$. Since f induces homotopy equivalences $\mathcal{F}(U_n) \to \mathcal{G}(U_n)$, it induces a homotopy equivalence between the geometric realizations $\mathcal{F}(C) \to \mathcal{G}(C)$.

Now consider the general case. As before, we choose a surjection $U \to C$, where $U \in \mathfrak{X}_0$. In this case, each fiber power $U \times_C U \times_C \ldots \times_C U$ is a subobject of $U \times U \times \ldots U$ which belongs to $\mathfrak{X}_0$. The above argument shows that $\mathcal{F}(U_n) \to \mathcal{F}'(U_n)$ is an equivalence for each n, and once again we obtain an equivalence $\mathcal{F}(C) \to \mathcal{F}'(C)$. □

Proposition 2.7.3. *Let $\mathfrak{X}$ be a topos. Then the ∞-category of stacks on $\mathfrak{X}$ is an ∞-topos.*

Proof. Let $\mathcal{C}$ be a small category with finite limits and equipped with a Grothendieck topology such that $\mathfrak{X}$ is equivalent to the category of sheaves of sets on $\mathcal{C}$. We will obtain the ∞-category of sheaves on $\mathfrak{X}$ as a left-exact localization of the ∞-category $\mathcal{P} = \mathcal{S}^{\mathcal{C}^{op}}$ of prestacks on $\mathcal{C}$.

Let $\mathcal{F}$ be a prestack on $\mathcal{C}$. We define a new presheaf $\mathcal{F}^+$ as follows: $\mathcal{F}^+(U) = \operatorname{colim}_{\mathfrak{S}} \lim_{S \in \mathfrak{S}} \mathcal{F}(S)$, where $\mathfrak{S}$ ranges over the (directed) collection of sieves in $\mathcal{C}$ which cover U. Note that there is a morphism $\mathcal{F} \to \mathcal{F}^+$ which is natural in $\mathcal{F}$. Consequently we may obtain a transfinite sequence of functors s_α by iterating the construction $\mathcal{F} \mapsto \mathcal{F}^+$. Let $s_0(\mathcal{F}) = \mathcal{F}$, $s_{\alpha+1}(\mathcal{F}) = s_\alpha(\mathcal{F})^+$, and $s_\lambda(\mathcal{F}) = \operatorname{colim}\{s_\alpha(\mathcal{F})\}_{\alpha<\lambda}$ when λ is a limit ordinal. One verifies easily by induction that each s_α is left-exact and accessible.

Let us call a prestack $\mathcal{F}$ on $\mathcal{C}$ a *stack* if, for any sieve $\mathfrak{S}$ in $\mathcal{C}$ covering U, the natural map

$$\mathcal{F}(U) \to \lim_{S \in \mathfrak{S}} \mathcal{F}(S)$$

is an equivalence. It follows easily by induction that for any prestack $\mathcal{F}$ and any stack $\mathcal{F}'$, the natural map

$$\operatorname{Hom}(s_\alpha \mathcal{F}, \mathcal{F}') \to \operatorname{Hom}(\mathcal{F}, \mathcal{F}')$$

is an equivalence for every ordinal α. We also note that for κ sufficiently large, $s_\kappa(\mathcal{F})$ is a stack for *any* prestack $\mathcal{F}$. It follows immediately that s_κ is a localization functor, whose essential image consists of the sheaves on $\mathcal{C}$. Since s_κ is left-exact, we deduce that the ∞-category $\mathcal{X}$ of stacks on $\mathcal{C}$ is an ∞-topos.

By construction, the discrete objects of $\mathcal{X}$ are precisely the sheaves of *sets* on $\mathcal{C}$. Thus $\tau_0 \mathcal{X}$ is equivalent to the topos $\mathfrak{X}$. Let $\mathcal{X}'$ denote the ∞-category of sheaves on $\mathfrak{X}$. The restricted Yoneda embedding induces a functor $F : \mathcal{X} \to \mathcal{X}'$. To complete the proof, it will suffice to show that F is an equivalence.

We next construct a left adjoint G to F. Given a sheaf $\mathcal{F}$ Given a sheaf $\mathcal{F}$ on X and an object $U \in \mathcal{C}$, let $G\,\mathcal{F}(U)$ denote the value of $\mathcal{F}$ on the sheafification of U. One verifies readily that $G : \mathcal{X}' \to \mathcal{X}$ is left adjoint to F. It now suffices to show that the adjunction morphisms $\mathrm{id}_{\mathcal{X}'} \to FG$ and $GF \to \mathrm{id}_{\mathcal{X}}$ are equivalences.

It is clear from the definitions that $GF \to \mathrm{id}_{\mathcal{X}}$ is an equivalence. For the other adjunction, we must show that if $\mathcal{F}$ is a sheaf on X and $\mathcal{F}'$ denotes the induced sheaf $G\,\mathcal{F}$ on $\mathcal{C}$, then $\mathcal{F}(W) \to \operatorname{Hom}_{\mathcal{X}}(W, \mathcal{F}')$ is an equivalence. The right hand side here may be written as a homotopy limit

$$\lim_{U \in \mathcal{C}} \operatorname{Hom}(W(U), \mathcal{F}'(U)) \simeq \lim_{(U,\eta)} \mathcal{F}'(U),$$

where the latter homotopy limit is taken over all pairs $U \in \mathcal{C}$, $\eta \in W(U)$. If W is the sheafification of an object U of $\mathcal{C}$, then this category has a final object and hence the homotopy limit is given by $\mathcal{F}'(U) = \mathcal{F}(W)$, as desired.

Now we have a natural morphism $\mathcal{F} \to FG(\mathcal{F})$, which induces equivalences when evaluated at any W which is the sheafification of a representable presheaf. Since every object of $\mathfrak{X}$ is a colimit of sheafifications of representable presheaves, we may conclude using Lemma 2.7.2. □

If $\mathfrak{X}$ is a topos, we will denote the ∞-topos of stacks on $\mathfrak{X}$ by $\Delta\mathfrak{X}$. We now formulate a universal property enjoyed by $\Delta\mathfrak{X}$:

Proposition 2.7.4. *Let $\mathfrak{X}$ be a topos and $\mathcal{Y}$ an ∞-topos. The category of geometric morphisms $\operatorname{Top}(\tau_0 \mathcal{Y}, \mathfrak{X})$ is naturally equivalent to the ∞-category of geometric morphisms $\operatorname{Top}^\infty(\mathcal{Y}, \Delta\mathfrak{X})$.*

Proof. Let $\mathcal{C}$ be a small category with finite limits, equipped with a Grothendieck topology such that $\mathfrak{X}$ is equivalent to the category of sheaves of sets on $\mathcal{C}$. Then the proof of Proposition 2.7.3 shows that we may identify $\Delta\mathfrak{X}$ with the ∞-category of sheaves on $\mathcal{C}$.

We first assume that the topology on $\mathcal{C}$ is discrete, in the sense that any sieve which covers an object $C \in \mathcal{C}$ actually contains the identity map $C \to C$. In this case, geometric morphisms from $\mathcal{Y}$ into $\Delta\mathfrak{X}$ are just given by left-exact functors $f^* : \mathcal{C} \to \mathcal{Y}$. A similar argument shows that $\mathrm{Top}(\tau_0\,\mathcal{Y}, \mathfrak{X})$ is equivalent to the category of left exact functors $\mathcal{C} \to \tau_0\,\mathcal{Y}$. Note that every object of $\mathcal{C}$ is discrete, and left-exact functors preserve discrete objects by Lemma 2.5.5. Therefore every left exact functor $\mathcal{C} \to \mathcal{Y}$ factors uniquely through $\tau_0\,\mathcal{Y}$, which completes the proof in this case.

Now we consider the general case. Let $\mathcal{P}$ denote the ∞-category of presheaves on $\mathcal{C}$, so that $\Delta\mathfrak{X}$ is a left-exact localization of $\mathcal{P}$ and $\mathfrak{X}$ is a left-exact localization of $\tau_0\,\mathcal{P}$. We will denote both localization functors by L. The case that we have just treated shows that $\mathrm{Top}(\tau_0\,\mathcal{Y}, \tau_0\,\mathcal{P}) \simeq \mathrm{Top}(\mathcal{Y}, \mathcal{P})$. On the other hand, $\mathrm{Top}(\mathcal{Y}, \Delta\mathfrak{X})$ can be identified with the full subcategory of $\mathrm{Top}(\mathcal{Y}, \mathcal{P})$ consisting of left-exact, colimit preserving functors $f^* : \mathcal{P} \to \mathcal{Y}$ which factor through L, in the sense that the natural morphism $f^*(E) \to f^*(LE)$ is an equivalence for all $E \in \mathcal{P}$. Similarly, $\mathrm{Top}(\tau_0\,\mathcal{Y}, \mathfrak{X})$ can be identified with the collection of all left-exact, colimit preserving functors $f^* : \tau_0\,\mathcal{P} \to \tau_0\,\mathcal{Y}$ which factor through L. To complete the proof, it suffices to show that a geometric morphism $f^* : \mathcal{P} \to \mathcal{Y}$ factors through L if and only if the restriction $f^*|\tau_0\,\mathcal{P}$ factors through L.

The "only if" part is trivial, so let us focus on the "if". Let S denote the collection of all morphisms of the form $U \to LU$, where $U \in \tau_0\,\mathcal{P}$. Since every morphism of S becomes invertible after applying f^*, there exists a factorization $\mathcal{P} \to S^{-1}\,\mathcal{P} \to \mathcal{Y}$. It now suffices to verify that $S^{-1}\,\mathcal{P} \to \mathcal{Y}$ is an equivalence. In other words, we must show that every S-local object of $\mathcal{P}$ is actual L-local. This follows immediately from the definition of L. □

Remark 2.7.5. It is possible to give a definition of n-topos for any $0 \leq n \leq \infty$: this is a particular kind of n-category, which "looks like" the n-category of $(n-1)$-truncated stacks on a topological space X. For $n = \infty$, one recovers Definition 2.2.5. For $n = 1$, one recovers the classical definition of a topos. For $n = 0$, one recovers the notion of a locale.

For each $n \geq m$, one has a "truncation" functor from n-topoi to m-topoi, which replaces an n-topos by its full subcategory of m-truncated objects. Each of these truncation functors has a right adjoint. We have discussed the situation here when $n = \infty$, $m = 1$. When $n = 1$, $m = 0$, the "truncation" replaces a topos by the locale consisting of its "open subsets" (that is, subobjects of the final object).

2.8. **Homotopy Groups.** In this section, we will discuss the homotopy groups of objects and morphisms in an ∞-topos. These will be needed in §2.9 on hyperdescent and in §4 on dimension theory, but are not needed for the main result of §3.

Let $\mathcal{X}$ be an ∞-topos, and let X be an object of $\mathcal{X}$. Fixing a base point $* \in S^n$, we get an "evaluation" map $X^{S^n} \to X$. Let $\pi_n(X) = \tau_0^{\mathcal{X}_{/X}}(X^{S^n})$, where the notation indicates that we consider X^{S^n} as an object of the ∞-topos $\mathcal{X}_{/X}$. This is a sheaf of pointed sets on X: the base point is induced by the retraction of S^n onto $*$. As in the classical case, the co-H-space structure on the sphere gives rise to a group structure on the sheaf $\pi_n(X)$ if $n \geq 1$, which is abelian if $n \geq 2$.

In order to work effectively with homotopy sets, it is convenient to define the homotopy sets $\pi_n(f)$ of a morphism $f : X \to Y$ to be the homotopy sets of X, considered via f as an object of $\mathcal{X}_{/Y}$. This construction gives again sheaves of pointed sets on X (groups if $n \geq 1$, abelian groups if $n \geq 2$). The intuition is that the stalk of these sheaves at a "point" $p \in X$ is the nth homotopy group of the mapping fiber of f, with base point p.

It will be useful to have the following recursive definition of homotopy groups. Regarding X as an object of the topos $\mathcal{X}_{/Y}$, we may take its 0th truncation $\tau_0^{\mathcal{X}_{/Y}}X$. This is a sheaf of sets on $\mathcal{X}_{/Y}$; by definition we set $\pi_0(f) = f^*\tau_0^{\mathcal{X}_{/Y}}(X) = X \times_Y \tau_0^{\mathcal{X}_{/Y}}(X)$. The natural map $X \to \tau_0^{\mathcal{X}_{/Y}}(X)$ gives a global section of $\pi_0(f)$. Note that in this case, $\pi_0(f)$ is the pullback of a sheaf on Y: this is because the definition of π_0 does not require a base point. If $n \geq 0$, then $\pi_n(f) = \pi_{n-1}(f')$, where $f' : X \to X \times_Y X$ is the diagonal map. Finally, in the case where $f = X \to 1$, we let $\pi_n(X) = \pi_n(f)$.

Remark 2.8.1. Let $f : \mathcal{X} \to \mathcal{Y}$ be a geometric morphism of ∞-topoi, and let $g : Y \to Y'$ be a morphism in $\mathcal{Y}$. Then $f^*(\pi_n(g)) \simeq \pi_n(f^*(g))$. This follows immediately from the definition and Proposition 2.5.10.

Remark 2.8.2. Given a pair of composable morphisms $X \xrightarrow{f} Y \xrightarrow{g} Z$, we have a sequence of pointed sheaves

$$\ldots \to f^*\pi_{i+1}(g) \to \pi_i(f) \to \pi_i(g \circ f) \to f^*\pi_i(g) \to \pi_{i-1}(f) \to \ldots$$

with the usual exactness properties. To see this, one first mimics the usual construction to produce the boundary maps in the above sequence. To prove exactness, one can embed the ∞-topos $\mathcal{X}$ in an ∞-topos $\mathcal{P}$ of presheaves. Using Remark 2.8.1, we reduce to the problem of showing that the corresponding sequence is exact in $\mathcal{P}$. This can be checked componentwise. We are thereby reduced to the case $\mathcal{X} = \mathcal{S}$, which is classical.

Remark 2.8.3. If $\mathcal{X} = \mathcal{S}$, and $\eta : 1 \to X$ is a pointed space, then $\eta^*\pi_n(X)$ is the nth homotopy group of X with base point η.

We now study the implications of the vanishing of homotopy groups.

Proposition 2.8.4. *Let $f : X \to Y$ be an n-truncated morphism. Then $\pi_k(f) = *$ for all $k > n$. If $n \geq 0$ and $\pi_n(f) = *$, then f is $(n-1)$-truncated.*

Proof. The proof goes by induction on n. If $n = -2$, then f is an equivalence and there is nothing to prove. Otherwise, we know that $f' : X \to X \times_Y X$ is $(n-1)$-truncated. The inductive hypothesis then allows us to infer that $\pi_k(f) = \pi_{k-1}(f') = *$ whenever $k > n$ and $k > 0$. Similarly, if $n \geq 1$ and $\pi_n(f) = \pi_{n-1}(f') = *$, then f' is $(n-2)$-truncated by the inductive hypothesis, so that f is $(n-1)$-truncated.

The case of small k and n requires special attention: we must show that if f is 0-truncated, then f is (-1)-truncated if and only if $\pi_0(f) = *$. The fact that f is 0-truncated implies that $\tau_0^{\mathcal{X}_Y} X = X$, so that $\pi_0(f) = X \times_Y X$. To say $\pi_0(f) = *$ is to assert that the map $f' : X \to \pi_0(f)$ is an equivalence, which is to say that X is (-1)-truncated. $\square$

Remark 2.8.5. The Proposition 2.8.4 implies that if f is n-truncated for $n \gg 0$, then we can test whether or not f is n-truncated for any particular value of n by computing the homotopy groups of f. In contrast to the classical situation, it is not always possible to drop the assumption that f is n-truncated for $n \gg 0$.

Lemma 2.8.6. *Let X be an object in an ∞-topos $\mathcal{X}$. Then the natural map $p : X \to \tau_n X$ induces isomorphisms $\pi_k(X) \to p^*\pi_k(\tau_n X)$ for all $k \leq n$.*

Proof. Let $\phi : \mathcal{X} \to \mathcal{Y}$ be a geometric morphism such that ϕ_* is fully faithful. By Proposition 2.5.10 and Remark 2.8.1, it will suffice to prove the lemma in the case where $\mathcal{X} = \mathcal{Y}$. By Proposition 2.4.1, we may assume that $\mathcal{Y}$ is an ∞-category of presheaves. In this case, homotopy groups and truncations are computed pointwise. Thus we may reduce to the case $\mathcal{X} = \mathcal{S}$, where the conclusion is evident. $\square$

Definition 2.8.7. Let $f : X \to Y$ be a morphism in an ∞-topos $\mathcal{X}$. We shall say that f is (-1)-*connected* if f is surjective. If $0 \leq n \leq \infty$, we shall say that f is n-*connected* if it is surjective and $\pi_k(f) = *$ for $k \leq n$. We shall say that the object X is n-*connected* if $f : X \to 1$ is n-connected.

Proposition 2.8.8. *Let X be an object in an ∞-topos $\mathcal{X}$. Then X is n-connected if and only if $\tau_n X \to 1$ is an equivalence.*

Proof. The proof goes by induction on $n \geq -1$. If $n = -1$, then the conclusion holds by definition. Suppose $n \geq 0$. Let $p : X \to \tau_n X$ denote the natural map. If $\tau_n X \simeq 1$, then $\pi_k X = p^*\pi_k(\tau_n X) = p^*(\pi_k(1)) = *$ for $k \leq n$ by Lemma 2.8.6. Conversely, suppose that X is n-connected. Then $p^*\pi_n(\tau_n X) = *$. Since p is surjective, we deduce that $\pi_n(\tau_n X) = *$. This implies that $\tau_n X$ is $(n-1)$-truncated, so that $\tau_n X \simeq \tau_{n-1} X$. Repeating this argument, we reduce to the case where $n = -1$ which was handled above. $\square$

Since $\mathrm{Hom}_{\mathcal{X}}(X, Y) \simeq \mathrm{Hom}_{\mathcal{X}}(\tau_n X, Y)$ when Y is n-truncated, we deduce that X is n-connected if and only if $\mathrm{Hom}_{\mathcal{X}}(1, Y) \to \mathrm{Hom}_{\mathcal{X}}(X, Y)$ is an equivalence for all n-truncated Y. From this, we can immediately deduce the following relative version of Proposition 2.8.8:

Corollary 2.8.9. *Let $f : X \to X'$ be a morphism in an ∞-topos $\mathcal{X}$. Then f is n-connected if and only if $\mathrm{Hom}_{X'}(X', Y) \to \mathrm{Hom}_{X'}(X, Y)$ is an equivalence, for any n-truncated $Y \to X'$.*

Corollary 2.8.10. *The class of n-connected morphisms in an ∞-topos is stable under cobase extension.*

Proof. This follows from the preceding corollary and the fact that n-truncated morphisms are stable under base extension. □

A morphism $f : X \to X'$ is ∞-connected if and only if it is n-connected for all n. By the preceding corollary, this holds if and only if $\mathrm{Hom}_{X'}(X', Y) \simeq \mathrm{Hom}_{X'}(X, Y)$ for any morphism $Y \to X'$ which is n-truncated for some n. In $\mathcal{S}$, this implies that f is an equivalence, but this is not true in general. We will investigate this issue further in the next section.

We conclude by noting the following stability properties of the class of n-connected morphisms:

Proposition 2.8.11. *Let $\mathcal{X}$ be an ∞-topos.*

(1) *Any n-connected morphism of $\mathcal{X}$ is m-connected for any $m \leq n$.*
(2) *Any equivalence is ∞-connected.*
(3) *The class of n-connected morphisms is closed under composition.*
(4) *Let $f : X \to Y$ be an n-connected morphism and let $Y' \to Y$ be any map. Then the induced map $f' : X' = X \times_Y Y' \to Y'$ is n-connected. The converse holds if $Y' \to Y$ is surjective.*

Proof. The first two claims are obvious. The third follows from the long exact sequence of Remark 2.8.2. The first assertion of (4) follows from the fact that homotopy groups are compatible with base change (a special case of Remark 2.8.1). The second assertion of (4) follows from Remark 2.8.1 and descent for equivalences ($\pi_k(f') \simeq Y'$ if and only if $\pi_k(f) \simeq Y$). □

2.9. **Hyperstacks.** If $\mathcal{C}$ is a small category equipped with a Grothendieck topology, then Jardine (see [14]) shows that the category of simplicial presheaves on $\mathcal{C}$ is equipped with the structure of a simplicial model category, which gives rise to an ∞-category. On the other hand, we have associated to $\mathcal{C}$ an ∞-category $\mathcal{X}$ of "stacks on $\mathcal{C}$". These two constructions are not equivalent: the weak equivalences in Jardine's model structure are the "local homotopy equivalences", which correspond to the ∞-connected morphisms in our setting. However, it turns out the Joyal-Jardine theory can be obtained from ours by inverting the ∞-connected morphisms. The resulting localization functor is left exact, so that the ∞-category underlying Jardine's model structure is also an ∞-topos.

Let us begin by studying the ∞-connected morphisms:

Proposition 2.9.1. *Let $\mathcal{X}$ be an ∞-topos, and let S denote the collection of ∞-connected morphisms of $\mathcal{X}$. Then S is saturated.*

Proof. This follows immediately from the description of ∞-connected morphisms given in Corollary 2.8.9. □

In order to construct a localization of $\mathcal{X}$ which inverts the ∞-connected morphisms, we will need to show that the class of ∞-connected morphisms is setwise generated. This follows immediately from the following:

Proposition 2.9.2. *Let $\mathcal{X}$ be an ∞-topos. Then there exists a cardinal κ such that any ∞-connected morphism $f : X \to Y$ may be written as a κ-filtered colimit of ∞-connected morphisms $f_\alpha : X_\alpha \to Y_\alpha$ between κ-compact objects.*

Proof. The proof is routine cardinality estimation, and may be safely omitted by the reader. We choose $\kappa_0 > \omega$ large enough that $\mathcal{X}$ is κ_0-coherent, the functor τ_0 is κ-continuous and preserves κ_0-compact objects. Since homotopy sets are constructed using finite limits and τ_0, it follows that the formation of homotopy sets in $\mathcal{X}$ preserves κ_0-compactness and commutes with κ_0-filtered colimits. Now choose $\kappa \gg \kappa_0$.

Since $\mathcal{X}$ is κ-coherent, we may write f as a κ-filtered colimit of morphisms $f_\alpha : X_\alpha \to Y_\alpha$ where each X_α and each Y_α is κ-compact. For simplicity, let us assume that the colimit is indexed by a partially ordered set $\mathcal{I}$. Let $\mathcal{P}$ denote the collection of all κ_0-filtered, κ-small subsets of $\mathcal{I}$ (ordered by inclusion). For each $S \in \mathcal{P}$, we let $f_S : X_S \to Y_S$ denote the colimit indexed by S. Let $\mathcal{P}_0 \subseteq \mathcal{P}$ denote the collection of all subsets S such that f_S is ∞-connected. It is easy to see that $\mathcal{P}$ is κ-filtered, and the colimit of the f_S is f. To complete the proof, it suffices to show that $\mathcal{P}_0$ is cofinal in $\mathcal{P}$.

Given any $\alpha \in \mathcal{I}$ and any $k \geq 0$, we know that $\pi_k(f_\alpha) \to \pi_k(f)$ factors through X (since $\pi_k(f) \simeq X$). Since the source is κ-compact, we deduce that there exists $\beta > \alpha$ such that $\pi_i(f_\alpha) \to \pi_i(f_\beta)$ factors through X_β.

Given $S \in \mathcal{P}$, we define a transfinite sequence of enlargements of S as follows. Let $S_0 = S$, let $S_\lambda = \bigcup_{\gamma<\lambda} S_\gamma$ when λ is a limit ordinal. Finally let $S_{\lambda+1}$ be some κ_0-filtered, κ-small subset of $\mathcal{I}$ containing S_i with the property that for any $\alpha \in S_i$ and any $k \geq 0$, there exists $\beta \in S_{i+1}$ such that $\beta \geq \alpha$ and $\pi_k(f_\alpha) \to \pi_k(f_\beta)$ factors through the base point. It is clear that S_λ is well-defined for $\lambda < \kappa$, and that S_{κ_0} is κ-filtered. Since the formation of homotopy sets commutes with κ-filtered colimits, we deduce that $f_{S_{\kappa_0}}$ is ∞-connected, so that $S_{\kappa_0} \in \mathcal{P}_0$ as desired. □

Let $\mathcal{X}$ be an ∞-topos, and let S be the class of ∞-connected morphisms of $\mathcal{X}$. Propositions 2.9.1 and 2.9.2 imply that S is saturated and setwise generated, so that we may construct a localization functor $L : \mathcal{X} \to \mathcal{X}$ which inverts precisely the morphisms in S.

Proposition 2.9.3. *Let $\mathcal{C}$ be a presentable ∞-category, and let S denote a setwise generated, saturated collection of morphisms of $\mathcal{C}$. Let $L : \mathcal{C} \to \mathcal{C}$ denote the corresponding localization functor. The following are equivalent:*

(1) *The class of morphisms S is stable under base change.*
(2) *The localization functor L is left exact.*

Proof. Since S is precisely the collection of morphisms f such that Lf is an equivalence, it is immediate that (2) implies (1). Let us therefore assume (1). Since the final object $1 \in \mathcal{X}$ is obviously S-local, we have $L1 \simeq 1$. Thus it will suffice to show that L commutes with pullbacks. In other words, we must show that $f : X \times_Y Z \to LX \times_{LY} LZ$ is an S-localization for any pair of objects X, Z over Y in $\mathcal{X}$. Since $LX \times_{LY} LZ$ is clearly S-local, it suffices to prove that f belongs to S. Now f is a composite of maps

$$X \times_Y Z \to X \times_{LY} Z \to LX \times_{LY} Z \to LX \times_{LY} LZ.$$

The last two maps are obtained from $X \to LX$ and $Z \to LZ$ by base change. It follows that they belong to S. Thus, it will suffice to show that $f' : X \times_Y Z \to X \times_{LY} Z$ belongs to S. Since this map is obtained from $Y \times_Y Y \to Y \times_{LY} Y$ by a base-change, it suffices to prove that $f'' : Y \to Y \times_{LY} Y$ belongs to S. Projection to the first factor gives a section $s : Y \times_{LY} Y \to Y$ of f'', so it suffices to prove that s belongs to S. But s is a base change of the morphism $Y \to LY$. □

Since the formation of homotopy sheaves is compatible with base change, we see that the class S of ∞-connected morphisms satisfies (1) of Proposition 2.9.3. It follows from Theorem 2.4.1 that the collection of S-local objects of $\mathcal{X}$ forms an ∞-topos, which we shall denote by $\mathcal{X}^{\mathrm{hyp}}$. In a moment we shall characterize $\mathcal{X}^{\mathrm{hyp}}$ by a universal property. First we need a simple lemma.

Lemma 2.9.4. *Let $f : X \to Y$ be a morphism in $\mathcal{X}^{hyp}$. We may form $\pi_n(f)$ in $\mathcal{X}$. Then $\pi_n(f)$ lies in $\mathcal{X}^{hyp}$, and is the nth homotopy sheaf of f in $\mathcal{X}^{hyp}$.*

Proof. Homotopy sheaves are constructed using finite limits and τ_0. Since $\mathcal{X}^{\mathrm{hyp}}$ is stable under finite limits, it suffices to show that for any $g : E \to E'$ in $\mathcal{X}^{\mathrm{hyp}}$, we have $\tau_0^{\mathcal{X}_{E'}} E \in \mathcal{X}^{\mathrm{hyp}}$. This follows from the fact that $\tau_0^{\mathcal{X}_{E'}}$ is discrete over E' and E' is S-local. □

Following [17], we shall say that an ∞-topos $\mathcal{X}$ *t-complete* if every ∞-connected morphism of $\mathcal{X}$ is an equivalence.

Lemma 2.9.5. *The ∞-topos $\mathcal{X}^{hyp}$ is t-complete.*

Proof. Since the formation of homotopy sheaves in $\mathcal{X}^{\mathrm{hyp}}$ is compatible with the formation of homotopy sheaves in $\mathcal{X}$, any morphism $f : X \to Y$ which is ∞-connected in $\mathcal{X}^{\mathrm{hyp}}$ is also ∞-connected in $\mathcal{X}$, and therefore an equivalence in $\mathcal{X}^{\mathrm{hyp}}$. Thus $\mathcal{X}^{\mathrm{hyp}}$ is t-complete. □

We can give the following universal characterization of $\mathcal{X}^{\mathrm{hyp}}$.

Proposition 2.9.6. *Let $\mathcal{X}$ be an ∞-topos and $\mathcal{Y}$ an ∞-topos which is t-complete. Then composition with the natural geometric morphism $\mathcal{X}^{hyp} \to \mathcal{X}$ induces an equivalence of ∞-categories*

$$\mathrm{Top}^{\infty}(\mathcal{Y}, \mathcal{X}^{hyp}) \to \mathrm{Top}^{\infty}(\mathcal{Y}, \mathcal{X}).$$

Proof. In other words, we must show that if $f : \mathcal{Y} \to \mathcal{X}$ is a geometric morphism, then f^* factors through L: that is, it carries ∞-connected morphisms into equivalences. Since $\mathcal{Y}$ is t-complete, it suffices to show that f^* preserves ∞-connected morphisms. This follows immediately from the fact that f^* commutes with the formation of homotopy sheaves. □

The objects of $\mathcal{X}^{\mathrm{hyp}}$ are the S-local objects of $\mathcal{X}$. It is of interest to describe these objects explicitly. This description was given in [27] in the case where $\mathcal{X}$ is a left exact localization of a category of prestacks on an ordinary category, and in [17] in general. We will summarize their results here. What follows will not be needed later in this paper and may be safely omitted by the reader.

We will need to employ a bit more terminology concerning simplicial objects in an ∞-category. Let Δ denote the category of combinatorial simplices, and let $\Delta_{\leq n}$ denote the full subcategory consisting of combinatorial simplices of dimension $\leq n$. If $\mathcal{C}$ is any ∞-category, then a simplicial object of $\mathcal{C}$ is a contravariant functor $E_\bullet : \Delta \to \mathcal{C}$. This induces a functor $\Delta_{\leq n} \to \mathcal{C}$ by restriction, called the *skeleton* of $E_\bullet$. If $\mathcal{C}$ admits finite limits, then the skeleton functor $\mathcal{C}^{\Delta^{op}} \to \mathcal{C}^{\Delta^{op}_{\leq n}}$ has a right adjoint, called the n-coskeleton. If $E_\bullet$ is a simplicial object of $\mathcal{C}$, we will let cosk^n denote the n-coskeleton of the n-skeleton of E; this is a new simplicial object equipped with a map $E_\bullet \to \mathrm{cosk}^n_\bullet(E_\bullet)$ which is an equivalence on the n-skeleton (and is universal with respect to this property). We say that $E_\bullet$ is *n-coskeletal* if $E_\bullet \to \mathrm{cosk}^n_\bullet(E_\bullet)$ is an equivalence.

Definition 2.9.7. Let $\mathcal{X}$ be an ∞-topos, and let $E_\bullet$ denote a simplicial object of $\mathcal{X}$. We shall say that $E_\bullet$ is a *hypercovering* if, for each $n \geq 0$, the natural map $E_n \to \mathrm{cosk}^{n-1}_n(E_\bullet)$ is surjective.

More generally, if $E_\bullet \to E$ is an augmented simplicial object, we shall say that $E_\bullet$ is a *hypercovering of* E if $E_\bullet$ is a hypercovering in the ∞-topos $\mathcal{X}_{/E}$.

Theorem 2.9.8. [17], [27] *Let $\mathcal{X}$ be an ∞-topos, let S denote the class of ∞-connected morphisms of $\mathcal{X}$, and let S' denote the class of morphisms having the form $|E_\bullet| \to E$, where $E_\bullet$ is a hypercovering of E in $\mathcal{X}$. An object $E' \in \mathcal{X}$ is S-local if and only if it is S'-local.*

Proof. Suppose first that E' is S'-local. Let $f : X \to Y$ be an ∞-connected morphism. Regard X as a constant simplicial object over Y. If we can show that X is a hypercovering, then $X \simeq |X| \to Y$ induces an equivalence $\mathrm{Hom}_{\mathcal{X}}(Y, E') \to \mathrm{Hom}_{\mathcal{X}}(X, E')$ and we are done. In other words, we need to show that for each $n \geq 0$, the natural map $X \to \mathrm{cosk}^{n-1}_n(X)$ is surjective (here the coskeleton is formed in $\mathcal{X}_{/Y}$). This is precisely the condition that $\pi_n(f)$ vanishes.

To prove the converse, we show that if $E_\bullet \to E$ is a hypercovering then $|E_\bullet| \to E$ is ∞-connected. Without loss of generality, we may replace $\mathcal{X}$ by $\mathcal{X}_E$ and thereby assume that $E = 1$. We need to prove that $|E_\bullet|$ is n-connected for every n. By Corollary 2.8.9, it suffices to show that $\mathrm{Hom}_{\mathcal{X}}(1, Y) \to \mathrm{Hom}_{\mathcal{X}}(|E_\bullet|, Y)$ is an equivalence whenever Y is n-truncated. In this case, $\mathrm{Hom}_{\mathcal{X}}(|E_\bullet|, Y)$ does not depend on E_k for $k > n+1$. Thus we may replace $E_\bullet$ by its $(n+1)$-coskeleton without loss of generality. One then shows, by induction on k, that $|\mathrm{cosk}^n_\bullet(E_\bullet)| \to E$ is an equivalence, using the fact that $E_\bullet$ is a hypercovering and the fact that all groupoids are effective in $\mathcal{X}$. Taking $k = n+1$, we obtain the desired result. □

In other words, the S-local objects of $\mathcal{X}$ are precisely the objects which *satisfy hyperdescent.*

2.10. **Stacks Versus Hyperstacks.** If $\mathfrak{X}$ is a Grothendieck topos, then we have seen how to construct two potentially different ∞-topoi associated to $\mathfrak{X}$. First, we have the ∞-topos $\mathcal{X} = \Delta\mathfrak{X}$ consisting of stacks on X. Second, we can form the ∞-topos $\mathcal{X}^{\mathrm{hyp}}$ consisting of stacks on X satisfying hyperdescent, which is a localization of $\mathcal{X}$. It seems that the majority of the literature is concerned with $\mathcal{X}^{\mathrm{hyp}}$, while $\mathcal{X}$ itself has received less attention (although there is some discussion in [27]). We would like to make the case that $\mathcal{X}$ is the more natural choice. Proposition 2.7.4 provides a formal justification for this claim: the functor which constructs $\mathcal{X}$ from $\mathfrak{X}$ is the adjoint of the forgetful functor from ∞-topoi to topoi. In this section, we would like to summarize some less formal reasons why it may be nicer to work with $\mathcal{X}$.

- Suppose that $\pi : X \to S$ and $\psi : S' \to S$ denote continuous maps of locally compact topological spaces, and let $X' = X \times_S S'$ with projections $\pi' : X' \to S'$ and $\psi' : X' \to X$. In classical sheaf theory, one has a natural transformation $\psi^* \pi_* \to \pi'_* \psi'^*$ of functors between the derived categories of left-bounded complexes of sheaves on X and on S'. The proper base change theorem asserts that this transformation is an equivalence whenever the map π is proper.

  The functors $\psi^* \pi_*$ and $\pi'_* \psi'^*$ may also be defined on ∞-categories of stacks and ∞-categories of hyperstacks, and one again has a base-change transformation as above. It is natural to ask if the base change transformation is an equivalence when π is proper. It turns out that this is *true* for the ∞-categories of stacks, but *false* for the ∞-categories of hyperstacks:

  Counterexample 2.10.1. Let Q denote the Hilbert cube $[0,1] \times [0,1] \times \ldots$, and let $\pi : X \to [0,1]$ be the projection onto the first factor. For each i, we let Q_i denote "all but the first i" factors of Q, so that $Q = [0,1]^i \times Q_i$.

  We construct a stack $\mathcal{F}$ on $X = Q \times [0,1]$ as follows. Begin with the empty stack. Adjoin to it two sections, defined over the open sets $[0,1) \times Q_1 \times [0,1)$ and $(0,1] \times Q_1 \times [0,1)$. These sections both restrict to give sections of $\mathcal{F}$ over the open set $(0,1) \times Q_1 \times [0,1)$. We next adjoin paths between these sections, defined over the smaller open sets $(0,1) \times [0,1) \times Q_2 \times [0, \frac{1}{2})$ and $(0,1) \times (0,1] \times Q_2 \times [0, \frac{1}{2})$. These paths are both defined on the smaller open set $(0,1) \times (0,1) \times Q_2 \times [0, \frac{1}{2})$, so we next adjoin two homotopies between these paths over the open sets $(0,1) \times (0,1) \times [0,1) \times Q_3 \times [0, \frac{1}{3})$ and $(0,1) \times (0,1) \times (0,1] \times Q_3 \times [0, \frac{1}{3})$. Continuing in this way, we obtain a stack $\mathcal{F}$. On the closed subset $Q \times \{0\} \subset X$, the stack $\mathcal{F}$ is ∞-connected by construction, and therefore the associate hyperstack admits a global section. However, the hyperstack associated to $\mathcal{F}$ does not admit a global section in any neighborhood of $Q \times \{0\}$, since such a neighborhood must contain $Q \times [0, \frac{1}{n})$ for $n \gg 0$ and the higher homotopies required for the construction of a section are eventually not globally defined.

  The same issue arises in classical sheaf theory if one wishes to work with unbounded complexes. In [23], Spaltenstein defines a derived category of unbounded complexes of sheaves on X, where X is a topological space. His definition forces all quasi-isomorphisms to become invertible, which is analogous to procedure of obtaining $\mathcal{X}^{\text{hyp}}$ from $\mathcal{X}$ by inverting the ∞-connected morphisms. Spaltenstein's work shows that one can extend the *definitions* of all of the basic objects and functors. However, it turns out that the *theorems* do not all extend: in particular, one does not have the proper base change theorem in Spaltenstein's setting (Counterexample 2.10.1 can easily be adapted to the setting of complexes of abelian sheaves). The problem may be rectified by imposing weaker descent conditions, which do not invert all quasi-isomorphisms.

  The proof of the proper base change theorem in the context of ∞-topoi will be given in a sequel to this paper.
- The ∞-topos $\mathcal{X}$ has better finiteness properties than $\mathcal{X}^{\text{hyp}}$. Suppose that X is a coherent topological space (that is, X has a basis of compact open sets which is stable under the formation of finite intersections), and let $\mathcal{X}$ denote the ∞-category of stacks on X (in other words, stacks on the ordinary topos of sheaves on X). We may construct the ∞-category $\mathcal{X}$ as a localization of the ∞-category $\mathcal{P}$ of prestacks on $\mathcal{C}$, where $\mathcal{C}$ denotes the partially ordered set of compact open subsets of X. The localization functor L is given by transfinitely iterating the construction $\mathcal{F} \mapsto \mathcal{F}^+$ of Proposition 2.7.3. Since every covering in X_ω may be refined to a finite covering, one can show that $\mathcal{F} \mapsto \mathcal{F}^+$ commutes with filtered colimits in $\mathcal{F}$. It follows easily that L commutes with filtered colimits. From this, we may deduce that L carries compact objects of $\mathcal{P}$ into compact objects of $\mathcal{X}$. Since $\mathcal{P}$ is generated by compact objects, we deduce that $\mathcal{X}$ is generated by compact objects.

  By contrast, one can give an example of a coherent topological space for which the "hypersheafification" of prestacks does not commute with filtered colimits. The construction is modeled on Counterexample 2.10.1, replacing each copy of the interval $[0,1]$ with a finite topological space having a generic point (analogous to the interior $(0,1)$) and two special points (analogous the the endpoints 0 and 1). The compatibility of sheafification with filtered colimits is a crucial feature of coherent topoi, which might be lost in the ∞-categorical setting if $\mathcal{X}^{\text{hyp}}$ is used.

Remark 2.10.2. In particular, we note that $\mathcal{X} \neq \mathcal{X}^{\mathrm{hyp}}$ when $\mathfrak{X}$ is a coherent topos, so that $\mathcal{X}$ does not necessarily have enough points. Thus, one does not expect an ∞-categorical version of the Gödel-Deligne theorem to hold.

- Although $\mathcal{X}$ and $\mathcal{X}^{\mathrm{hyp}}$ are not equivalent in general, they do coincide whenever certain finite-dimensionality conditions are satisfied (see Corollary 4.1.6). These conditions are satisfied in most of the situations in which Jardine's model structure is commonly used, such as the Nisnevich topology on a scheme of finite Krull dimension.
- Let X be a paracompact space, $\mathcal{X}$ the ∞-topos of stacks on X, K a homotopy type, and $p : \mathcal{X} \to \mathcal{S}$ the natural geometric morphism. Then $\pi_0 p_* p^* K$ is a natural definition of the "sheaf cohomology" of X with coefficients in K, and this agrees with the definition $[X, K]$ given by classical homotopy theory (Theorem 3.0.5). As we mentioned in the introduction, this fails if we replace $\mathcal{X}$ by $\mathcal{X}^{\mathrm{hyp}}$.
- If the topos $\mathfrak{X}$ has enough points, then the local equivalences of $\mathcal{X}$ are precisely the morphisms which induce weak equivalences on each stalk. It follows immediately that $\mathcal{X}^{\mathrm{hyp}}$ has enough points, and that if $\mathcal{X}$ has enough points then $\mathcal{X} \simeq \mathcal{X}^{\mathrm{hyp}}$. The possible failure of the Whitehead theorem in $\mathcal{X}$ may be viewed either as a bug or a feature. While this failure has the annoying consequence that $\mathcal{X}$ may not have "enough points" even though $\mathfrak{X}$ does, it also means that $\mathcal{X}$ might detect certain global phenomena which cannot be properly understood by restricting to points. Let us consider an example from classical geometric topology. If $f : X \to Y$ is a continuous map between compact ANRs, then f is said to be *cell-like* if for any $y \in Y$, the inverse image $f^{-1}(y)$ is contractible in arbitrarily small neighborhoods of itself. In the finite dimensional case, this holds if and only if each fiber $f^{-1}(y)$ has trivial shape, but this fails in general.

Let us indicate how the notion of an ∞-topos can shed some light on the situation. First, we note the following:

Proposition 2.10.3. *Let $f : X \to Y$ be a map of compact ANRs, and let f also denote the associated geometric morphism from $\mathcal{X}$ to $\mathcal{Y}$, where $\mathcal{X}$ and $\mathcal{Y}$ denote the ∞-topoi of stacks on X and Y, respectively. The following conditions are equivalent:*

(1) *The map f is cell-like.*
(2) *For every $E \in \mathcal{Y}$, the adjunction morphism $E \to f_* f^* E$ is an equivalence.*
(3) *The functor f^* is fully faithful.*

Proof. A proof will be given in a sequel to this paper. □

Since f is proper, the stalk of $f_* f^* E$ at a point y is equal to the cohomology of $f^{-1}(y)$ with coefficients in E_y (by the proper base change theorem promised above), which is equivalent to E_y if $f^{-1}(y)$ has trivial shape (essentially by the definition of the strong shape). It follows that if all of the fibers of f have trivial shape, then the adjunction map $E \to f_* f^* E$ induces an equivalence on all stalks, and is therefore ∞-connected, for any $E \in \mathcal{Y}$. If Y is finite dimensional, we can conclude that f is cell-like, but not in general.

It is crucial that we work with stacks, rather than hyperstacks, in the above discussion. If we chose instead to work with hyperstacks, then we could take Proposition 2.10.4 as the definition of a concept analogous to, but distinct from, that of a cell-like map. However, this analogue is of more limited geometric interest.

3. Paracompact Spaces

Let X be a topological space, and let $p : X \to *$ be the projection. Let K be a homotopy type. Then we may regard K as a stack on $*$, and form the stack $p_* p^* K$. Alternatively, we can form an actual topological space $|K|$ (which we assume to be nice, say a metric ANR) and study the collection of homotopy classes of maps $[X, |K|]$. The goal of this section is to prove the following:

Theorem 3.0.4. *If X is paracompact, then there is a bijection*

$$\phi : [X, |K|] \to \pi_0 p_* p^* K.$$

In fact, the map ϕ always exists and is natural in both K and X; only the bijectivity of ϕ requires X to be paracompact.

Let us outline the proof of Theorem 3.0.5. In order to do calculations with stacks on X, we will need to introduce some models for the stacks. It is most common to use simplicial presheaves on X, but this is inconvenient for two reasons. The first is that general open subsets of X can behave badly; we will therefore restrict our attention to presheaves defined with respect to a nice basis for the topology of X. The second problem is more subtle, but for technical reasons it will be convenient for us to introduce a category $\mathcal{K}$ of Ind-compact metric spaces for use as a model of homotopy theory, rather than simplicial sets. We do not require much from the category $\mathcal{K}$: it does not need to be a Quillen model category, for example. The main properties of $\mathcal{K}$ which we shall need are as follows:

- There is a good formal relationship between $\mathcal{K}$-valued presheaves on X and topological spaces *over* X.
- Any $\mathcal{K}$-valued presheaf on X gives rise to a prestack on X, and all prestacks arise in this way.
- The homotopy limit constructions needed to "stackify" a prestack may be carried out easily on the level of $\mathcal{K}$-valued presheaves.

After establishing these properties, we will be in a position to prove Theorem 3.5.2, which is closely related to Theorem 3.0.5 and has a few other interesting implications. Finally, in §3.6 we will show how to deduce Theorem 3.0.5 from Theorem 3.5.2.

3.1. **Some Point-Set Topology.** Let X denote a paracompact topological space. In order to prove Theorem 3.5.2, we will need to investigate the homotopy theory of prestacks on X. We then encounter the following technical obstacle: an open subset of a paracompact space need not be paracompact. Because we wish to deal only with paracompact spaces, it will be convenient to restrict our attention to presheaves which are defined only with respect to a particular basis $\mathcal{B}$ for X consisting of paracompact open sets. More precisely, we need the following:

Lemma 3.1.1. *Let X be a paracompact topological space. There exists a collection $\mathcal{B}$ of open subsets of X with the following properties:*

- *The elements of $\mathcal{B}$ form a basis for the topology of X.*
- *Each element of $\mathcal{B}$ is paracompact.*
- *The elements of $\mathcal{B}$ are stable under finite intersections.*

Proof. If X is a metric space, then we may take $\mathcal{B}$ to consist of the collection of *all* open subsets of X. In general, this does not work because the property of paracompactness is not inherited by open subsets. A more general solution is to take $\mathcal{U}$ to consist of all open F_σ subsets of X, together with X itself. Recall that a subset of X is called an F_σ if it is a countable union of closed subsets of X. It is then clear that $\mathcal{B}$ is stable under finite intersections. The paracompactness of each element of $\mathcal{B}$ follows since each is a countable union of paracompact spaces (see for example [20]). To see that $\mathcal{B}$ forms a basis for X, we argue as follows. Given any $x \in V \subseteq X$, we may construct a sequence of closed subsets $K_i \subseteq X$ with $K_0 = \{x\}$ and such that $K_{i+1} \subseteq V$ contains some neighborhood of K_i (by normality). Then $U = \bigcup K_i$ is an open F_σ containing x and contained in V. □

In the sequel, we will frequently suppose some collection of open sets $\mathcal{B}$ has been chosen so as to have the above properties. Then $\mathcal{B}$ may be viewed as a category with finite limits, and is equipped with a natural Grothendieck topology. The topos of sheaves of sets on X is equivalent to the topos of sheaves of sets on $\mathcal{B}$. Thus, the proof of Proposition 2.7.3 shows that the ∞-topos of X^∞ is equivalent to the ∞-topos of stacks on $\mathcal{B}$ (in the sense of Proposition 2.7.3).

3.2. **Presheaves of Spaces.** Throughout this section, we will assume that X is a topological space and that $\mathcal{B}$ is a collection of open subsets of X which is stable under finite intersections. We will abuse notation by identifying X with the topos of sheaves on X.

In order to prove Theorem 3.0.5, we will need to do some calculations in the ∞-topos ΔX. For this, we will need to use some concrete models for objects in ΔX, and to exhibit these we need to choose a model for homotopy theory. The homotopy theory of "sheaves of spaces" is usually formalized in terms of simplicial presheaves, as in [14]. However, for our purposes it will be more convenient to use a somewhat obscure model for homotopy types.

Let $\mathcal{K}_0$ denote the (ordinary) category of compact metrizable spaces (and continuous maps). We let $\mathcal{K}$ denote the category $\mathrm{Ind}(\mathcal{K}_0)$. Passing to filtered colimits gives a functor F from $\mathcal{K}$ to topological spaces. Inside of $\mathcal{K}_0$ one may find the ordinary category of finite CW-complexes, and inside of $\mathcal{K}$ one may find the ordinary category of *all* CW-complexes. On these subcategories the functor F is fully faithful. In general, the functor F is neither full nor faithful, and its essential image contains many "bad" spaces. However, for our purposes it will be convenient to use the category $\mathcal{K}$ as a model for homotopy theory.

Remark 3.2.1. Since the category $\mathcal{K}_0$ is essentially small and has finite colimits, the category $\mathcal{K}$ is presentable, and may therefore be identified with the category of colimit-preserving presheaves on $\mathcal{K}$, which are in turn determined by their restriction to $\mathcal{K}_0$. In other words, we may view an object $Z \in \mathcal{K}$ as being given by a contravariant functor $\mathrm{Hom}(\bullet, Z)$ on the category of compact metrizable spaces. This functor is required to be compatible with finite colimits in $\mathcal{K}_0$, and every such functor is representable by an object of $\mathcal{K}$.

We let $\mathcal{K}_{\mathcal{B}}$ denote the (ordinary) category of presheaves on $\mathcal{B}$ with values in $\mathcal{K}$. Since $\mathcal{K}_0$ contains the category of finite simplices, every object of $\mathcal{K}$ determines a (fibrant) simplicial set and we obtain a functor $\Psi : \mathcal{K} \to \mathcal{S}$, and therefore a functor $\Psi_{\mathcal{B}} : \mathcal{K}_{\mathcal{B}} \to \mathcal{S}^{\mathcal{B}^{op}}$. Proposition 3.4.1, which we shall prove later, implies that $\Psi_{\mathcal{B}}$ is essentially surjective. We wish to use the source category $\mathcal{K}_{\mathcal{B}}$ as "models" for the target ∞-category $\mathcal{S}^{\mathcal{B}^{op}}$. We use this terminology only to suggest intuition; we neither need nor use the language of model categories.

We will let Sp_X denote the (ordinary) category of *Hausdorff topological spaces over* X; that is, an object of Sp_X is a Hausdorff space Y equipped with a continuous map $f : Y \to X$. The proof of Theorem 3.0.5 will make use of a pair of adjoint functors

$$S : \mathrm{Sp}_X \to \mathcal{K}_{\mathcal{B}}$$
$$|\,| : \mathrm{Sp}_X \leftarrow \mathcal{K}_{\mathcal{B}}$$

which we now describe.

We will begin with the functor S, whose definition is more intuitive. Let $f : Y \to X$ be an object of Sp_X. We define an object $S(Y) \in \mathcal{K}_{\mathcal{B}}$ by the following formula:

$$\mathrm{Hom}_{\mathcal{K}}(Z, S(Y)(U)) = \mathrm{Hom}_X(Z \times U, Y)$$

where Z is any compact metrizable space, U is any element of $\mathcal{B}$, and the right hand side denotes the space of continuous maps *over* the space X. Using the fact that Y is Hausdorff, one readily checks that this formula is compatible with colimits in Z, and therefore uniquely defines $S(Y)(U) \in \mathcal{K}$ (which is automatically functorial in U). We note that S is a functor from Sp_Y to $\mathcal{K}_{\mathcal{B}}$.

The functor S has a left adjoint, which we will denote by

$$\mathcal{F} \mapsto |\,\mathcal{F}\,|.$$

We shall call $|\,\mathcal{F}\,|$ the *geometric realization* of $\mathcal{F}$. The construction of $|\,\mathcal{F}\,|$ is straightforward: one begins by gluing together the spaces $\mathcal{F}(U) \times U$ in the obvious way, and then takes a maximal Hausdorff quotient. We leave the details to the reader.

3.3. **Homotopy Limits.** Let X be a topological space, $\mathcal{B}$ a basis for the topology of X which is closed under finite intersections.

Suppose that $\mathcal{F} \in \mathcal{K}_{\mathcal{B}}$ is a $\mathcal{K}$-valued presheaf on the basis $\mathcal{B}$. Then $\Psi_{\mathcal{B}}\,\mathcal{F}$ is a prestack on $\mathcal{B}$. In the process of sheafifying $\Psi_{\mathcal{B}}\,\mathcal{F}$, we encounter the limit

$$\lim_{V \in \mathfrak{S}} \Psi_{\mathcal{B}}(\mathcal{F})(V)$$

where $\mathfrak{S}$ is a sieve on an open set $U \in \mathcal{B}$. This limit is canonically determined in the ∞-category $\mathcal{S}$ of spaces. However, for concrete computations, we would like to represent this limit by an element of $\mathcal{K}$ which depends functorially (in the usual, 1-categorical sense) on $\mathcal{F}$. For this, we shall fix an open cover $\mathcal{U} = \{U_\alpha\}_{\alpha \in A}$ of U which generates the sieve $\mathfrak{S}$.

Now define an object $\operatorname{holim}_{\mathcal{U}} \mathcal{F} \in \mathcal{K}$ by the following universal property:

- For any compact metrizable space K, $\operatorname{Hom}(K, \operatorname{holim}_{\mathcal{U}} \mathcal{F})$ is equal to the set of all collections of compatible maps $\phi_v : K \times \Delta^n \to \mathcal{F}_\bullet(U_{v(0)} \cap U_{v(1)} \cap \ldots \cap U_{v(n)})$, where v ranges over all functions $[0,n] \to A$. The compatibility requirement is that $\phi_v \circ (\mathrm{id}_K \times \Delta^f) = \phi_{v \circ f}$, for any factorization $[0,m] \xrightarrow{f} [0,n] \xrightarrow{v} \mathcal{U}$.

Remark 3.3.1. The existence of $\operatorname{holim}_{\mathcal{U}} \mathcal{F}$ depends on the representability of the functor described above. It suffices to show that the functor carries finite (colimits) in K to finite limits, which is obvious.

If $\mathcal{F} \in \mathcal{K}_{\mathcal{B}}$, then $\operatorname{holim}_{\mathcal{U}} \mathcal{F}$ actually represents the limit

$$\lim_{V \in \mathfrak{S}} \Psi \mathcal{F}(V) \in \mathcal{S}.$$

In particular, it is independent of the choice of $\mathcal{U}$, up to canonical homotopy equivalence. This is straightforward and left to the reader.

If $\mathcal{U}$ is an open cover of U generating a sieve $\mathfrak{S}$, then the functor

$$\mathcal{F} \mapsto \operatorname{holim}_{\mathcal{U}} \mathcal{F}$$

is co-represented by an object $N(\mathcal{U}) \in \mathcal{K}_{\mathcal{B}}$. More precisely, let $N(\mathcal{U})(V)$ denote the object of $\mathcal{K}_{\mathcal{B}}$ whose n-simplices are indexed by $(n+1)$-element subsets of $\{\alpha \in A : V \subseteq U_\alpha\}$. Then we have

$$\operatorname{Hom}_{\mathcal{K}}(K, \operatorname{holim}_{\mathcal{U}} \mathcal{F}) = \operatorname{Hom}_{\mathcal{K}_{\mathcal{B}}}(K \times N(\mathcal{U}), \mathcal{F})$$

where on the right side, we regard K as a constant presheaf in $\mathcal{K}_{\mathcal{B}}$.

Now we have the following easy result, which explains the usefulness of a basis of *paracompact* open sets.

Lemma 3.3.2. *Suppose that $U \in \mathcal{B}$ is paracompact, and let $\mathfrak{S}$ be a sieve on U, generated by an open cover $\mathcal{U} = \{U_\alpha\}_{\alpha \in A}$. The natural map $\pi : |N(\mathcal{U})| \to U$ is a fiberwise homotopy equivalence of spaces over X. In other words, there exists a map $s : U \to |N(\mathcal{U})|$ such that $\pi \circ s$ is the identity and $s \circ \pi$ is fiberwise homotopic to the identity.*

Proof. Any partition of unity subordinate to the cover $\mathcal{U}$ gives rise to a section s. (Indeed, a section s is almost the same thing as a partition of unity subordinate to $\mathcal{U}$; there is only a slight difference in the local finiteness conditions.) To show that $s \circ \pi$ is fiberwise homotopic to the identity, use a "straight line" homotopy. □

Proposition 3.3.3. *Suppose that $\mathcal{B}$ is a basis for a topological space X which is closed under finite intersections and consists of paracompact sets. Then for any space $Y \in \mathrm{Sp}_X$, the prestack $\Psi_{\mathcal{B}} S(Y)$ on $\mathcal{B}$ is a stack.*

Proof. Let $\mathcal{F} = S(Y)$. Choose any $U \in \mathcal{B}$ and any sieve $\mathfrak{S}$ which covers U; we must show that

$$f : \mathcal{F}(U) \to \lim_{\mathcal{U}} \mathcal{F}$$

is a weak homotopy equivalence (that is, it induces an equivalence after applying Ψ). It will suffice to prove the following statement for any inclusion of finite simplicial complexes $K_0 \subseteq K$:

- Given any map $g : K_0 \to \mathcal{F}(U)$, and any extension h of $f \circ g$ to K, there exists an extension $\widetilde{g}$ of g to K_0 such that $f \circ \widetilde{g}$ is homotopic to h by a homotopy fixed on K_0.

Since K_0 and K are compact metric spaces, the above assertion is equivalent to the following:

- Given any map $g : K_0 \times U \to Y$ which is a section of $Y \to X$ and any extension h of the induced map $K_0 \times N(\mathcal{U}) \to Y$ to $K \times N(\mathcal{U}) \to Y$, there exists an extension $\widetilde{g} : K \times U \to Y$ such that the induced map $K \times N(\mathcal{U}) \to Y$ is homotopic to h by a homotopy fixed on $K_0 \times N(\mathcal{U})$.

This follows immediately from Lemma 3.3.2 and the fact that $K_0 \to K$ is a cofibration. □

Remark 3.3.4. The prestack $\Psi_{\mathcal{B}} S(Y)$ is *not necessarily a hyperstack.* This observation was one of the main motivations for this paper.

3.4. **Simplicial Complexes.** In this section, we discuss the use of simplicial complexes as a model for homotopy theory. Recall that a *combinatorial simplicial complex* consists of the following data:

- A set V of vertices.
- A collection $J(V)$ of finite subsets of V, containing the empty set, and with the property that $S \in J(V)$ and $S' \subseteq S$ implies $S' \in J(V)$

A morphism between combinatorial simplicial complexes $(V, J(V))$ and $(V', J'(V'))$ is a map $\alpha : V \to V'$ such that $\alpha(S) \in J'(V')$ for any $S \in J(V)$. Let $\mathbf{CX}$ denote the category of combinatorial simplicial complexes.

Given a combinatorial simplicial complex $K = (V, J(V))$, we can associate an object $rK \in \mathcal{K}$, the *geometric realization* of K. If V is finite, then we take rK to be the full subcomplex of the simplex Δ^V spanned by the vertices in V, consisting of all faces $F \subseteq \Delta^V$ such that $\{v \in V : v \in F\} \in J(V)$. In the general case, we define rK to be a formal filtered colimit of the spaces rK_α, where $K_\alpha = (V_\alpha, \{S \in J(V) : S \subseteq V_\alpha\})$ and V_α ranges over the finite subsets of V.

In particular, after composition with Ψ, we obtain a geometric realization functor $\mathbf{CX} \to \mathcal{S}$. It is well-known that this functor is essentially surjective (in other words, every topological space has the weak homotopy type of a simplicial complex). In this section, we sketch a proof of a relative version of this statement.

Let $\mathcal{B}$ be a partially ordered set (the particular case we have in mind is that $\mathcal{B}$ is a basis of paracompact open sets for a paracompact space X, but that will play no role in this section). We let $\mathbf{CX}_{\mathcal{B}}$ denote the category of $\mathbf{CX}$-valued presheaves on $\mathcal{B}$ for which the induced presheaf of vertex sets is constant. In other words, an object of $\mathbf{CX}_{\mathcal{B}}$ consists of a vertex set V together with a presheaf of families of finite subsets of V, which satisfy the definition of a combinatorial simplicial complex over each $U \in \mathcal{B}$.

Applying the realization functor R pointwise, we obtain a functor $\mathbf{CX}_{\mathcal{B}} \to \mathcal{K}_{\mathcal{B}}$, which we shall denote also by r. The main result we will need is the following:

Proposition 3.4.1. *The functor $\Psi r : \mathbf{CX}_{\mathcal{B}} \to \mathcal{S}^{\mathcal{B}^{op}}$ is essentially surjective.*

Proof. Let $\mathcal{F}$ be a stack on $\mathcal{B}$. If we model objects of $\mathcal{S}$ by simplicial sets, then we may view $\mathcal{F}$ as a homotopy coherent diagram of simplicial sets, indexed by $\mathcal{B}$. Using standard techniques, we may replace this homotopy coherent diagram by a strictly commutative diagram, which we shall also denote by $\mathcal{F}$.

Not every simplicial set $S_\bullet$ corresponds to a combinatorial simplicial complex: roughly speaking, this requires that simplices of $S_\bullet$ are determined by their vertices. However, this condition is always satisfied if we replace $S_\bullet$ by its first barycentric subdivision, after which we may functorially extract a combinatorial simplicial complex whose geometric realization is homeomorphic to the geometric realization of $S_\bullet$. Applying this procedure to the functor $\mathcal{F}$, we may replace $\mathcal{F}$ with a strictly commutative diagram in the category $\mathbf{CX}$, which we shall denote by $\mathcal{F}'$.

Let $\mathcal{F}'(U) = (V_U, J_U(V_U))$. The problem now is that the sets V_U may vary; we need to replace $\mathcal{F}'$ by a diagram in $\mathbf{CX}$ which has a constant vertex set. This may be achieved as follows: let $V = \coprod_{U \in \mathcal{B}} V_U$. For each $S \in J_U(V_U)$, we let

$$S' = \{v \in V_{U'} : (U \subseteq U') \wedge (v|U \in S)\},$$

and we let $J'_U(V)$ denote the collection of all finite subsets of V which are contained in S' for some $S \in J_U(V_U)$. Then the collection $\{(V, J'_U(V))\}_{U \in \mathcal{B}}$ gives an object of $\mathbf{CX}_{\mathcal{B}}$, regarded as a $\mathcal{B}$-indexed diagram $\mathcal{F}'$ in $\mathbf{CX}$. By construction, there is a natural transformation $\mathcal{F} \to \mathcal{F}'$. One can easily check that it induces homotopy equivalences at each level $U \in \mathcal{B}$. □

3.5. **The Main Result.** The main goal of this section is to prove:

Lemma 3.5.1. *Let X be a paracompact topological space, and $\mathcal{B}$ a basis for the topology of X which is stable under finite intersections and consists of paracompact sets. Let $K \in \mathbf{CX}_{\mathcal{B}}$. Then the induced map*

$$(\Psi_{\mathcal{B}} rK)^+ \to (\Psi_{\mathcal{B}} S|rK|)^+$$

is an equivalence of presheaves of spaces on $\mathcal{B}$. Here $r : \mathbf{CX}_{\mathcal{B}} \to \mathcal{K}_{\mathcal{B}}$ denotes the functor obtained by applying $r : \mathbf{CX} \to \mathcal{K}$ componentwise, and $\mathcal{F} \rightsquigarrow \mathcal{F}^+$ denotes the partial sheafification functor of Proposition 2.7.3.

Proof. We must show that the induced map $p : (\Psi rK)^+(U) \to (\Psi S|rK|)^+(U)$ is an equivalence for each $U \in \mathcal{B}$. Replacing X by U and all other objects by their restrictions to U, we may reduce to the case where $U = X$.

We prove the following: given any map $f : S^n \to (\Psi rK)^+(X)$, and any extension g of $p \circ f$ to D^{n+1}, there exists an extension $\widetilde{f}$ of f to D^{n+1} such that $p \circ \widetilde{f}$ is homotopic to g by a homotopy fixed on S^n. Since $(\Psi rK)^+(X)$ and $(\Psi S|rK|)^+$ are defined as filtered colimits, and S^n and D^{n+1} are finite complexes, we may assume given factorizations $S^n \to \lim_{V \in \mathfrak{S}}(\Psi rK)(V)$ and $D^{n+1} \to \lim_{V \in \mathfrak{S}}(\Psi S|rK|)(V)$. Now choose an open cover $\mathcal{U}$ of X which generates the sieve $\mathfrak{S}$; we may then represent f and g by maps $S^n \times N(\mathcal{U}) \to RK$ and $D^{n+1} \times N(\mathcal{U}) \to S|rK|$, which we shall again denote by f and g.

We will produce a map $\widetilde{f} : D^{n+1} \times N(\mathcal{U}) \to rK$, possibly after passing to some refinement of the original covering $\mathcal{U}$, which extends f and induces a map $D^{n+1} \times N(\mathcal{U}) \to S|rK|$ which is homotopic to g by a homotopy fixed on $S^n \times N(\mathcal{U})$. Let $\mathcal{U} = \{U_\alpha\}_{\alpha in A}$. We first choose a covering $\mathcal{U}' = \{U'_\alpha\}_{\alpha \in A}$ such that $\mathcal{U}'$ is locally finite and $\overline{U'_\alpha} \subseteq U_\alpha$ for each $\alpha \in A$.

Let $K = (V, J(V))$. Then we may regard g as determining a continuous map $D^{n+1} \times |N(\mathcal{U})| \to |rK| \subseteq \Delta^V \times X$, where Δ^V denotes the (possibly infinite) simplex spanned by the vertices in V. For each $v \in V$, let Δ^V_v denote the open subset of Δ^V consisting of points with nonzero v-coordinate, so that the open sets Δ^V_v cover Δ^V.

Choose any $\alpha \in A$ and any $x \in U'_\alpha$. Let $A_x = \{\alpha' \in A : x \in \overline{U'_{\alpha'}}\}$, so that g determines a map $D^{n+1} \times \Delta^{A_x} \times W_{x,\alpha} \to \Delta^V$ for some small neighborhood $W_{x,\alpha} \subseteq U'_\alpha$ of x. Note that A_x is a finite set. Using the compactness of $D^{n+1} \times \Delta^{A_x}$ and a tube-lemma argument, we can find (after possibly shrinking W_x) a finite set $V_x \subseteq V$ and open subsets $Z_v \subseteq D^{n+1} \times \Delta^{A_x}$ for $v \in V_x$ such that $g(Z_v \times W_{x,\alpha}) \subseteq \Delta^V_v$. Choose a partition of unity dominated by the cover $\{Z_v : v \in V_x\}$, and use this partition of unity to define a map $g_{x,\alpha} : D^{n+1} \times \Delta^{A_x} \times W_x \to \Delta^{V_x}$.

Now let $B = \{(x, \alpha) : (\alpha \in A) \wedge x \in U'_\alpha$, and consider the open cover $\mathcal{W} = \{W_{x,\alpha}\}_{(x,\alpha) \in B}$. This cover refines $\mathcal{U}$ via the projection $\pi : B \to A$. We define a map $\widetilde{f}' : D^{n+1} \times N(\mathcal{W}) \to RK$ as follows. Given $\beta_0, \ldots, \beta_m \in B$, we let $\widetilde{f}'(D^{n+1} \times \Delta^{\{\beta_0, \ldots, \beta_n\}} \to RK(W_{\beta_0} \cap \ldots \cap W_{\beta_m})$ by the formula

$$\widetilde{f}(z, \sum_{0 \leq i \leq m} c_i \beta_i) = \sum_{0 \leq i \leq m} c_i g_{\beta_i}(x, \sum_{0 \leq j \leq m} c_j \pi(\beta_j)).$$

One readily checks that $\widetilde{f}|S^n \times N(\mathcal{W})$ is homotopic to f (restricted to $N(\mathcal{W})$) and that $p \circ \widetilde{f}'$ is homotopic to g (restricted to $N(\mathcal{W})$) via straight-line homotopies (which are obviously compatible with one another). A standard argument, using the fact that $S^n \to D^{n+1}$ is a cofibration, allows us to replace $\widetilde{f}'$ by a function $\widetilde{f}$ with the desired properties. ☐

Now, the hard work is done and we can merely collect up the consequences.

Theorem 3.5.2. *Let X be a paracompact topological space, and $\mathcal{B}$ a basis for the topology of X which consists of paracompact open sets and is stable under intersections. Then:*

- *For any prestack $\mathcal{F}$ on $\mathcal{B}$, $\mathcal{F}^+$ is a stack.*
- *If $\mathcal{F}$ is a prestack on $\mathcal{B}$ which is represented by an object $K \in \mathbf{CX}_{\mathcal{B}}$, then $S|rK|$ represents the sheafification of $\mathcal{F}$.*

Proof. The claim follows immediately from Lemmas 3.3.3 and 3.5.1. ☐

Remark 3.5.3. The first part of the theorem asserts that on paracompact topological spaces, sheafification requires only one step. This generalizes the well-known fact that Čech resolutions can be used to compute cohomology on paracompact spaces.

3.6. **The Proof of Theorem 3.0.5.** Now that we have Theorem 3.5.2 in hand, it is easy to give a proof of Theorem 3.0.5.

Let K be a combinatorial simplicial complex, X a paracompact topological space, and $\mathcal{X}$ the ∞-topos of stacks on X. Choose a basis $\mathcal{B}$ for X consisting of paracompact open sets and stable under finite intersections, and let $\widetilde{K} \in \mathcal{K}_{\mathcal{B}}$ denote the corresponding presheaf of spaces on $\mathcal{B}$. Let $p : X \to *$ be the canonical map, which induces a functor $p^* : \mathcal{S} \to \mathcal{X}$. By Theorem 3.5.2, the set $\pi_0 \operatorname{Hom}_{\mathcal{X}}(1, p^* RK)$ may be identified with the set of homotopy classes of sections of the projection $f : |\widetilde{K}| \to X$. To complete the proof, it will suffice to show the following:

Proposition 3.6.1. *There is a natural bijection between the set of homotopy classes of sections of f and homotopy classes of maps $X \to |K|$, where $|K|$ denotes the geometric realization of K.*

Proof. There is a continuous bijection $b : |\widetilde{K}| \to X \times |K|$, so that any section of f gives a map $X \to |K|$. The only problem is that b is not necessarily a homeomorphism, so it is not clear that every continuous map $X \to |K|$ arises in this way. However, we will show that this is always true "up to homotopy". In order to see that the corresponding statement is also true for homotopies *between* continuous maps $X \to |K|$, it is useful to formulate a slightly more general assertion.

If Z is a compact metrizable space, let us call a continuous map $X \times Z \to |K|$ *completely continuous* if it lifts to a continuous map $X \times Z \to |\widetilde{K}|$ (over X). We will prove the following claim:

- If $Z_0 \subseteq Z$ is a cofibration of compact metrizable spaces, and $g : Z \times X \to |K|$ is such that $g|Z_0 \times X$ is completely continuous, then g is homotopic to a completely continuous map by a homotopy which is fixed on Z_0.

Assuming the claim for the moment, let us complete the proof. Applying the claim in to the case where Z is a point, we deduce that every continuous map $X \to |K|$ is homotopic to one induced by a section of f. The injectivity assertion is proved by applying the claim in the case where $Z = [0,1]$, $Z_0 = \{0,1\}$.

It remains to establish the claim. Since the inclusion of Z_0 into Z is a cofibration, a standard argument shows that it will suffice to produce *any* completely continuous map homotopic to g; one can then correct the continuous map and the homotopy so that it is fixed on $Z_0 \times X$. In other words, we may assume that Z_0 is empty.

Let $K = (V, J(V))$. Then we may regard $|K|$ as a subset of the infinite simplex spanned by the elements of V; a map $g : Z \times X \to |K|$ is determined by component maps $g_v : Z \times X \to [0,1]$ which satisfy the following conditions:

(1) Each g_v is continuous.
(2) For any point $(z,x) \in Z \times X$, the set $\{v \in V : g_v(z,x) \neq 0\}$ lies in $J(V)$. In particular, it is finite.
(3) For each $(z,x) \in Z \times X$, the sum $\sum_v g_v(z,v)$ is equal to 1.

Every collection of functions $\{g_v\}$ satisfying the above condition gives a map of *sets* $g : Z \times X \to |K|$, but the function g is not necessarily continuous: this is equivalent to a complicated local finiteness condition which is slightly stronger than condition (2) above. The complete continuity of g is equivalent to an even stronger local finiteness condition. However, both of these local finiteness conditions are satisfied if the collection $\{g_v\}$ satisfies the following condition:

- For any $(z,x) \in Z \times X$, there is an open neighborhood $U \subseteq Z \times X$ of (z,x) such that $g_v|U = 0$ for almost every $v \in V$.

Given any continuous $g : Z \times X \to |K|$, we let $U_v = \{(z,x) \in Z \times X : g_v(z,x) \neq 0\}$. Since $Z \times X$ is paracompact, we may find a locally finite refinement $\{U'_v\}_{v \in V}$ of the cover $\{U_v\}_{v \in V}$. Let $\{g'_v\}_{v \in V}$ denote any partition of unity subordinate to the cover $\{U'_v\}_{v \in V}$. The family $\{g'_v\}$ satisfies all of the conditions enumerated above, including the strongest local finiteness assumption, so that it induces a completely continuous map $g' : Z \times X \to |K|$. To complete the proof, we need only to show that g' is homotopic to g. For this, one uses a "straight-line" homotopy from g to g'. □

4. Dimension Theory

In this section, we will discuss the dimension theory of topological spaces from the point of view of ∞-topoi. We introduce the *homotopy dimension* of an ∞-topos, and explain how it relates to various classical notions: covering dimension for paracompact spaces, cohomological dimension, and Krull dimension of Noetherian spaces. We also show that finiteness of the homotopy dimension of an ∞-topos $\mathcal{X}$ has nice consequences: it implies that every object is the inverse limit of its Postnikov tower, which proves that $\mathcal{X}$ is t-complete. It follows that for topological spaces of finite homotopy dimension, our theory coincides with the Joyal-Jardine theory.

We will conclude by proving a generalization of Grothendieck's vanishing theorem for the cohomology of abelian sheaves on Noetherian topological spaces.

4.1. **Homotopy Dimension.** Let $\mathcal{X}$ be an ∞-topos. We shall say that $\mathcal{X}$ has *homotopy dimension* $\leq n$ if every $(n-1)$-connected object $E \in \mathcal{X}$ has a global section $1 \to E$. We say that $\mathcal{X}$ has *finite homotopy dimension* if there exists $n \geq 0$ such that $\mathcal{X}$ has homotopy dimension $\leq n$.

We shall say that $\mathcal{X}$ is *locally of homotopy dimension* $\leq n$ if every object of $\mathcal{X}$ can be constructed as a colimit of objects E with the property that $\mathcal{X}_{/E}$ has homotopy dimension $\leq n$. We say that $\mathcal{X}$ is *locally of finite homotopy dimension* if every object of $\mathcal{X}$ can be constructed as a colimit of objects E with the property that $\mathcal{X}_{/E}$ has finite homotopy dimension.

Remark 4.1.1. An ∞-topos $\mathcal{X}$ is (locally or globally) of homotopy dimension ≤ -1 if and only if $\mathcal{X}$ is equivalent to $*$, the ∞-category with a single object and a contractible space of endomorphisms (the ∞-category of stacks on the empty space).

Remark 4.1.2. If $\mathcal{X}$ is a coproduct of ∞-topoi $\mathcal{X}_\alpha$, then $\mathcal{X}$ is of homotopy dimension $\leq n$ (locally of homotopy dimension $\leq n$, locally of finite homotopy dimension) if and only if each $\mathcal{X}_\alpha$ is of homotopy dimension $\leq n$ (locally of homotopy dimension $\leq n$, locally of finite homotopy dimension).

Remark 4.1.3. The ∞-topos $\mathcal{S}$ is of homotopy dimension ≤ 0 and locally of homotopy dimension ≤ 0. For any object $E \in \mathcal{S}$, the slice ∞-topos $\mathcal{S}_{/E}$ is of homotopy dimension $\leq n$ if E can be represented by a CW complex with cells only in dimensions $\leq n$.

Lemma 4.1.4. *Let $\mathcal{X}$ be an ∞-topos of homotopy dimension $\leq n$, and let $E \in \mathcal{X}$ be k-connected. Then $\mathrm{Hom}_{\mathcal{X}}(1, E)$ is $(k-n)$-connected. In particular, if E is ∞-connected, then $\mathrm{Hom}_{\mathcal{X}}(1, E)$ is contractible.*

Proof. The proof goes by induction on k. If $k < n-1$ there is nothing to prove, and if $k = n-1$ then the assertion follows immediately from the definition of homotopy dimension. In the general case, we note that $\mathrm{Hom}_{\mathcal{X}}(1, E)$ is $(k-n)$-connected if and only if for any pair of points $p, q \in \mathrm{Hom}_{\mathcal{X}}(1, E)$, the space of paths joining p to q is $(k-n-1)$-connected. Since this space of paths is given by $\mathrm{Hom}_{\mathcal{X}}(1, 1 \times_E 1)$, the result follows from the inductive hypothesis since $1 \times_E 1$ is $(k-n-1)$-connected. □

If E is an object in an ∞-topos X, then we have seen that there is a Postnikov tower

$$\to \tau_n E \to \tau_{n-1} E \to \ldots \to \tau_0 E \to \tau_{-1} E.$$

Let $\tau_\infty E$ denote the inverse limit of this tower. There is a natural map $E \to \tau_\infty E$.

Proposition 4.1.5. *Let $\mathcal{X}$ be an ∞-topos which is locally of finite homotopy dimension. Then the natural map $E \to \tau_\infty E$ is an equivalence for any $E \in \mathcal{X}$.*

Proof. We must show that $\mathrm{Hom}_{\mathcal{X}}(E', E) \to \mathrm{Hom}_{\mathcal{X}}(E', \tau_\infty E)$ is an equivalence for every object $E' \in \mathcal{X}$. Since $\mathcal{X}$ is locally of finite homotopy dimension, it will suffice to prove this in the case where $\mathcal{X}_{E'}$ has homotopy dimension $\leq k$.

Let $\eta : E' \to \tau_\infty E$ be an arbitrary map. For each $n \geq 0$, let η_n denote the induced map $E' \to \tau_n E$, and let $E_n = E' \times_{\tau_n E} E$. Since $E \to \tau_n E$ is n-connected, we deduce that $E_n \to E'$ is n-connected. The space $\mathrm{Hom}_{\tau_\infty E}(E', E)$ is equal to the homotopy inverse limit of the sequence of spaces $\mathrm{Hom}_{\tau_n E}(E', E) = \mathrm{Hom}_{E'}(E', E_n)$. Since $\mathrm{Hom}_{E'}(E', E_n)$ is $(n-k)$-connected by Lemma 4.1.4, we deduce that $\mathrm{Hom}_{\tau_\infty E}(E', E)$ is contractible, as desired. □

Corollary 4.1.6. *If $\mathcal{X}$ is locally of homotopy dimension $\leq n$, then $\mathcal{X}$ satisfies hyperdescent.*

Proof. Let S denote the collection of ∞-connected morphisms of $\mathcal{X}$. Every n-truncated object of $\mathcal{X}$ is S-local by Corollary 2.8.9. It follows that any limit of truncated objects of $\mathcal{X}$ is S-local. Since any object $E \in \mathcal{X}$ is the inverse limit of its Postnikov tower, we deduce that E is S-local. Thus $\mathcal{X} \to S^{-1}\mathcal{X}$ is an equivalence of categories. □

4.2. **Cohomological Dimension.** Let $\mathcal{X}$ be an ∞-topos. Given any sheaf A of groups on $\mathcal{X}$ (that is, an abelian group object in the ordinary category $\mathcal{X}_1$), we may construct a groupoid $X_\bullet$ with $X_n = G^n$. Let $K(G,1) = |X_\bullet|$; this is a 1-truncated object of $\mathcal{X}$. A similar construction permits us to construct Eilenberg-MacLane objects $K(G,n)$ for $n \geq 2$ when G is abelian. We can then define the *cohomology* of the ∞-topos $\mathcal{X}$ with coefficients in G by the formula

$$H^n(\mathcal{X}, G) = \pi_0(\mathrm{Hom}_{\mathcal{X}}(1, K(G,n))).$$

Remark 4.2.1. If $\mathcal{X}$ is the ∞-topos of stacks on some ordinary topos X, then $H^n(\mathcal{X}, G) \simeq H^n(X, G)$, where the right hand side denotes sheaf cohomology, defined using an injective resolution of G in the category of abelian sheaves on X. For some discussion we refer the reader to [14].

Definition 4.2.2. Let $\mathcal{X}$ be an ∞-topos. We will say that $\mathcal{X}$ has *cohomological dimension* $\leq n$ if, for any sheaf of abelian groups G on $\mathcal{X}$, we have $H^k(\mathcal{X}, G) = *$ for $k > n$.

Remark 4.2.3. For small values of n, some authors prefer to require a stronger vanishing condition which applies also when G is a non-abelian coefficient system. The appropriate definition requires the vanishing of cohomology for coefficient groups which are defined only up to inner automorphisms, as in [18]. With the appropriate modifications, Theorem 4.2.8 below remains valid for $n < 2$.

In order to study the cohomological dimension, we will need to be able to recognize Eilenberg-MacLane objects $K(G,n)$. By construction, we note that $K(G,n)$ is n-truncated, $(n-1)$-connected, and possesses a global section. We will next prove that these properties characterize the Eilenberg-MacLane objects. We will deduce this from the following more general statement:

Proposition 4.2.4. *Let $\mathcal{X}$ be an ∞-topos, and let E and E' be objects of $\mathcal{X}$ with base points $\eta : 1 \to E$ and $\eta' : 1 \to E'$. Suppose that $n \geq 1$, E is n-connected, and E' is n-truncated. Then the space of pointed morphisms from E to E' is equivalent to the* (*discrete*) *space of maps of sheaves of groups $\eta^* \pi_n(E) \to {\eta'}^* \pi_n E'$.*

Proof. Suppose first that $\mathcal{X} = \mathcal{S}$. In this case, the result follows from classical obstruction theory. If $\mathcal{X}$ is an ∞-category of presheaves, then the result can be proved by working componentwise. In the general case, we apply Theorem 2.4.1 to realize $\mathcal{X}$ as a left-exact localization of some ∞-category $\mathcal{P}$ of presheaves. Let $f : \mathcal{X} \to \mathcal{P}$ denote the natural geometric morphism.

The object $f_* E$ is not necessarily n-connected in $\mathcal{P}$. However, if we let F denote the mapping fiber $1 \times_{\tau_{n-1} f_* E} f_* E$, then F is a pointed, n-connected object of $\mathcal{P}$ and the natural map $F \to f_* E$ becomes an equivalence upon applying f^*. Consequently, we have $\mathrm{Hom}_*(E, E') = \mathrm{Hom}_*(f^* F, E') = \mathrm{Hom}_*(F, f_* E')$, where Hom_* denotes the base-point compatible morphisms. Since $f_* E'$ is n-truncated in $\mathcal{P}$, we deduce that the latter space is equivalent to the space of morphisms of groups from $\xi^* \pi_n F$ to ${\eta'}^* \pi_n f_* E'$, where $\xi : 1 \to F$ denotes the natural base point.

Since E is an n-truncated object of $\mathcal{X}$, we see that the space of pointed maps from an n-sphere into E is already discrete. Consequently, no truncation is required in the definition of $\pi_n E$, and we see that the formation of $\pi_n E$ is compatible with left-exact functors. It follows that $\pi_n f_* E = f_* \pi_n E$, so that the set of group homomorphisms from $\xi^* \pi_n F$ into $(\eta')^* \pi_n f_* E$ is equivalent to the space of group homomorphisms from $\xi^* \pi_n f_* E \simeq f_* \eta^* \pi_n E$ into $f_* {\eta'}^* \pi_n E'$. Since f_* is fully faithful and compatible with products, this is equivalent to the set of group homomorphisms from $\eta^* \pi_n E$ to ${\eta'}^* \pi_n E'$, as desired. □

We can now give a characterization of the Eilenberg-MacLane objects in an ∞-topos.

Corollary 4.2.5. *Let $\mathcal{X}$ be an ∞-topos, and let $n \geq 1$. Suppose that E is an object of $\mathcal{X}$ which is n-truncated, $(n-1)$-connected, and equipped with a section $\eta : 1 \to E$. Then there exists a canonical equivalence $E \simeq K(\eta^* \pi_n E, n)$.*

Before we can proceed further, we need a lemma.

Lemma 4.2.6. *Let $\mathcal{X}$ be an ∞-topos, $n \geq 0$, and E an n-connected object of $\mathcal{X}$, and $f : \mathcal{X}_{/E} \to \mathcal{X}$ the canonical geometric morphism. Then f^* induces an equivalence of ∞-categories between the ∞-category of $(n-2)$-truncated objects of $\mathcal{X}$ and the $(n-2)$-truncated objects of $\mathcal{X}_E$.*

Proof. We first prove that f^* is fully faithful on the ∞-category of $(n-1)$-truncated objects of $\mathcal{X}$. Let $X, Y \in \mathcal{X}$ be objects, where Y is $(n-1)$-truncated. Then $\mathrm{Hom}_{\mathcal{X}_E}(f^*X, f^*Y) = \mathrm{Hom}_{\mathcal{X}}(E \times X, Y) = \mathrm{Hom}_{\mathcal{X}}(E, Y^X)$. Since Y is $(n-1)$-truncated, Y^X is also $(n-1)$-truncated so that the n-connectedness of E implies that $\mathrm{Hom}_{\mathcal{X}}(E, Y^X) \simeq \mathrm{Hom}_{\mathcal{X}}(1, Y^X) = \mathrm{Hom}_{\mathcal{X}}(X, Y)$.

Now suppose that $X_E \in \mathcal{X}_{/E}$ is $(n-2)$-truncated; we must show that $X_E \simeq f^*X$ for some $X \in \mathcal{X}$ (automatically $(n-2)$-connected by descent, so that X is canonically determined by the first part of the proof). The natural candidate for X is the object $f_* X_E$. We will show that the adjunction $f^* f_* X_E \to X_E$ is an equivalence; then $f_* X_E$ will be $(n-2)$-connected by descent and the proof will be complete.

Let $\pi_0, \pi_1 : E \times E \to E$ denote the two projections. We first claim that there exists an equivalence $\pi_0^* X_E \simeq \pi_1^* X_E$. To see this, we note that both sides become equivalent to X_E after pulling back along the diagonal $\delta : E \to E \times E$. Since E is n-connected, the diagonal map $E \to E \times E$ is $(n-1)$-connected, so that δ^* is fully faithful on $(n-2)$-truncated objects (by the first part of the proof). Thus $\pi_0^* X_E$ and $\pi_1^* X_E$ must be equivalent to begin with (in fact canonically, but we shall not need this).

We wish to show that the adjunction morphism $p : f^* f_* X_E$ is an equivalence; then we may take $X = f_* X_E$. It suffices to show that p is an equivalence after (surjective) base change to E; in other words, we must show that

$$p_E : \pi_0^*(\pi_0)_* \pi_1^* X_E \to \pi_1^* X_E$$

is an equivalence. Since π_0^* is fully faithful (on $(n-1)$-truncated objects), we see immediately that the adjunction morphism $\pi_0^*(\pi_0)_* Y \to Y$ is an equivalence whenever Y lies in the essential image of π_0^*. Since $\pi_1^* X_E \simeq \pi_0^* X_E$, it lies in the essential image of π_0^* and the proof is complete. □

Suppose that E is an $(n-1)$-connected, n-truncated object of $\mathcal{X}$ for $n \geq 2$. Then $\pi_n(E)$ is a sheaf of abelian groups on E. By the preceding lemma, there is a unique sheaf of abelian groups G on $\mathcal{X}$ such that $\pi_n(E) = E \times G$. In this situation, we shall say that E is an *(n, G)-bundle*.

Proposition 4.2.7. *Let $\mathcal{X}$ be an ∞-topos, $n \geq 2$, and G a sheaf of abelian groups on $\mathcal{X}$. Then the (n, G)-bundles on X are classified, up to equivalence, by the abelian group $H^{n+1}(\mathcal{X}, G)$. Under this equivalence, the identity element of $H^{n+1}(\mathcal{X}, G)$ corresponds to $K(G, n)$. An (n, G) bundle is equivalent to $K(G, n)$ if and only if it admits a global section.*

Proof. Given $\eta \in H^{n+1}(\mathcal{X}, G)$, we may represent η by a morphism $\eta : 1 \to K(G, n+1)$, and then form the fiber product $E_\eta = 1 \times_{K(G,n+1)} 1$ using the map η and the base point of $K(G, n+1)$. Since $K(G, n+1)$ is 0-connected, η is locally equivalent to the base point and E is locally equivalent to $K(G, n)$. It follows by descent that E is $(n-1)$-connected and n-truncated. The long exact homotopy sequence gives a canonical identification of $\pi_n E_\eta$ with $G \times E_\eta$. Thus, E_η is an (n, G)-bundle.

For each object $E \in \mathcal{X}$, let $\mathcal{C}_E$ denote the ∞-category of (n, G)-bundles in $\mathcal{X}_{/E}$ (where we also regard G as a sheaf of groups on $\mathcal{X}_{/E}$, via the pullback construction). Morphisms in $\mathcal{C}_E$ are required induce the identity on G. It is easy to see that all morphisms in $\mathcal{C}_E$ are equivalences, and with a little bit of work one can show that $\mathcal{C}_E$ is essentially small. Consequently we may view $\mathcal{C}_E$ as a space which varies contravariantly in E (via the pullback construction). A descent argument shows that the functor $E \mapsto \mathcal{C}_E$ carries colimits into limits, so that by Theorem 1.4.3 this functor is representable by some object $C \in \mathcal{X}$. Then the construction of the first part of the proof gives a map of spaces

$$\phi : K(G, n+1) \to C$$

which we shall show to be an equivalence.

Since every (n, G)-bundle is locally trivial, we deduce that ϕ is surjective. Hence, to show that ϕ is an equivalence it suffices to prove that $K(G, n+1) \times_C K(G, n+1)$ is equivalent to $K(G, n+1)$. In other words, we must show that given two maps $\eta, \eta' : E \to K(G, n+1)$, the space of paths from η to η' in $\mathrm{Hom}_{\mathcal{X}}(E, K(G, n+1))$ is equivalent to the space of identifications between the associated (n, G)-bundles on E. Both of these spaces form stacks on E, so to prove that they are equivalent we are free to replace E by any object which surjects onto E. Thus, we may assume that η and η' are trivial, in which case the associated (n, G) bundles are both equivalent to $E \times K(G, n)$. In other words, we are reduced to proving that the natural map

$$\mathrm{Hom}_{\mathcal{X}}(E, K(G,n)) \to \mathrm{Hom}_E(K(G,n) \times E, K(G,n) \times E)$$

is a homotopy equivalence of the left hand side onto the collection of components on the right hand side consisting of maps which induce the identity on G.

Replacing $\mathcal{X}$ by $\mathcal{X}_{/E}$, we may assume that $E = 1$. We note that $K(G,n)$ has an infinite loop structure and in particular we have a splitting $\mathrm{Hom}_{\mathcal{X}}(K(G,n), K(G,n)) \simeq \mathrm{Hom}_*(K(G,n), K(G,n)) \times \mathrm{Hom}_{\mathcal{X}}(1, K(G,n))$, where the first factor denotes the space of pointed morphisms. By Proposition 4.2.4, this is a discrete space consisting of all morphisms of sheaves of groups from G to G. In particular, there is only one point corresponding to the identity map of G, which proves that ϕ is an equivalence onto its image as desired. □

Theorem 4.2.8. *Let $\mathcal{X}$ be an ∞-topos and $n \geq 2$. Then $\mathcal{X}$ has cohomological dimension $\leq n$ if and only if it satisfies the following condition: any $(n-1)$-connected, truncated object of $\mathcal{X}$ admits a global section.*

Proof. Suppose that $\mathcal{X}$ has the property that every $(n-1)$-connected, truncated object of $\mathcal{X}$ admits a global section. It follows that for any truncated, n-connected object E of $\mathcal{X}$, $\mathrm{Hom}_{\mathcal{X}}(1, E)$ is connected. Let $k > n$, and let G be a sheaf of abelian groups on $\mathcal{X}$. Then $K(G,k)$ is n-connected, so that $H^k(\mathcal{X}, G) = *$. Thus $\mathcal{X}$ has cohomological dimension $\leq n$.

For the converse, let us assume that $\mathcal{X}$ has cohomological dimension $\leq n$ and let E denote an $(n-1)$-connected, k-truncated object of $\mathcal{X}$. We will show that E admits a global section by induction on k. If $k \leq n-1$, then $E = 1$ and there is nothing to prove. For the inductive step, we may assume that $\tau_{k-1}E$ admits a global section η. We may replace E by $1 \times_{\tau_{k-1}E} E$, and thereby assume that E is $(k-1)$-connected. Let $G = \pi_k E$; then by Lemma 4.2.6, $G = E \times G'$ for some sheaf of groups G' on $\mathcal{X}$. Since $H^{k+1}(\mathcal{X}, G) = *$, we deduce that the $E \simeq K(G', k)$ and therefore possesses a global section. □

Corollary 4.2.9. *Let $\mathcal{X}$ be an ∞-topos. If $\mathcal{X}$ has homotopy dimension $\leq n$, then $\mathcal{X}$ has cohomological dimension $\leq n$. The converse holds provided that $\mathcal{X}$ has finite homotopy dimension and $n \geq 2$.*

Proof. Only the last claim requires proof. Suppose that $\mathcal{X}$ has cohomological dimension $\leq n$ and homotopy dimension $\leq k$. We must show that every $(n-1)$-connected object E of $\mathcal{X}$ has a global section. The object $\tau_{k-1}E$ is truncated and $(n-1)$-connected, so it has a global section by Theorem 4.2.8. Replacing E by $E \times_{\tau_{k-1}E} 1$, we can reduce to the case where E is $(k-1)$-connected, which follows from the definition of homotopy dimension. □

In the next two sections, we will examine classical conditions which give bounds on the cohomological dimension and prove that they also give bounds on the homotopy dimension. We do not know if every ∞-topos of finite cohomological dimension also has finite homotopy dimension (though this seems unlikely). In particular, we do not know if the ∞-topos of stacks on $B\hat{Z}$, the classifying topos of the profinite completion of $\mathbb{Z}$, has finite homotopy dimension. This topos is known to have cohomological dimension 2; see for example [26].

4.3. **Covering Dimension.**

In this section, we will review the classical theory of covering dimension for paracompact spaces, and then show that the covering dimension of a paracompact space X coincides with its homotopy dimension.

Let X be a paracompact space. Recall that X has *covering dimension* $\leq n$ if the following condition is satisfied: for any open covering $\{U_\alpha\}$ of X, there exists an open refinement $\{V_\alpha\}$ of X such that each intersection $V_{\alpha_0} \cap \ldots \cap V_{\alpha_{n+1}} = \emptyset$ provided the α_i are pairwise distinct.

Remark 4.3.1. If X is paracompact, the above definition is equivalent to the (a priori weaker) requirement that such a refinement exist when $\{U_\alpha\}$ is a finite covering of X. This weaker condition gives a good notion whenever X is a normal topological space. Moreover, if X is normal, then the covering dimension of X (by this new definition) is equal to the covering dimension of the Stone-Čech compactification of X. Thus, the dimension theory of normal spaces is controlled by the dimension theory of compact Hausdorff spaces.

Remark 4.3.2. Suppose that X is a compact Hausdorff space, which is written as a filtered inverse limit of compact Hausdorff spaces $\{X_\alpha\}$, each of which has dimension $\leq n$. Then X has dimension $\leq n$. Conversely, any compact Hausdorff space of dimension $\leq n$ can be written as a filtered inverse limit of finite simplicial complexes having dimension $\leq n$. Thus, the dimension theory of compact Hausdorff spaces is controlled by the (completely straightforward) dimension theory of finite simplicial complexes.

Remark 4.3.3. There are other approaches to classical dimension theory. For example, a topological space X is said to have *small (large) inductive dimension* $\leq n$ if every point of X (every closed subset of X) has arbitrarily small open neighborhoods U such that ∂U has small inductive dimension $\leq n-1$. These notion are well-behaved for separable metric spaces, where they coincides with the covering dimension (and with each other). In general, the covering dimension has better formal properties.

We now give an alternative characterization of the covering dimension of a paracompact space. First, we need a technical lemma.

Lemma 4.3.4. *Let X be a paracompact space, $k \geq 0$, $\{U_\alpha\}_{\alpha \in A}$ be a covering of X. Suppose that for any $J \subseteq U_\alpha$ of size $k+1$, we are given a covering $\{V_{J,\beta}\}_{\beta \in B_J}$ of the intersection $U_J = \bigcap_{\alpha \in J} U_\alpha$. Then there exists a covering $\{W_\alpha\}_{\alpha \in \widetilde{A}}$ of X and a map $\pi : \widetilde{A} \to A$ with the following properties:*

- *For $\alpha \in \widetilde{A}$, $W_\alpha \subseteq U_{\pi(\alpha)}$.*
- *Suppose that $\alpha_0, \ldots, \alpha_k \in \widetilde{A}$ have the property that $J = \{\pi(\alpha_0), \ldots, \pi(\alpha_k)\}$ has cardinality $k+1$. Then there exists $\beta \in B_J$ such that $W_{\alpha_0} \cap \ldots \cap W_{\alpha_k} \subseteq V_{J,\beta}$.*

Proof. Since X is paracompact, we may find a locally finite refinement $\{U'_\alpha\}_{\alpha \in A}$ which covers X, such that the each closure $\overline{U'_\alpha}$ is contained in U_α. Let S denote the set of all subsets $J \subseteq A$ having size $k+1$. For $J \in S$, let $K_J = \bigcap_{\alpha \in J} \overline{U_\alpha}$. Now let

$$\widetilde{A} = \{(\alpha, J, \beta) : \alpha \in A, J \in S, \alpha \in J, \beta \in B_J\} \coprod A.$$

For $(\alpha, J, \beta) \in \widetilde{A}$, we set $\pi(\alpha, J, \beta) = \alpha$ and $W_{\alpha,J,\beta} = (U'_\alpha - \bigcup_{\alpha \in J'} K_{J'}) \cup (V_{J,\beta} \cap U'_\alpha)$. If $\alpha \in A$, we let $\pi(\alpha) = \alpha$ and $W_\alpha = (U'_\alpha - \bigcup_{\alpha \in J'} K_{J'})$. It is easy to check that this assignment has the desired properties. $\square$

Theorem 4.3.5. *Let X be a paracompact topological space of covering dimension $\leq n$. Then X^∞ has homotopy dimension $\leq n$.*

Proof. Suppose that $\mathcal{F}$ is an $(n-1)$-connected stack on X. We must invoke the notation (but not the conclusions) of Section §3. Let $\mathcal{B}$ denote a basis for X satisfying the conclusions of Lemma 3.1.1. Then $\mathcal{F}$ may be represented by some $\widetilde{\mathcal{F}} \in \mathbf{CX}_{\mathcal{B}}$. We wish to produce a global section of $\mathcal{F}$. It will be sufficient (and also necessary, by Theorem 3.5.2) to produce a map $N(\mathcal{U}) \to r\widetilde{\mathcal{F}}$ for some open cover $\mathcal{U}$ of X consisting of elements of $\mathcal{B}$.

We will prove the following statement by induction on i, $-1 \leq i \leq n$:

- There exists an open cover $\{U_\alpha\}$ of X (consisting of elements of the basis $\mathcal{B}$) and a map from the i-skeleton of the nerve of this cover into $r\widetilde{\mathcal{F}}$ (in the category $\mathcal{K}_{\mathcal{B}}$).

Assume that this statement holds for $i = n$. Passing to a refinement, we may assume that the cover $\{U_\alpha\}$ has the property that no more than $n+1$ of its members intersect (this is the step where we shall use the assumption on the covering dimension of X). It follows that the n-skeleton of the nerve is the entire nerve, so that we obtain a global section of $\mathcal{F}$.

To begin the induction in the case $i = -1$, we use the cover $\{X\}$; the (-1)-skeleton of the nerve of this cover is empty so there is no data to provide.

Now suppose that we have exhibited the desired cover $\{U_\alpha\}_{\alpha \in A}$ for some value $i < n$. Suppose $J \subseteq A$ has cardinality $i+2$. Over the open set $U_J = \bigcap_{\alpha \in J} U_\alpha$, our data provides us with a map $f_J : \partial \Delta^{i+1} \times U_J \to \mathcal{F}$. By assumption this map is locally trivial, so that we may cover U_J by open sets $\{V_{J,\beta}\}_{\beta \in B_J}$ over which the map f_J extends to a map $f'_{J,\beta} : \Delta^{i+1} \times V_{J,\beta} \to \mathcal{F}$. We apply Lemma 4.3.4 to this data, to obtain an new open cover $\{W_\alpha\}_{\alpha \in \widetilde{A}}$ which refines $\{U_\alpha\}_{\alpha \in A}$. Refining the cover further if necessary, we may assume that each of its members belongs to $\mathcal{B}$. By functoriality, we obtain a map f from the i-skeleton of the nerve of $\{W_\alpha\}_{\alpha\in\widetilde{A}}$ to $\mathcal{F}$. To complete the proof, it will suffice to extend f to the $(i+1)$-skeleton of the nerve of $\{W_\alpha\}_{\alpha\in\widetilde{A}}$. Let $\pi : \widetilde{A} \to A$ denote the map of Lemma 4.3.4, and consider any $(i+2)$-element subject $J \subseteq \widetilde{A}$. If $\pi(J)$ has size $< i+2$, then composition with π gives us a canonical extension of f to $\Delta^{i+1} \times \bigcap_{j\in J} W_j$ (since the corresponding simplex becomes degenerate with respect to the covering $\{U_\alpha\}$). If $\pi(J)$ has size $i+2$, then Lemma 4.3.4 assures us that $\bigcap_{j\in J} W_j$ is contained in some $V_{\pi(J),\beta}$, so that the desired extension exists. □

Remark 4.3.6. In fact, the inequality of Theorem 4.3.5 is an equality: if the homotopy dimension of X is $\leq n$, then the cohomological dimension of X is $\leq n$, so that the paracompact space X has covering dimension $\leq n$.

4.4. **Heyting Dimension.** Let X be a topological space. Recall that X is *Noetherian* if the collection of closed subsets of X satisfies the descending chain condition. In particular, we see that the irreducible closed subsets of X form a well-founded set. Consequently there is a unique ordinal-valued rank function r, defined on the irreducible closed subsets of X, having the property that $r(C)$ is the smallest ordinal which is larger than $r(C_0)$ for any proper subset $C_0 \subset C$. We call $r(C)$ the *Krull dimension* of C, and we call $\sup_{C\subseteq X} r(C)$ the *Krull dimension* of X.

Let X be a topological space. We shall say that X is a *Heyting space* if satisfies the following conditions:

- The space X has a basis consisting of compact open sets.
- The compact open subsets of X are stable under finite intersections.
- If U and V are compact open subsets of X, then the interior of $U \cup (X - V)$ is compact.
- Every irreducible closed subset of X has a unique generic point (in other words, X is a *sober* topological space).

Remark 4.4.1. The last condition is not very important. Any topological space which does not satisfy this condition can be replaced by a topological space which does, without changing the lattice of open sets.

Remark 4.4.2. Recall that a *Heyting algebra* is a distributive lattice L with the property that for any $x, y \in L$, there exists a maximal element z with the property that $x \wedge z \subseteq y$. It follows immediately from our definition that the lattice of compact open subsets of a Heyting space forms a Heyting algebra. Conversely, given any Heyting algebra one may form its spectrum, which is a Heyting space. This sets up a duality between the category of Heyting spaces and the category of Heyting algebras, which is a special case of a more general duality between coherent topological spaces and distributive lattices. We refer the reader to [21] for more details.

Remark 4.4.3. Suppose that X is a Noetherian topological space in which every irreducible closed subset has a unique generic point. Then X is a Heyting space, since every open subset of X is compact.

Remark 4.4.4. If X is a Heyting space and $U \subseteq X$ is a compact open subset, then X and $X - U$ are also Heyting spaces. In this case, we say that $X - U$ is a *cocompact* closed subset of X.

We next define the dimension of a Heyting space. The definition is recursive. Let α be an ordinal. A Heyting space X has *Heyting dimension* $\leq \alpha$ if and only if, for any compact open subset $U \subseteq X$, the boundary of U has Heyting dimension $< \alpha$ (we note that the boundary of U is also a Heyting space); a Heyting space has dimension < 0 if and only if it is empty.

Remark 4.4.5. A Heyting space has dimension ≤ 0 if and only if it is Hausdorff. The Heyting spaces of dimension ≤ 0 are precisely the compact, totally disconnected Hausdorff spaces. In particular, they are also paracompact spaces and their Heyting dimension coincides with their covering dimension.

Proposition 4.4.6. (1) *Let X be a Heyting space of dimension $\leq \alpha$. Then for any compact open subset $U \subseteq X$, both U and $X - U$ have Heyting dimension $\leq \alpha$.*
(2) *Let X be a Heyting space which is a union of finitely many compact open subsets U_α of dimension $\leq \alpha$. Then X has dimension $\leq \alpha$.*

(3) *Let X be a Heyting space which is a union of finitely many cocompact closed subsets K_α of Heyting dimension $\leq \alpha$. Then X has Heyting dimension $\leq \alpha$.*

Proof. All three assertions are proven by induction on α. The first two are easy, so we restrict our attention to (3). Let U be a compact open subset of X, having boundary B. Then $U \cap K_\alpha$ is a compact open subset of K_α, so that the boundary B_α of $U \cap K_\alpha$ in K_α has dimension $\leq \alpha$. We see immediately that $B_\alpha \subseteq B \cap K_\alpha$, so that $\bigcup B_\alpha \subseteq B$. Conversely, if $b \notin \bigcup B_\alpha$ then, for every β such that $b \in K_\beta$, there exists a neighborhood V_β containing b such that $V_\beta \cap K_\beta \cap U = \emptyset$. Let V be the intersection of the V_β, and let $W = V - \bigcup_{b \notin K_\gamma} K_\gamma$. Then by construction, $b \in W$ and $W \cap U = \emptyset$, so that $b \in B$. Consequently, $B = \bigcup B_\alpha$. Each B_α is closed in K_α, thus in X and also in B. The hypothesis implies that B_α has dimension $< \alpha$. Thus the inductive hypothesis guarantees that B has dimension $< \alpha$, as desired. □

Remark 4.4.7. It is not necessarily true that a Heyting space which is a union of finitely many *locally closed* subsets of dimension $\leq \alpha$ is also of dimension $\leq \alpha$. For example, a topological space with 2 points and a nondiscrete, nontrivial topology has Heyting dimension 1, but is a union of two locally closed subsets of Heyting dimension 0.

Proposition 4.4.8. *If X is a sober Noetherian topological space, then the Krull dimension of X coincides with the Heyting dimension of X.*

Proof. We first prove, by induction on α, that if the Krull dimension of a sober Noetherian space is $\leq \alpha$, then the Heyting dimension of a sober Noetherian space is $\leq \alpha$. Let X be a Noetherian topological space of Krull dimension $\leq \alpha$. Using Proposition 4.4.6, we may assume that X is irreducible. Consider any open subset $U \subseteq X$, and let Y be its boundary. We must show that Y has Heyting dimension $\leq \alpha$. Using Proposition 4.4.6 again, it suffices to prove this for each irreducible component of Y. Now we simply apply the inductive hypothesis and the definition of the Krull dimension.

For the reverse inequality, we again use induction on α. Assume that X has Heyting dimension $\leq \alpha$. To show that X has Krull dimension $\leq \alpha$, we must show that every irreducible closed subset of X has Krull dimension $\leq \alpha$. Without loss of generality we may assume that X is irreducible. Now, to show that X has Krull dimension $\leq \alpha$, it will suffice to show that any *proper* closed subset $K \subseteq X$ has Krull dimension $< \alpha$. By the inductive hypothesis, it will suffice to show that K has Heyting dimension $< \alpha$. By the definition of the Heyting dimension, it will suffice to show that K is the boundary of $X - K$. In other words, we must show that $X - K$ is dense in X. This follows immediately from the irreducibility of X. □

We now prepare the way for our vanishing theorem. First, we introduce a modified notion of connectivity:

Definition 4.4.9. Let X be a Heyting space and k any integer. A stack $\mathcal{F}$ on a compact open set $V \subseteq X$ is *strongly k-connected* if the following condition is satisfied: for any $m \geq -1$, any compact open $U \subseteq V$, and any map $\phi : S^m \to \mathcal{F}(U)$, there exists a cocompact closed subset $K \subseteq U$ such that $\overline{K} \subseteq X$ has dimension $< m - k$ and an open cover $\{V_\alpha\}$ of $V - K$ such that the restriction of ϕ to $\mathcal{F}(V_\alpha)$ is nullhomotopic for each α (if $m = -1$, this means that $\mathcal{F}(V_\alpha)$ is nonempty).

Remark 4.4.10. Since the definition involves taking the closure of K in X, rather than in V, we note that the strong connectivity of $\mathcal{F}$ may increase if we replace X by V.

Remark 4.4.11. Strong k-connectivity is an unstable analogue of the connectivity conditions on complexes of sheaves, associated to a the dual of the standard perversity (which is well-adapted to the pushforward functor). For a discussion of perverse sheaves in the abelian context, see for example [24].

Remark 4.4.12. Suppose X has Heyting dimension $\leq n$. If $\mathcal{F}$ is $(k+n)$-connected, then for $m \leq k+n$ we may take $K = \emptyset$ and for $m > k+n$ we may take $K = V$. It follows that $\mathcal{F}$ is strongly k-connected. Conversely, it is clear from the definition that strong k-connectivity implies k-connectivity.

The strong k-connectivity of $\mathcal{F}$ is, by construction, a local property. The key to our vanishing result is that this is equivalent to a stronger *global* property.

Theorem 4.4.13. *Let X be a Heyting space of dimension $\leq n$, let $W \subseteq X$ be a compact open set, and let $\mathcal{F}$ be a stack on W. The following conditions are equivalent:*

1. *For any compact open $U \subseteq V \subseteq W$, any $m \geq -1$, map $\zeta : S^m \to \mathcal{F}(V)$ and any nullhomotopy η of $\eta|U$, there exists an extension of η to $V - K$, where $K \subseteq V$ is a cocompact closed subset and $\overline{K} \subseteq X$ has dimension $< m - k$.*
2. *For any compact open $V \subseteq W$, any $m \geq -1$, and any map $\zeta : S^m \to \mathcal{F}(V)$, there exists a nullhomotopy of ζ on $V - K$, where $K \subseteq V$ is a cocompact closed subset and $\overline{K} \subseteq X$ has dimension $< m - k$.*
3. *The stack $\mathcal{F}$ is strongly k-connected.*

Proof. It is clear that (1) implies (2) (take U to be empty) and that (2) implies (3) (by definition). We must show that (3) implies (1). Replacing W by V and $\mathcal{F}$ by the stack $\mathcal{F}|V \times_{\mathcal{F}|V \otimes S^m} 1$, we may reduce to the case where $W = V$ and $m = -1$.

The proof goes by induction on k. For our base case, we take $k = -n-2$, so that there is no connectivity assumption on the stack $\mathcal{F}$. We are then free to choose $K = X - U$ (it is clear that $\overline{K}$ has dimension $\leq n$).

Now suppose that the theorem is known for strongly $(k-1)$-connected stacks on any compact open subset of X; we must show that for any strongly k-connected $\mathcal{F}$ on V and any $\eta \in \mathcal{F}(U)$, there exists an extension of η to $V - K$ where $\overline{K} \subseteq X$ has dimension $< -1 - k$.

Since $\mathcal{F}$ is strongly k-connected, we deduce that there exists an open cover $\{V_\alpha\}$ of some open subset $V - K_0$, where K_0 has dimension $< -1 - k$ in X, together with elements $\psi_\alpha \in \mathcal{F}(V_\alpha)$. Replacing V by $V - K_0$ and U by $U - K_0 \cap U$, we may suppose $K_0 = \emptyset$.

Since V is compact, we may assume that there exist only finitely many indices α. Proceeding by induction on the number of indices, we may reduce to the case where $V = U \cup V_\alpha \cup K_0$ for some α. Let $\mathcal{F}'$ denote the sheaf on $U \cap V_\alpha$ of *paths from ψ_α to η* (in other words, $\mathcal{F}' = 1 \times_{\mathcal{F}} 1$, where the maps $1 \to \mathcal{F}$ are given by ψ_α and η). Then $\mathcal{F}'$ is strongly $(k-1)$-connected, so by the inductive hypothesis (applied to $\emptyset \subseteq U \cap V_\alpha \subseteq X$), there exists a closed subset $K \subset U \cap V_\alpha$ having dimension $< -k$ in X, such that ψ_α and η are equivalent on $(U \cap V_\alpha) - K$. Since $\overline{K}$ has dimension $< -k$ in X, the boundary ∂K of K has codimension $< -k-1$ in X. Let $W = V_\alpha \cap (X - \overline{K})$. Then we can glue η and ψ_α to obtain a global section of $\mathcal{F}$ over $W \cup U$, which contains $V - \partial K = X - \partial K$. □

Corollary 4.4.14. *Let $\pi : X \to Y$ be a continuous map between Heyting spaces of finite dimension. Suppose that π has the property that for any cocompact closed subset $K \subseteq X$ of dimension $\leq n$, $\pi(K)$ is contained in a cocompact closed subset of dimension $\leq n$. Then the functor π_* carries strongly k-connected stacks into strongly k-connected stacks.*

Proof. This is clear from the characterization (2) of Theorem 4.4.13. □

Corollary 4.4.15. *Let X be a Heyting space of finite dimension, and let $\mathcal{F}$ be a strongly k-connected stack on X. Then $\mathcal{F}(X)$ is k-connected.*

Proof. Apply Corollary 4.4.14 in the case where Y is a point. □

Corollary 4.4.16. *If X is a Heyting space of Heyting dimension $\leq n$, then the ∞-topos X^∞ has homotopy dimension $\leq n$.*

Remark 4.4.17. Since every compact open subset of X is also a Heyting space of Heyting dimension $\leq n$, we can deduce also that the ∞-topos ΔX is locally of homotopy dimension $\leq n$, and therefore is t-complete.

In particular, we obtain Grothendieck's vanishing theorem:

Corollary 4.4.18. *Let X be a Noetherian topological space of Krull dimension $\leq n$. Then X has cohomological dimension $\leq n$.*

Example 4.4.19. Let V be a real algebraic variety (defined over the real numbers, say). Then the lattice of open subsets of V that can be defined by polynomial equations and inequalities is a Heyting algebra, and the spectrum of this Heyting algebra is a Heyting space X having dimension at most equal to the dimension of V. The results of this section therefore apply to X.

More generally, let T be an o-minimal theory (see for example [19]), and let S_n denote the set of complete n-types of T. We endow S_n with the following topology generated by the sets $U_\phi = \{p : \phi \in p\}$, where ϕ ranges over formula with n free variables such that T proves that the set of solutions to ϕ is an open set. Then S_n is a Heyting space of dimension $\leq n$.

Remark 4.4.20. The methods of this section can be adapted to slightly more general situations, such as the Nisnevich topology on a Noetherian scheme of finite Krull dimension. It follows that the ∞-topoi associated to such sites have finite homotopy dimension and hence our theory is equivalent to the Joyal-Jardine theory.

5. Appendix

5.1. **Homotopy Limits and Colimits.** In this appendix, we will summarize (with sketches of proofs) some basic facts regarding limits and colimits in ∞-categories. These are analogous to well-known facts concerning limits and colimits in ordinary categories. We will limit the discussion to colimits; the corresponding statements for limits may be obtained by passing to opposite categories.

In this paper, we make use of colimits where the diagrams are indexed by ∞-categories. Many authors use only ordinary categories to index their colimits, which seems at first to give a less general notion. In fact, this is not the case:

Proposition 5.1.1. *Let κ be an infinite regular cardinal. Then all κ-small homotopy colimits can be constructed using:*

- *Coequalizers and κ-small sums.*
- *Pushouts and κ-small sums.*

Proof. The coequalizer of a pair of morphisms $p, q : X \to Y$ coincides with the pushout of X and Y over $X \coprod X$ (which maps to Y via $p \coprod q$). Given morphisms $X \to Y$ and $X \to Z$, the pushout $Y \coprod_X Z$ may be constructed as the coequalizer of the induced maps $X \to Y \coprod Z$. Thus, coequalizers and coproducts can be constructed in terms of one another (and finite sums); it therefore suffices to prove the assertion regarding pushouts.

We next note that from finite sums and pushouts, one can construct $K \otimes X$ where K is any finite cell complex. We will need this in the case where K is the sphere S^n. If $n = -1$, then $K \otimes S^n$ is the initial object (the empty sum). For $n \geq 0$, we note that $S^n \otimes K$ is the pushout of K with itself over $S^{n-1} \otimes K$.

Let $\mathcal{C}$ be an ∞-category, and let $F : \mathcal{D} \to \mathcal{C}$ be a diagram in $\mathcal{C}$ with $\mathcal{D}$ κ-presented. We may view $\mathcal{D}$ as obtained by a process of "attaching n-cells" as n goes from 0 to ∞, where a 0-cell is an object of $\mathcal{D}$, a 1-cell is a morphism of $\mathcal{D}$, and more generally adjoining an n-cell to $\mathcal{D}$ means adjoining an $(n-1)$-cell to some mapping space $\mathrm{Hom}_{\mathcal{D}}(x, y)$ (together with a mess of other morphisms which are generated by new cell). The hypothesis that $\mathcal{D}$ is κ-generated means that we may construct $\mathcal{D}$ using fewer than κ n-cells for each n, and if $\kappa = \omega$ then we use only finitely many cell in total.

If $\mathcal{D}$ is obtained from a "subcategory" $\mathcal{D}_0$ by adjoining an $(n-1)$-cell to $\mathrm{Hom}(x, y)$, then the colimit of F in $\mathcal{C}$ may be obtained as a pushout $X \coprod_{S^{n-2} \otimes Fx} Fx$, where X is the colimit of $F|\,\mathcal{D}_0$. If $\mathcal{D}$ is obtained from $\mathcal{D}_0$ by adjoining a new object x, then the colimit of F is obtained from the colimit of $F|\,\mathcal{D}_0$ by taking the sum with x. By induction on the number of cells needed to construct $\mathcal{D}$, we can complete the proof in the case where $\kappa = \omega$.

For $\kappa > \omega$, we must work a little bit harder. First of all, we note that if $\mathcal{D}$ is obtained from $\mathcal{D}_0$ by adjoining any κ-small collection of cells simultaneously, then we may construct the colimit of F as a pushout $X \coprod_Y Z$ as above, where X is the colimit of $F|\,\mathcal{D}_0$ and the pair (Y, Z) is a sum of pairs of the form $(S^n \otimes Fx, Fx)$ considered above. In general, $\mathcal{D}$ can be obtained by an infinite sequence of simultaneous cell-attachments, so that the colimit of F is may be written as a direct limit of objects $X_n = \mathrm{colim}\, F|\,\mathcal{D}_n$, and each X_n may be constructed using pushouts and κ-small sums by the argument sketched above.

It therefore suffices to show that we may construct the direct limit of a sequence $X_0 \to X_1 \to \ldots$. But this is easy: the direct limit of the $\{X_i\}$ can be written as a coequalizer of the identity and the shift map

$$\coprod_i X_i \to \coprod_i X_i.$$

It is permissible to form this countably infinite coproduct since $\kappa > \omega$. □

Remark 5.1.2. If $\kappa = \omega$, then finite sums may be constructed by taking pushouts over an initial object. Consequently, all finite colimits may be constructed using an initial object and pushouts.

5.2. **Filtered Colimits.** We now sketch a proof that the general notion of a filtered colimit is not really any more general than "directed colimits", which the reader will surely find familiar.

More specifically, we sketch a proof of the following result, which is used in the body of the paper:

Proposition 5.2.1. *Let κ be a regular cardinal, and let $\mathcal{D}$ be a small κ-filtered ∞-category. Then there exists a κ-directed partially ordered set (which we may regard as an ordinary category, and therefore an ∞-category) $\mathcal{D}_0$, and a functor $F : \mathcal{D}_0 \to \mathcal{D}$ having the following property: for any functor $G : \mathcal{D} \to \mathcal{C}$, the diagrams G and $G \circ F$ have the same colimits.*

Proof. We will sketch a construction of $\mathcal{D}_0$ and leave the property to the reader. The construction uses the Boardman-Vogt approach to ∞-categories based on simplicial sets satisfying a weak Kan condition (see [6] and [7]). According to this approach, we may model $\mathcal{D}$ by a simplicial set D (satisfying certain extension conditions).

Suppose that $D' \subseteq D$ is a simplicial subset and $v \in D'_0$. We shall say that v is *final* in D' if every map $\partial \Delta^n \to D'$ which carries the last vertex of Δ^n into v admits an extension $\Delta^n \to D'$. We call $D' \subseteq D$ *κ-small* if it has fewer than κ nondegenerate vertices.

We let $\mathcal{D}_0$ denote the collection of all κ-small simplicial subsets $D' \subseteq D$ which possess a final vertex. The set $\mathcal{D}_0$ is partially ordered by inclusion. We claim that it is κ-directed: in other words, any subset $S \subseteq \mathcal{D}_0$ of size $< \kappa$ has an upper bound. It suffices to show that *any* κ-small simplicial subset $D' \subseteq D$ can be enlarged to a κ-small simplicial subset with a final vertex. This is more or less equivalent to the assertion that $\mathcal{D}$ is κ-directed.

The functor $F : \mathcal{D}_0 \to \mathcal{D}$ is obtained by choosing a final vertex from each $D' \in \mathcal{D}_0$. Although the final vertex is not necessarily unique, it is unique up to a contractible space of choices (essentially by definition) so that the functor F is well-defined. □

5.3. **Enriched ∞-Categories and $(\infty, 2)$-Categories.** In this section, we discuss $(\infty, 2)$-categorical limits and colimits, which were mentioned several times in the body of the paper. Before we can describe these, we need to decide what an $(\infty, 2)$-category *is*. Just as an ordinary 2-category may be described as a category "enriched over categories", we can obtain the theory of $(\infty, 2)$-categories as a special case of enriched ∞-category theory.

Given any ∞-category $\mathcal{E}$, one can define a notion of "∞-categories enriched over $\mathcal{E}$". A $\mathcal{E}$-enriched ∞-category $\mathcal{C}$ consists of a collection of objects, together with an object $\mathrm{Hom}_{\mathcal{C}}(X, Y) \in \mathcal{E}$ for every pair of objects $X, Y \in \mathcal{C}$. Finally, these objects of $\mathcal{E}$ should be equipped with *coherently associative* composition products

$$\mathrm{Hom}_{\mathcal{C}}(X_0, X_1) \times \ldots \times \mathrm{Hom}_{\mathcal{C}}(X_{n-1}, X_n) \to \mathrm{Hom}_{\mathcal{C}}(X_0, X_n).$$

Remark 5.3.1. More generally, one can replace the Cartesian product $\times$ by any coherently associative monoidal structure $\otimes$ on $\mathcal{E}$. If $\mathcal{E}$ is an ordinary category, then we arrive at the usual notion of an $\mathcal{E}$-enriched category.

Taking $\mathcal{E}$ to be the ∞-category $\mathfrak{S}$ of spaces, we recover the notion of an $(\infty, 1)$-category. On the other hand, if we take $\mathcal{E}$ to be the ∞-category of $(\infty, 1)$-categories (regarded as an $(\infty, 1)$-category by ignoring noninvertible natural transformations), then we obtain the notion of an $(\infty, 2)$-category. Iterating this procedure leads to a notion of (∞, n)-category for all n. The formalization of this process leads to the Simpson-Tamsamani theory of higher categories: see [28].

Remark 5.3.2. The "∞-category of ∞-categories" is much more natural than its 1-categorical analogue, the "category of categories". In the latter case, one should really have a 2-category: it is unnatural to ask for two functors to be equal, and we lose information by ignoring the natural transformations. In the ∞-categorical case, much less information is lost since we discard only the noninvertible natural transformations.

Now, if $\mathcal{C}$ is any ∞-category enriched over $\mathcal{E}$, then any object of $\mathcal{C}$ represents a $\mathcal{E}$-valued functor on $\mathcal{C}$. Thus we obtain an "enriched Yoneda embedding" $\mathcal{C} \to \mathcal{E}^{\mathcal{C}^{op}}$, where the target denotes the $\mathcal{E}$-enriched ∞-category of $\mathcal{E}$-valued presheaves on $\mathcal{C}$ (we consider only presheaves which are compatible with the enrichment; that is, functors $\mathcal{C}^{op} \to \mathcal{E}$ between $\mathcal{E}$-enriched categories).

Given another (small) ∞-category $\mathcal{D}$ which is enriched over $\mathcal{E}$, and a functor $F : \mathcal{D} \to \mathcal{E}$, we obtain a pullback functor $F^* : \mathcal{E}^{\mathcal{C}^{op}} \to \mathcal{E}^{\mathcal{D}^{op}}$ between the ∞-categories of $\mathcal{E}$-valued presheaves. Composing this with Yoneda embedding, we obtain a functor $T : \mathcal{C} \to \mathcal{E}^{\mathcal{D}^{op}}$. Finding colimits in $\mathcal{C}$ amounts to finding a left-adjoint to T. That is, given any $\mathcal{E}$-valued presheaf $\mathcal{F}$ on $\mathcal{D}$, the colimit of $(F, \mathcal{F})$ is an object $X \in \mathcal{C}$ equipped with a map $\mathcal{F} \to TX$ which is universal, in the sense that it induces equivalences

$$\mathrm{Hom}_{\mathcal{C}}(X, Y) \to \mathrm{Hom}_{\mathcal{E}^{\mathcal{D}^{op}}}(\mathcal{F}, TY)$$

where both sides are regarded as living in the ∞-category $\mathcal{E}$.

Remark 5.3.3. The above notion of colimit is apparently more general than the notion we have considered in the body of the paper for the case where $\mathcal{E} = \mathcal{S}$. This generality is only apparent, however: we can replace a pair $(\mathcal{D}, \mathcal{F})$ by the ∞-category $\mathcal{D}_{\mathcal{F}}$ "fibered in spaces" over $\mathcal{D}$; then a colimit of $(F, \mathcal{F})$ in the sense just defined is the same thing as a colimit of $F| \, \mathcal{D}_{\mathcal{F}}$. This does not work in the $\mathcal{E}$-enriched case because there is no analogous "fibered category" construction, so we must work with the more general notion.

Using a variant of the argument of the proof of Proposition 5.1.1, one can show that all $\mathcal{E}$-colimits in $\mathcal{C}$ can be constructed using pushouts, sums, and tensor products $\mathcal{E} \otimes \mathcal{C} \to \mathcal{C}$.

5.4. **Free Algebras.** Granting a good theory of $(\infty, 2)$-categories, we can discuss *monoidal ∞-categories*. As in classical category theory, these may be thought of either as $(\infty, 2)$-categories with a single (specified) object, or as $(\infty, 1)$-categories equipped with a coherently associative product which we shall denote by $\otimes$. We shall adopt the second point of view.

A *monoidal functor* $F : \mathcal{C} \to \mathcal{C}'$ between two monoidal ∞-categories is a functor which is compatible with the associative products $\otimes$ on $\mathcal{C}$ and $\mathcal{C}'$, up to equivalence. (Alternatively, from the $(\infty, 2)$-categorical point of view, we may view a monoidal functor as simply a functor F between the underlying 2-categories *together with* an identification $* \simeq F(*)$, where $*$ is used to denote the specified objects.)

Let $\mathcal{C}$ be a monoidal ∞-category. Suppose that $\mathcal{C}$ is presentable and $A \in \mathcal{C}$. Using Theorem 1.4.3, we see that the following conditions are equivalent:

- The functor $X \mapsto A \otimes X$ commutes with the formation of colimits.
- For each $B \in \mathcal{C}$, there exists an object $\mathrm{Hom}(A, B) \in \mathcal{C}$ and a natural equivalence

$$\mathrm{Hom}_{\mathcal{C}}(\bullet, \mathrm{Hom}(A, B)) \simeq \mathrm{Hom}_{\mathcal{C}}(A \otimes \bullet, B)$$

If these equivalent conditions and their duals (involving products where A appears on the *right*) are satisfied for each $A \in \mathcal{C}$, then we shall say that $\otimes$ is *colimit-preserving*.

If $\mathcal{C}$ is any monoidal ∞-category, one may define an ∞-category $M(\mathcal{C})$ of *monoid objects* of $\mathcal{C}$. The objects of $M(\mathcal{C})$ are objects $A \in \mathcal{C}$ together with multiplications $A^{\otimes n} \to A$ $(n \geq 0)$, together with various homotopies and higher homotopies which explicate the idea that A should be associative up to coherent homotopy.

The purpose of this appendix is to sketch a proof of the following fact which is needed in the body of the paper:

Proposition 5.4.1.

- *Let $\mathcal{C}$ be a presentable monoidal ∞-category, such that the monoidal structure $\otimes$ is colimit-preserving. Let $A \in M(\mathcal{C})$, $N \in \mathcal{C}$, and let $f : A \to N$ be a morphism in $\mathcal{C}$. Then there exists an object $B(f) \in M(\mathcal{C})$ and an equivalence*

$$\mathrm{Hom}_{M(\mathcal{C})}(B(f), \bullet) \simeq \mathrm{Hom}_{M(\mathcal{C})}(A, \bullet) \times_{\mathrm{Hom}_{\mathcal{C}}(A, \bullet)} \mathrm{Hom}_{\mathcal{C}}(M, \bullet).$$

- *Suppose that $\mathcal{C}$ and $\mathcal{C}'$ are both presentable monoidal ∞-categories with colimit-preserving monoidal structures. Let $F : \mathcal{C} \to \mathcal{C}'$ be a monoidal, colimit-preserving functor. Let $A \in M(\mathcal{C})$, $N \in \mathcal{C}$, and $f : A \to M$ a morphism in $\mathcal{C}$. Then the natural map $B(Ff) \to FB(f)$ is an equivalence.*

Proof. To prove the first assertion, we will sketch a construction of $B(f)$. Moreover, in our construction, we will make use only of the monoidal structure on $\mathcal{C}$ and colimits in $\mathcal{C}$. The second part will then follow immediately.

For the construction, we will need to make use of a monoidal ∞-category which is, in some sense, freely generated by a monoid object, a second object, and a map between them. This free ∞-category turns out to be an ordinary category, which is easy to construct explicitly. We will proceed by giving an explicit description of this free category, but we note that this description is largely irrelevant to the proof, for which only the existence is needed.

Let $\mathcal{D}$ denote the ordinary category whose objects are finite, linearly ordered sets with a distinguished subset of marked points. If $J, J' \in \mathcal{D}$, then the morphisms from J to J' in $\mathcal{D}$ are maps $f : J \to J'$ satisfying the following conditions:

- If $j \leq k$, then $f(j) \leq f(k)$.
- If $j \in J$ is marked, then $f(j)$ is marked and $f^{-1}f(j) = \{j\}$.

We define a functor $G : \mathcal{D} \to \mathcal{C}$ as follows. Given an object $J = \{j_1, \ldots, j_n\} \in \mathcal{D}$, with marked points $\{j_{i_1}, \ldots, j_{i_k}\}$, we set

$$G(J) = A_1 \otimes \ldots \otimes A_{j_{i_1}-1} \otimes M_{j_1} \otimes \ldots \otimes M_{j_{i_k}} \otimes A_{j_{i_k}+1} \otimes \ldots \otimes A_n.$$

In words, we simply take a large tensor product, indexed by J, of copies of A (for the unmarked points) and copies of M (for the marked points). Here the subscripts are introduced to index several copies of A and M that are used in the tensor product. The definition of G on morphisms is encoded in the monoid structure on A and the map $f : A \to M$.

Note that $\mathcal{D}$ has the structure of a monoidal category, with the monoidal structure given by concatenation. With respect to this monoidal structure, G is a monoidal functor.

Let G^n denote the induced functor $G \otimes \ldots \otimes G : \mathcal{D}^n \to \mathcal{C}$, and let $B(f)^{\otimes n}$ denote the colimit of G^n, regarded as a diagram in G (in the exceptional case where $n = 0$, we instead define $B(f)^{\otimes 0}$ to be the unit in $\mathcal{C}$). Using the fact that $\otimes$ is colimit-preserving in $\mathcal{C}$, we deduce that $B(f)^{\otimes n}$ is actually a tensor product of n copies of $B(f) = B(f)^{\otimes 1}$. Moreover, using the fact that G is a monoidal functor, we deduce that $G = G^n \circ \delta$ (where δ denotes the diagonal $\mathcal{D} \to \mathcal{D}^n$), so that we get a family of maps $B(f)^{\otimes n} \to B(f)$; further analysis of the situation gives coherence data for this family of maps, so that $B(f)$ has the structure of a monoid in $\mathcal{C}$.

Repeating the above discussion using the subcategory $\mathcal{D}_0 \subseteq \mathcal{D}$ consisting of those finite linearly ordered sets with *no* marked points, we can construct another monoid object of $\mathcal{C}$. Since $\mathcal{D}_0$ has a final object, the colimit is trivial and we recover the monoid A from this construction. Moreover, the inclusion $\mathcal{D}_0 \subseteq \mathcal{D}$ gives a map $A \to B(f)$ of monoids.

On the other hand, consider the object $J = \{*\}$ consisting of a single marked point. By construction we have a map $M = G(J) \to B(f)$, through which the monoid morphism $A \to B(f)$ factors. To complete the proof, it suffices to show that $B(f)$ is universal with respect to this property. This (ultimately) comes down the universal property of $\mathcal{D}$. □

References

[1] Adámek and Rosicky. *Locally Presentable and Acessible Categories.* Cambridge University Press, Cambridge, 1994.

[2] Mardesic and Segal. *Shape Theory.* North-Holland, Amsterdam, 1982.

[3] Makkai and Pare. *Accessible Categories.* Contemporary Mathematics 104. American Mathematical Society, Rhode Island, 1989.

[4] Simpson, Carlos. "A Giraud-type characterization of the simplicial categories associated to closed model categories as ∞-pretopoi".

[5] Artin, M. *Théorie des Topos et cohomologie étale des schémas.* SGA 4. Lecture Notes in Mathematics 269, Springer-Verlag, Berlin and New York, 1972.

[6] Boardman, J.M. and R.M. Vogt. *Homotopy Invariant Structures on Topological Spaces.* Lecture Notes in Mathematics 347. Springer-Verlag, Berlin and New York, 1973.

[7] Günther, Bernd. "The use of semisimplicial complexes in strong shape theory." Glas. Mat. Ser. III 27 (47) ,1992, no. 1, 101-144.

[8] Cordier, Jean Marc and Timothy Porter. "Homotopy Coherent Category Theory". Transactions of the American Mathematical Society 349, 1997, no. 1, 1-54.

[9] "Čech and Steenrod Homotopy Theories with Applications to Geometric Topology". Lecture Notes in Mathematics 542. Springer-Verlag, Berlin and New York, 1976.

[10] Lazard, Daniel. "Sur les modules plats." C.R. Acad. Sci. Paris 258, 1964, 6313-6316.

[11] Bousfield, A. K. "The localization of spaces with respect to homology." Topology 14, 1975, 133-150.

[12] Gordon, R., Power, A.J., and R. Street. "Coherence for tricategories." Memoirs of the American Mathematical Society 116, 1995, no. 558.

[13] Stasheff, J. "Homotopy Associativity of H-spaces I,II." Transactions of the American Mathematical Society 108, 1963, 275-312.

[14] Jardine, J.F. "Simplicial Presheaves." Journal of Pure and Applied Algebra 47, 1987, no.1, 35-87.

[15] Street, Ross. "Two-Dimensional Sheaf Theory." Journal of Pure and Applied Algebra 23, 1982, no. 3, 251-270.

[16] Bourn, Dominique. "Sur les Ditopos." C.R. Acad. Sci. Paris Ser. A 279, 1974, 911-913.

[17] Toën, Bertrand and Gabriele Vezzosi. "Segal topoi and stacks over Segal categories."

[18] Giraud, Jean. *Cohomologie non abelienne.* Springer Verlag, Berlin and New York, 1971

[19] van den Dries, Lou. *Tame Topology and O-minimal Structures.* Cambridge University Press, Cambridge, 1998.

[20] Munkres, James. *Topology.* Prentice Hall, Inc., Englewood Cliffs, New Jersey, 1975.

[21] Johnstone, Peter. *Stone Spaces.* Cambridge University Press, Cambridge, 1982.

[22] Grothendieck, Alexandre. "Sur Quelques points d'algebra homologique." Tohoku Math Journal 9, 1957, 119-221.

[23] Spaltenstein, N. "Resolutions of Unbounded Complexes". Compositio Math. 65, 1988, no. 2, 121-154.

[24] Beilinson, A. A., Bernstein, J. and P. Deligne. *Faisceaux Pervers.* Asterisuqe 100, Volume 1, 1982.

[25] Dwyer, W.G. and D.M. Kan. "Homotopy Theory of Simplicial Groupoids." Nederl. Akad. Wetensch. Indag. Math. 46 (1984) no. 4 379-385.

[26] Serre, J.P. *Cohomologie Galoisienne.* Lecture Notes in Mathematics 5, Springer-Verlag, Berlin and New York, 1964.

[27] Dugger, D., Hollander, S. and D. Isaksen. *Hypercovers and simplicial presheaves.*

[28] Tamsamani, Zouhair. *On non-strict notions of n-category and n groupoid via multisimplicial sets.*