PULLBACKS OF LOCAL OPERATORS AND TIDY MORPHISMS

P. JOHNSTONE

REPORT No. 35, 2000/2001

ISSN 1103-467X
ISRN IML-R- -35-00/01- -SE

INSTITUT MITTAG-LEFFLER
THE ROYAL SWEDISH ACADEMY OF SCIENCES

Pullbacks of Local Operators and Tidy Morphisms

P.T. Johnstone*
Department of Pure Mathematics, University of Cambridge

May 2001

Abstract

We present a new proof that the class of (bounded) tidy geometric morphisms is stable under (bounded) pullback in the 2-category of toposes and geometric morphisms, which is simpler than that given by Moerdijk and Vermeulen [10] in that it does not require any appeal to an axiom of infinity. On the way to it, we investigate a hitherto neglected 'initiality property' of pullbacks of inclusions: we do not have a proof that this property holds in complete generality, but we show that it holds in a great many particular cases — in particular, for all pullbacks of inclusions along bounded morphisms.

Introduction

Suppose one wishes to prove that some property P of geometric morphisms is stable under pullback; that is, whenever

$$\begin{array}{ccc} \mathcal{H} & \xrightarrow{k} & \mathcal{G} \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle g} \\ \mathcal{F} & \xrightarrow{f} & \mathcal{E} \end{array}$$

is a pullback square of toposes and geometric morphisms and f satisfies P, then k also satisfies P. There are two basic 'paradigms' that one can follow in seeking such a proof. In order to give a description of the pullback topos $\mathcal{H}$, one needs to assume that at least one of f and g is bounded (cf. [4], 4.47); but one may choose to place the boundedness assumption on the same morphism f which has the property P under investigation, or on the morphism g along which one seeks to pull it back.

In the first paradigm, one seeks a 'site characterization' of the property P: that is, a property P′ of internal sites (C, T) in a topos such that a bounded morphism $f\colon \mathcal{F} \to \mathcal{E}$ satisfies P iff the topos $\mathcal{F}$ is representable as the topos $\mathbf{Sh}_{\mathcal{E}}(\mathrm{C}, T)$ of $\mathcal{E}$-valued sheaves on an internal site (C, T) satisfying P′. Then one seeks to prove that P′ is 'stable under change of base', in the sense that if (C, T) satisfies it then so does the site $(g^*\mathrm{C}, g^\# T)$ in $\mathcal{G}$ which corresponds to the pullback $k\colon \mathcal{H} \to \mathcal{G}$. A well known example of this paradigm is the proof by Joyal and Tierney [8] that open maps and open surjections are stable under pullback, using the fact that bounded open maps correspond to sites in which every covering family is inhabited.

However, this paradigm has some disadvantages. One is that the process of change of base for sites is rather a delicate one: in particular the 'local character' property of a coverage can be lost in the passage from T to $g^\# T$, and so the site characterization of P must be 'robust enough' to withstand this. Perhaps more seriously, for many properties P, in order to obtain a suitable site characterization one needs to make assumptions (such as closure under finite limits) about the underlying category C of the site which presuppose that the base topos $\mathcal{E}$ has a natural number object, and so the pullback-stability

* *Most of the research in this paper was undertaken while the author was visiting the Institut Mittag-Leffler in Stockholm, as a participant in their Mathematical Logic year 2000–01.*

theorem one obtains is valid only for geometric morphisms between toposes with natural number objects. This is not the case for the Joyal–Tierney result on open maps, mentioned above; but it does apply to their proof (also in [8]) that bounded hyperconnected morphisms (and hence the hyperconnected–localic factorization of an arbitrary bounded morphism) are stable under pullback.

For this reason, it seems worth considering the second paradigm, in which one places the boundedness assumption on the morphism g. In this case it is natural to factor g as

$$\mathbf{Sh}_{\mathcal{E}}(\mathrm{C}, T) \longrightarrow [\mathrm{C}^{\mathrm{op}}, \mathcal{E}] \longrightarrow \mathcal{E}$$

where the first factor is an inclusion, and then to prove the pullback-stability result separately for pullbacks along morphisms of the form $[\mathrm{C}^{\mathrm{op}}, \mathcal{E}] \to \mathcal{E}$ (or $[\mathrm{C}, \mathcal{E}] \to \mathcal{E}$), and for pullbacks along inclusions.

The first of these is almost always trivial, at least if property P can be expressed in terms of properties of f^* and f_* as $\mathcal{E}$-indexed functors: the reason is that any $\mathcal{E}$-indexed category can be canonically extended to a $\mathbf{Cat}(\mathcal{E})$-indexed category, where $\mathbf{Cat}(\mathcal{E})$ is the category of internal categories in $\mathcal{E}$, and the diagram category $[\mathrm{C}, \mathcal{E}]$ may be identified with the full subcategory of $\mathbf{Cat}(\mathcal{E})/\mathrm{C}$ whose objects are discrete opfibrations ([4], 2.15). Moreover, the direct and inverse images of the geometric morphism $[\mathrm{C}, f]$ which is the pullback of f along $[\mathrm{C}, \mathcal{E}] \to \mathcal{E}$ are obtained simply by extending f_* and f^* to $\mathbf{Cat}(\mathcal{E})$-indexed functors, slicing over C and restricting to the subcategory of discrete opfibrations; so, provided we can show that the property P is stable under these constructions, its stability under pullback along $[\mathrm{C}, \mathcal{E}] \to \mathcal{E}$ is immediate. Proving stability of P under pullback along inclusions can be rather more delicate; but, once again, we have a reasonably explicit description of the pullback topos $\mathcal{H}$ in this context (which we shall recall at the beginning of section 2 below), and this can be exploited in many particular cases.

Following this second paradigm, and independently of the work of Joyal and Tierney, the present author in [5] gave a proof that arbitrary open maps and open surjections are stable under bounded pullback. In [6], I also gave a proof that arbitrary hyperconnected morphisms are stable under bounded pullback, which does not require the assumption of a natural number object in the base topos.

More recently, Moerdijk and Vermeulen [10] have investigated proper maps of toposes, which are in a natural sense 'dual' to open maps, and the stronger notion of tidy morphism which is similarly dual to local connectedness. They gave 'first paradigm' proofs of pullback-stability for both these notions; but these proofs depend on site characterizations which start from the assumption that the underlying category C is a pretopos, and in particular closed under finite limits when regarded as an ($\mathcal{E}$-indexed) subcategory of $\mathcal{F}$. They are therefore applicable only to toposes with natural number objects. Moreover, the actual site characterizations (particularly in the tidy case) are extremely delicate, and proving that they are stable under change of base is not at all easy.

It therefore seemed worth investigating whether these pullback-stability theorems could alternatively be proved by the second paradigm. For proper maps this is not difficult, because of the fact that arbitrary hyperconnected morphisms are proper; so, using the result from [6] already mentioned, one can reduce to the case when f is localic (in which case the hard work was already done in Vermeulen's earlier paper [11] on proper maps of locales). However, when I investigated the tidy case, I found that the first step (pullback along $[\mathrm{C}, \mathcal{E}] \to \mathcal{E}$) went straightforwardly as usual, but the second step required a property of pullbacks of inclusions which appeared not to have been noticed previously, and which I was unable to prove in general. On the other hand, this property seems to hold in a great many cases. For want of a more descriptive term, let us call a geometric morphism *good* if all its pullbacks along inclusions have the property in question; we shall show in this paper that all bounded morphisms, all atomic morphisms (whether bounded or not) and all morphisms which can be factored through a Boolean topos are good. Indeed, it seems highly likely that all geometric morphisms are good.

The layout of the paper is as follows. In section 1 we recall some facts about local operators and their sheaves: in particular, the construction of the associated sheaf functor via iterated internal colimits, first introduced in [3]. In section 2 we recall the basic facts about pullbacks of inclusions, and introduce the notion of goodness: the results that we have so far been able to prove about goodness are contained in this section and section 3. Finally, in section 4 we introduce tidy morphisms, and give the 'second paradigm' proof of their stability under bounded pullback, and their associated characterization in terms of Beck–Chevalley conditions: as explained above, these results are already proved via the 'first paradigm' in [10], but our proof has the advantage that it does not require the assumption that the toposes involved have natural number objects.

1 Preliminaries

By a *local operator* on a topos $\mathcal{E}$, we mean what is often called a (Lawvere–Tierney) topology; that is, an endomorphism $j\colon \Omega \to \Omega$ of the subobject classifier satisfying $j\top = \top$, $jj = j$ and $j\wedge = \wedge(j \times j)$ where $\wedge$ is the binary meet operation on Ω. We shall use standard notation for local operators: given such an operator j, we shall write J æ Ω for the subobject classified by j and Ω_j æ Ω for the equalizer of j and the identity. We shall also use the notation $\Xi(D)$, introduced in [4], for the class of monomorphisms in $\mathcal{E}$ whose classifying maps factor through a given subobject D æ Ω; thus $\Xi(J)$ denotes the class of j-dense monomorphisms, and $\Xi(\Omega_j)$ the class of j-closed monomorphisms. We recall the following result ([4], Lemma 3.18):

Lemma 1.1 *Let J be a subobject of Ω in a topos $\mathcal{E}$. Then the classifying map of J æ Ω is a local operator iff $\Xi(J)$ is a saturated multiplicative system of monomorphisms in $\mathcal{E}$ (i.e., it contains all isomorphisms, and contains a composite monomorphism A'' æ A' æ A iff it contains both A'' æ A' and A' æ A).* Λ

Given a local operator j on $\mathcal{E}$, we write $\mathbf{sh}_j(\mathcal{E})$ for the full subcategory of j-sheaves (that is, objects of $\mathcal{E}$ which are orthogonal to j-dense monomorphisms). We recall that the inclusion $\mathbf{sh}_j(\mathcal{E}) \to \mathcal{E}$ is a direct image functor; that is, it has a left adjoint (the *associated j-sheaf functor*) which preserves finite limits. There are many constructions of this left adjoint in the literature; we shall need to make use of the one via directed internal colimits, introduced in [3] (and also described in [4]), and we briefly recall it here.

Let $d\colon 1$ æ J denote the generic j-dense monomorphism, i.e. the factorization of $\top\colon 1$ æ Ω through J æ Ω. Given an object A of $\mathcal{E}$, we write $\hat{A} \to J$ for the object $\Pi_d(A \to 1)$ of $\mathcal{E}/J$; it is easily seen that $\hat{A}$ classifies partial maps with j-dense domain, in the sense that for any object B morphisms $B \to \hat{A}$ correspond bijectively to diagrams

$$\begin{array}{ccc} B' & \longrightarrow & A \\ \downarrow & & \\ B & & \end{array}$$

where B' æ B is dense. Moreover, composing a morphism $B \to \hat{A}$ with $\hat{A} \to J$ corresponds to 'forgetting the morphism $B' \to A$' and remembering only the dense subobject B' æ B.

Using this description, one may easily make $\hat{A} \to J$ into a contravariant internal diagram on the internal poset $\mathbf{J}$ obtained by restricting the canonical order-relation on Ω to J; we write A^+ for the colimit of this internal diagram. There is a canonical morphism $A \to A^+$, obtained by composing the top edge of the pullback square

$$\begin{array}{ccc} A & \longrightarrow & \hat{A} \\ \downarrow & & \downarrow \\ 1 & \xrightarrow{\ d\ } & J \end{array}$$

with the 'colimit cone' $\hat{A} \to A^+$. Moreover, it is easy to see that $A \mapsto A^+$ is a functor $\mathcal{E} \to \mathcal{E}$, and that the canonical morphism just described is natural in A.

The following results are all proved in [3] and in [4], section 3.3:

Lemma 1.2 (i) *A is a j-sheaf iff $A \to A^+$ is an isomorphism.*

(ii) *Any morphism from A to a j-sheaf factors uniquely through $A \to A^+$.*

(iii) *For any object A, A^+ is j-separated.*

(iv) *If A is j-separated, then A^+ is a sheaf.* Λ

Corollary 1.3 *For any object A, the associated j-sheaf of A may be constructed as A^{++}.* Λ

¿From the point of view of this construction, the fact that the associated sheaf functor preserves finite limits is a reflection of the fact that $\mathbf{J}^{\mathrm{op}}$ is a filtered category. Indeed, $\mathbf{J}$ is a meet-semilattice, since the first and third conditions in the definition of a local operator ensure that finite intersections of j-dense subobjects are j-dense.

2 Pullbacks of Local Operators

A geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ is called an *inclusion* if it is equivalent to that induced by a local operator on $\mathcal{E}$ — equivalently, if the direct image functor f_* is full and faithful. Given an arbitrary geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ and local operators j, k on $\mathcal{E}, \mathcal{F}$ respectively, we recall that the composite $\mathbf{sh}_k(\mathcal{F}) \to \mathcal{F} \to \mathcal{E}$ factors (uniquely up to canonical isomorphism) through the inclusion $\mathbf{sh}_j(\mathcal{E}) \to \mathcal{E}$ iff f^* maps j-dense monomorphisms to k-dense monomorphisms, iff the composite f^*J æ $f^*\Omega_{\mathcal{E}} \to \Omega_{\mathcal{F}}$ (whose second factor is the canonical comparison map, i.e. the classifying map of $f^*(\top)$) factors through K æ $\Omega_{\mathcal{F}}$ (cf. [4], 3.47). For any f and j, there is a unique smallest local operator k on $\mathcal{F}$ for which this condition holds, and the choice of this particular k makes the resulting square

$$\begin{array}{ccc} \mathbf{sh}_k(\mathcal{F}) & \longrightarrow & \mathbf{sh}_j(\mathcal{E}) \\ \downarrow & & \downarrow \\ \mathcal{F} & \xrightarrow{\ f\ } & \mathcal{E} \end{array}$$

into a (pseudo-)pullback in the 2-category Top of toposes and geometric morphisms. (N.B.: when we are working in this 2-category, any reference to limits should be interpreted as a reference to pseudo-limits.) We shall loosely refer to this particular k as the 'pullback local operator' corresponding to f and j.

The construction of the smallest local operator k for which K contains a given subobject D æ Ω may be described in various ways: see [4], 3.58 for one possibility. However, none of the methods applicable for an arbitrary D is particularly simple or explicit. It appears that, at least in a very large proportion of cases, when D is the image of a morphism f^*J æ $f^*\Omega_{\mathcal{E}} \to \Omega_{\mathcal{F}}$ as in the previous paragraph, then we can greatly simplify the description: K is simply the upward closure of D in the canonical partial ordering on $\Omega_{\mathcal{F}}$. Note that this is equivalent to saying that the comparison map $f^*\mathbf{J} \to \mathbf{K}$ is an initial functor, where $\mathbf{J}$ and $\mathbf{K}$ denote the internal posets in $\mathcal{E}$ and $\mathcal{F}$ obtained by equipping J and K with their canonical orderings.

Definition 2.1 We shall call a geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ *good* if, for any local operator j on $\mathcal{E}$, the upward closure of the image of the composite f^*J æ $f^*\Omega_{\mathcal{E}} \to \Omega_{\mathcal{F}}$ is classified by a local operator on $\mathcal{F}$ (which is then necessarily the pullback local operator as described earlier, since the subobject classified by a local operator is always upward-closed). We shall say that f is *very good* if we do not even have to take the upward closure, i.e. if the image of $f^*J \to \Omega_{\mathcal{F}}$ is already classified by a local operator for any j.

We do not know any example of a geometric morphism which is not good. Indeed, in the course of this and the next section, we shall show that the class of good geometric morphisms (properly) includes the class of all bounded morphisms. However, we have been unable to find a 'uniform' proof that all morphisms are good: the proofs that we give for goodness of different classes of morphisms employ a wide range of different techniques. Nevertheless, it seems worth putting them all on record, in the hope that one of them may eventually be found to contain the seeds of a uniform proof.

We begin by noting that goodness is stable under composition.

Lemma 2.2 *A composite of good (resp. very good) geometric morphisms is good (resp. very good).*

Proof We deal with the first case; the second is similar but easier. Let $f\colon \mathcal{F} \to \mathcal{E}$ and $g\colon \mathcal{G} \to \mathcal{F}$ be good; let j be a local operator on $\mathcal{E}$, and let k, m be the pullback local operators on $\mathcal{F}, \mathcal{G}$ respectively. By assumption, K æ $\Omega_{\mathcal{F}}$ is the upward closure of the image of f^*J æ $f^*\Omega_{\mathcal{E}} \to \Omega_{\mathcal{F}}$; but all constructions involved in forming this upward closure (i.e., finite limits and image factorizations) are preserved by g^*, so g^*K is the upward closure of the image of g^*f^*J æ $g^*f^*\Omega_{\mathcal{E}} \to g^*\Omega_{\mathcal{F}}$. If we further compose the latter map with $g^*\Omega_{\mathcal{F}} \to \Omega_{\mathcal{G}}$ and take the upward closure of its image, we obtain a subobject (M', say) of $\Omega_{\mathcal{G}}$ whose pullback along $g^*\Omega_{\mathcal{F}} \to \Omega_{\mathcal{G}}$ is upward-closed (since this map is order-preserving) and hence contains g^*K. Hence M' contains the upward closure of the image of $g^*K \to \Omega_{\mathcal{G}}$, which equals M by assumption; but the reverse inclusion is trivial, so $M' = M$ and in particular its classifying map is a local operator. Λ

Goodness also 'descends' along open surjections:

Lemma 2.3 *Let $f\colon \mathcal{F} \to \mathcal{E}$ and $g\colon \mathcal{G} \to \mathcal{F}$ be geometric morphisms such that g is an open surjection and the composite fg is good. Then f is good.*

Proof As before, let j be a local operator on $\mathcal{E}$, and write k, m for the pullback local operators on $\mathcal{F}$ and $\mathcal{G}$ respectively. We know that the composite $g^*f^*\mathbf{J} \to g^*\mathbf{K} \to \mathbf{M}$ is an initial functor; but since g is open the comparison map $g^*\Omega_{\mathcal{F}} \to \Omega_{\mathcal{G}}$ is monic (cf. [5]) and hence order-reflecting, and so is $g^*K \to M$. It follows easily that $g^*f^*\mathbf{J} \to g^*\mathbf{K}$ is initial; but since g^* is faithful, it reflects the property of being an initial functor. Λ

Now we start on our repertoire of proofs that particular classes of morphisms are (very) good.

Lemma 2.4 *Any geometric morphism f whose inverse image preserves Ω (i.e., for which the comparison map $f^*\Omega_{\mathcal{E}} \to \Omega_{\mathcal{F}}$ is an isomorphism) is very good.*

Proof If the comparison map is an isomorphism, then the image of $f^*J \to \Omega_{\mathcal{F}}$ is simply f^*J itself; and since f^* preserves pullbacks, its classifying map is identified under the isomorphism with $f^*(j)\colon f^*\Omega_{\mathcal{E}} \to f^*\Omega_{\mathcal{E}}$. It is immediate that the latter inherits the commtutativity conditions for a local operator from j. Λ

Lemma 2.4 applies in particular if f is an atomic geometric morphism, i.e. if its inverse image functor is logical. It also applies to the inclusion $\mathbf{sh}_{\neg\neg}(\mathcal{E}) \to \mathcal{E}$ for any $\mathcal{E}$; but the latter case is also covered by the next result.

Lemma 2.5 *If $\mathcal{F}$ is Boolean, then any geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ is very good.*

Proof In a Boolean topos, Ω is simply the coproduct $1 \amalg 1$, from which it follows easily that any subobject of Ω which contains $\top\colon 1$ æ Ω is classified by a local operator. But this condition clearly holds for the image of $f^*J \to \Omega_{\mathcal{F}}$, for any local operator j on $\mathcal{E}$. Λ

As a companion for Lemma 2.5, we note

Lemma 2.6 *If $\mathcal{E}$ is Boolean, then any geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ is good.*

Proof In a Boolean topos, every local operator is open ([4], Ex. 5.1). But if j is the open local operator associated with a subobject U æ 1, then J is simply the upward closure of the element $u\colon 1 \to \Omega$ which is the classifying map of U æ 1. If we apply f^* to this, and take the upward closure of its image, we clearly get the upward closure of the element of $\Omega_{\mathcal{F}}$ which classifies f^*U æ 1 — that is, the open local operator corresponding to this subobject of 1. Λ

It is easy to find examples to show that we cannot strengthen 'good' to 'very good' in Lemma 2.6. Combining the last two results by means of 2.2, we note that any geometric morphism which can be factored through a Boolean topos is good.

If f is itself the inclusion of a sheaf subtopos $\mathbf{sh}_m(\mathcal{E})$ in $\mathcal{E}$, then we may give an explicit description of the class $\Xi(D)$, where D is the image of f^*J æ $f^*\Omega \to \Omega_m$: a monomorphism B' æ B in $\mathbf{sh}_m(\mathcal{E})$ belongs to $\Xi(D)$ iff it can be factored in $\mathcal{E}$ as B' æ B'' æ B where B' æ B'' is j-dense and B'' æ B is m-dense. To see this, note that D is simply the m-closure in Ω_m of the image (I, say) in $\mathcal{E}$ of J æ Ω ı Ω_m, since f^* preserves images; and we may obtain B'' by pulling back the m-dense monomorphism I æ D along the classifying map of B' æ B. Using this dscription, we may now prove

Lemma 2.7 *Closed inclusions are very good.*

Proof If m is the closed local operator corresponding to a subobject U æ 1, then a subobject A' æ A is m-closed iff it contains the subobject $A \times U$ æ A; hence any subobject containing an m-closed one is m-closed. Since all monomorphisms in $\mathbf{sh}_m(\mathcal{E})$ are m-closed, the description of $\Xi(D)$ simplifies still further: it consists of those monomorphisms in $\mathbf{sh}_m(\mathcal{E})$ which are j-dense. As such, it clearly satisfies the conditions of Lemma 1.1. Λ

3 Bounded Morphisms are Good

The results of the previous section are all 'elementary' in the sense that their proofs do not depend on a representation of $\mathcal{F}$ as the topos of sheaves on an internal site in $\mathcal{E}$. However, for our next result, we shall need to assume such a representation. We recall that a geometric morphism $f: \mathcal{F} \to \mathcal{E}$ is said to be *locally connected* if f^* has an $\mathcal{E}$-indexed left adjoint $f_!$, as well as its (automatically $\mathcal{E}$-indexed) right adjoint f_*. Local connectedness of f says that every object of $\mathcal{F}$ has a canonical representation as an $\mathcal{E}$-indexed coproduct of 'connected objects', the latter being defined as those B for which $f_!B \cong 1$. If f is also bounded (that is, $\mathcal{F}$ is representable as a topos of sheaves on an internal site in $\mathcal{E}$), then we may choose an internal site (C, T) for it such that C consists of connected objects of $\mathcal{F}$; then the coverage T has the property that every T-covering sieve on an object U of C is connected (as a subcategory of C/U). A site with this property is naturally called a *locally connected site.*

Lemma 3.1 *Bounded locally connected morphisms are good.*

Proof For clarity, in what follows, we shall use set-theoretic notation and terminology for objects and morphisms of $\mathcal{E}$; but our arguments are constructive, and a reader familiar with the interpretation of higher-order logic in a topos should have no difficulty in translating them into the internal logic of $\mathcal{E}$.

Suppose $\mathcal{F} \simeq \mathbf{Sh}_{\mathcal{E}}(\mathrm{C}, T)$, where (C, T) is a locally connected site, and let j be a local operator on $\mathcal{E}$. Since covers in T are connected, it is easy to see that any constant diagram of shape C^{op} is a T-sheaf, and hence in particular that $f^*\Omega_{\mathcal{E}}$ and f^*J are simply the appropriate constant diagrams on C^{op}. On the other hand, $\Omega_{\mathcal{F}}$ is the internal diagram corresponding to the functor $\mathrm{C}^{\mathrm{op}} \to \mathcal{E}$ which sends an object U of C to the set of T-closed sieves on U, where a sieve S is defined to be T-*closed* iff, given a morphism $a: V \to U$ and a T-covering sieve R on V such that the composite ab belongs to S for all $b \in R$, we necessarily have $a \in S$.

The U-component of the comparison map $f^*\Omega_{\mathcal{E}} \to \Omega_{\mathcal{F}}$ sends a truth-value $p \in \Omega_{\mathcal{E}}$ to the sieve $\{a: V \to U \mid p\}$ (note that such sieves are indeed T-closed, since T-covering sieves are inhabited). Hence, if we write K for the upward closure of the image of f^*J, then a sieve R belongs to $K(U)$ iff there exists a J-dense truth-value p such that $(p \leq [\![a \in R]\!])$ holds for all a with codomain U; but since R is a sieve (that is, $[\![a \in R]\!] \leq [\![ab \in R]\!]$ for all composable pairs a, b), this is equivalent to saying that the truth-value $[\![1_U \in R]\!]$ is j-dense. It now follows easily that a subobject B' æ B in $\mathcal{F}$ belongs to the class $\Xi(K)$ iff the sentence

$$(\forall U \in C_0)(\forall x \in B(U))([\![x \in B'(U)]\!] \in J)$$

is valid in $\mathcal{E}$ — that is, each $B'(U)$ æ $B(U)$ is a j-dense monomorphism. Using this, it is trivial to verify that $\Xi(K)$ satisfies the conditions of 1.1, and so the classifying map of K is a local operator. Λ

It is conceivable that a direct proof along the lines of 3.1 could be given to show that more general bounded morphisms are good; but the extra complications arising from the fact that, in general, $f^*\Omega_{\mathcal{E}}$ and f^*J are not constant diagrams (but rather the associated T-sheaves of constant diagrams) would be considerable. Fortunately, we do not have to wrestle with these complications, since we shall see shortly that the goodness of arbitrary bounded morphisms follows from 2.2, 2.7 and 3.1.

Following M. Bunge and J. Funk [1], we say that a geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ is *totally connected* if it has a right adjoint when considered as a morphism $f \to 1_{\mathcal{E}}$ in the 2-category $\mathsf{Top}/\mathcal{E}$. It is easy to see that this happens iff f is connected (that is, f^* is full and faithful) and locally connected, and the left adjoint $f_!$ of f^* preserves finite limits (so that f^* is the *direct* image of the geometric morphism $d\colon \mathcal{E} \to \mathcal{F}$ which is right adjoint to f, and $f_!$ is its inverse image).

In [7], I. Moerdijk and the present author investigated local morphisms of toposes, which are those $f\colon \mathcal{F} \to \mathcal{E}$ which have left adjoints in $\mathsf{Top}/\mathcal{E}$. Among the results established in that paper, we observed that, by the use of Artin glueing along a direct image functor, it is possible to factor an arbitrary geometric morphism as an open inclusion followed by a local morphism (see [7], 1.2(e)). Since totally connected morphisms have properties dual to those of local morphisms, it should come as no surprise that we may factor an arbitrary geometric morphism as a closed inclusion followed by a totally connected morphism, nor that this factorization is constructed by Artin glueing along an inverse image functor; but since this result does not (as far as we know) appear in the existing literature, we give the proof here.

Lemma 3.2 *Any geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ may be factored (up to isomorphism) as a composite gh, where g is totally connected and h is a closed inclusion.*

Proof Let $\mathcal{G}$ be the topos obtained by glueing along f^*, that is the comma category $(\mathcal{F}\downarrow f^*)$ whose objects are triples (A, B, α) with $A \in \mathrm{ob}\,\mathcal{E}$, $B \in \mathrm{ob}\,\mathcal{F}$ and $\alpha\colon B \to f^*A$ in $\mathcal{F}$. It is well known that we have complementary open and closed inclusions $d\colon \mathcal{E} \to \mathcal{G}$ and $h\colon \mathcal{F} \to \mathcal{G}$: the inverse images of these two inclusions are the obvious projections, and we have $d_*(A) = (A, f^*A, 1_{f^*A})$ and $h_*(B) = (1, B, (B \to 1))$. That much is true if we replace f^* by any finite-limit-preserving functor $\mathcal{E} \to \mathcal{F}$; but we may use the right adjoint of f^* to construct a further right adjoint $d^\#$ for d_*. Specifically, $d^\#(A, B, \alpha)$ is the pullback of

$$\begin{array}{ccc} & & f_*(B) \\ & & \Big\downarrow{\scriptstyle f_*(\alpha)} \\ A & \xrightarrow{\eta_A} & f_*f^*(A) \end{array}$$

in $\mathcal{E}$, where η is the unit of $(f^* \dashv f_*)$. Now if we define g_*, g^* and $g_!$ to be $d^\#$, d_* and d^* respectively, it is straightforward to verify that g is a totally connected morphism $\mathcal{G} \to \mathcal{E}$. Also, the composite h^*g^* is equal to f^* by definition, and it is easy to see that g_*h_* is at least canonically isomorphic to f_*; so we have constructed the required factorization of f. Λ

Corollary 3.3 *In addition to the hypotheses of* 3.2, *suppose that f is bounded. Then so is g.*

Proof It would be possible to prove this by constructing a site of definition for $\mathcal{G}$ over $\mathcal{E}$ from one for $\mathcal{F}$, but we shall instead work from the original definition of a bounded morphism as one which possesses a bound. Recall that, given $f\colon \mathcal{F} \to \mathcal{E}$, a *bound* (also called an object of generators) for $\mathcal{F}$ over $\mathcal{E}$ is an object G of $\mathcal{F}$ such that, for every object B of $\mathcal{F}$, we can find a diagram of the form

$$\begin{array}{ccc} \overline{B} & \rightarrowtail & f^*A \times G \\ \Big\downarrow & & \\ B & & \end{array}$$

where A is an object of $\mathcal{E}$. Suppose given such an object of $\mathcal{F}$; then we claim that $h_*(G) = (1, G, (G \to 1))$ is a bound for $\mathcal{G}$ over $\mathcal{E}$.

To prove this, suppose given an arbitrary object (A, B, α) of $\mathcal{G}$, and choose a diagram

$$\begin{array}{ccc} \overline{B} & \stackrel{m}{\rightarrowtail} & f^*C \times G \\ \Big\downarrow{\scriptstyle e} & & \\ B & & \end{array}$$

where we may assume (on replacing C by $C \amalg 1$, if necessary) that $C \to 1$ is an epimorphism. Now we have a diagram

$$\begin{array}{ccc} (A \times C, \overline{B}, (\alpha e, \pi_1 m)) & \stackrel{(1, (\alpha e, m))}{\rightarrowtail} & (A \times C, f^*(A \times C) \times G, \pi_1) \cong g^*(A \times C) \times h_*(G) \\ \Big\downarrow{\scriptstyle (\pi_1, e)} & & \\ (A, B, \alpha) & & \end{array}$$

of the required form in $\mathcal{G}$. Λ

Corollary 3.4 *Any bounded geometric morphism is good.*

Proof Let f be a bounded geometric morphism. By 3.2 and 3.3, we may factor it as gh, where g is totally connected (in particular, locally connected) and bounded, and h is a closed inclusion. Then g is good by 3.1, and h is good by 2.7; so the result follows from 2.2. Λ

4 Pullbacks of Tidy Morphisms

A geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ is said to be *tidy* if its direct image functor f_* preserves filtered colimits; however, in order to state this in a suitably '$\mathcal{E}$-indexed' form, we have to interpret 'filtered colimits' as meaning 'colimits indexed by filtered internal categories in arbitrary slices of $\mathcal{E}$'. Formally, we have

Definition 4.1 A geometric morphism $f\colon \mathcal{F} \to \mathcal{E}$ is *tidy* if, for any object A of $\mathcal{E}$ and any filtered internal category $\mathbb{I}$ in $\mathcal{E}/A$, the canonical 2-cell in the diagram

$$\begin{array}{ccc} [f^*\mathbb{I}, \mathcal{F}/f^*A] & \xrightarrow{[\mathbb{I}, f/A]_*} & [\mathbb{I}, \mathcal{E}/A] \\ \Big\downarrow{\scriptstyle \varinjlim} & \Downarrow & \Big\downarrow{\scriptstyle \varinjlim} \\ \mathcal{F}/f^*A & \xrightarrow{(f/A)_*} & \mathcal{E}/A \end{array}$$

is an isomorphism.

Tidy morphisms were first studied by T. Lindgren [9], who called them 'proper maps', in the context of Beck–Chevalley conditions for pullback squares of geometric morphisms. (The work of K.R. Edwards [2] was an important precursor.) Because the phrase 'proper map' has an alternative, less restrictive, meaning, I. Moerdijk and J. Vermeulen [10] chose to introduce the name 'tidy morphism' for the concept defined in 4.1, and we shall follow their terminology.

The connection with Beck–Chevalley conditions is as follows. Given a commutative square

$$\begin{array}{ccc} \mathcal{H} & \xrightarrow{k} & \mathcal{G} \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle g} \\ \mathcal{F} & \xrightarrow{f} & \mathcal{E} \end{array}$$

of toposes and geometric morphisms, the natural isomorphism $h^*f^* \cong k^*g^*$ transposes across the adjunctions $(f^* \dashv f_*)$ and $(k^* \dashv k_*)$ to yield a canonical natural transformation $g^*f_* \to k_*h^*$. We say that the square *satisfies the Beck–Chevalley condition* if this transformation is an isomorphism. In general (though there are some exceptions) one should expect a commutative square which satisfies the Beck–Chevalley condition to be a pullback. We shall say that a geometric morphism f is a *right Beck–Chevalley morphism* if, for any bounded morphism g whose codomain matches that of f, the pullback square of f against g satisfies the Beck–Chevalley condition; and we shall say that f is a *stable right Beck–Chevalley morphism* if every pullback of f along a bounded morphism is a right Beck–Chevalley morphism. (There is a dual notion of (stable) left Beck–Chevalley morphism, in which the rôles of f and g are reversed; Lindgren also investigated these, and showed that the stable left Beck–Chevalley morphisms are exactly the locally connected morphisms.)

If I is a filtered internal category in $\mathcal{E}/A$ for some A, then the colimit functor $[\mathrm{I}, \mathcal{E}/A] \to \mathcal{E}/A$ is the inverse image of a geometric morphism (over $\mathcal{E}/A$, and hence) over $\mathcal{E}$; following Moerdijk and Vermeulen, we shall denote this morphism by ∞. It is further easy to see that, for any $f : \mathcal{F} \to \mathcal{E}$, the square

$$\begin{array}{ccc} \mathcal{F}/f^*A & \xrightarrow{f/A} & \mathcal{E}/A \\ {\scriptstyle \infty}\downarrow & & \downarrow{\scriptstyle \infty} \\ [f^*\mathrm{I}, \mathcal{F}/f^*A] & \xrightarrow{[\mathrm{I}, f/A]} & [\mathrm{I}, \mathcal{E}/A] \end{array}$$

is a pullback, and that its horizontal edges are pullbacks of f. So we deduce:

Lemma 4.2 *Any stable right Beck–Chevalley morphism is tidy.* Λ

To prove the converse of 4.2, we clearly need to prove (a) that tidy morphisms are stable under pullback and (b) that they are right Beck–Chevalley morphisms. Lindgren [9] was the first to seek such a converse; but the proof he found was non-constructive (involving a transfinite iteration), and so only applicable to Grothendieck toposes (that is, toposes defined and bounded over the classical topos of sets). Moerdijk and Vermeulen [10] gave a proof of the converse which is constructive, but which relies on a site characterization of tidy morphisms; the latter in turn requires the underlying category of the site to have finite limits (and colimits), and so the Moerdijk–Vermeulen proof is applicable only to bounded morphisms between toposes with natural number objects.

The present author's desire to eliminate these restrictions led him to seek a proof using the 'second paradigm' mentioned in the Introduction: such a proof does indeed exist, and it is (in his opinion, at least) appreciably simpler than the Moerdijk–Vermeulen one. Unfortunately, we have not entirely eliminated the restrictions on the class of geometric morphisms to which it applies: for Lemma 4.4 below requires the hypothesis that f is good in the sense of Definition 2.1, and we have not succeeded in proving that all geometric morphisms are good. However, as we have seen in the last section, all bounded geometric morphisms are good; and since our proof has no need of natural number objects, it does apply to a strictly larger class of morphisms than the Moerdijk–Vermeulen one.

As usual with a 'second paradigm' pullback-stability proof, the first step (pulling back along a morphism of the form $[\mathrm{C}, \mathcal{E}] \to \mathcal{E}$) is essentially trivial.

Lemma 4.3 *Let $f : \mathcal{F} \to \mathcal{E}$ be a tidy geometric morphism, and let* C *be an internal category in $\mathcal{E}$. Then the induced morphism* $[\mathrm{C}, f] : [f^*\mathrm{C}, \mathcal{F}] \to [\mathrm{C}, \mathcal{E}]$ *is tidy, and the pullback square*

$$\begin{array}{ccc} [f^*\mathrm{C}, \mathcal{F}] & \xrightarrow{[\mathrm{C}, f]} & [\mathrm{C}, \mathcal{E}] \\ \downarrow & & \downarrow \\ \mathcal{F} & \xrightarrow{f} & \mathcal{E} \end{array}$$

satisfies the Beck–Chevalley condition.

Proof Since every slice category of $[\mathrm{C}, \mathcal{E}]$ is of the form $[\mathrm{D}, \mathcal{E}]$ for some internal category D, it suffices to deal with filtered internal categories in $[\mathrm{C}, \mathcal{E}]$ itself, rather than slices thereof. Given such a category I, we let J denote its underlying category in $\mathcal{E}/C_0$ (where C_0 denotes the object of objects of C); then J is also filtered, since the forgetful functor $[\mathrm{C}, \mathcal{E}] \to \mathcal{E}/C_0$ preserves the validity of the first-order formulae which assert that a category is filtered. Now we may form the cube

$$\begin{array}{ccccccc} [f^*\mathrm{I}, [f^*\mathrm{C}, \mathcal{F}]] & & \longrightarrow & & [\mathrm{I}, [\mathrm{C}, \mathcal{E}]] & & \\ \downarrow & \searrow & & & \downarrow & \searrow & \\ & & [f^*\mathrm{C}, \mathcal{F}] & \longrightarrow & & & [\mathrm{C}, \mathcal{E}] \\ [f^*\mathrm{J}, \mathcal{F}/f^*C_0] & & \downarrow & \longrightarrow & [\mathrm{J}, \mathcal{E}/C_0] & & \downarrow \\ & \searrow & & & & \searrow & \\ & & \mathcal{F}/f^*C_0 & & \longrightarrow & & \mathcal{E}/C_0 \end{array}$$

where the horizontal arrows are all induced by f_*, the diagonal arrows are colimit functors and the vertical arrows are forgetful functors. It is easily verified that each of the vertical faces commutes up to isomorphism, and the forgetful functors reflect isomorphisms; so, from the fact that the bottom face commutes up to isomorphism, we deduce that the top face does so too. The Beck–Chevalley condition is easily verified by direct calculation, without using the tidiness of f (which is not surprising, since $[\mathrm{C}, \mathcal{E}] \to \mathcal{E}$ is locally connected). Λ

Lemma 4.4 *Let $f : \mathcal{F} \to \mathcal{E}$ be a good tidy morphism; let j be a local operator on $\mathcal{E}$, and let j' be the pullback local operator on $\mathcal{F}$. Then the square*

$$\begin{array}{ccc} \mathbf{sh}_{j'}(\mathcal{F}) & \xrightarrow{k} & \mathbf{sh}_j(\mathcal{E}) \\ \downarrow h & & \downarrow g \\ \mathcal{F} & \xrightarrow{f} & \mathcal{E} \end{array}$$

satisfies the Beck–Chevalley condition.

Proof The functor k_* is simply the restriction of f_* to sheaves; so the Lemma is equivalent to the statement that f_* commutes up to isomorphism with the associated sheaf functors for j and j'. But, as we saw in 1.3, the associated j-sheaf of an object A may be constructed as A^{++}, where A^+ is the colimit of an internal diagram $\hat{A}$ of shape $\mathbf{J}^{\mathrm{op}}$. So it suffices to show that the diagram

$$\begin{array}{ccc} \mathcal{F} & \xrightarrow{f_*} & \mathcal{E} \\ \Big\downarrow + & & \Big\downarrow + \\ \mathcal{F} & \xrightarrow{f_*} & \mathcal{E} \end{array}$$

commutes up to isomorphism. As we observed at the end of section 1, $\mathbf{J}^{\mathrm{op}}$ is a filtered internal category; hence this is almost immediate from the definition of a tidy morphism. The only reason why it is not quite immediate is that, to compute the functor $+$ in $\mathcal{F}$, we have to form a colimit over the internal category $\mathbf{J}'^{\mathrm{op}}$, rather than $f^*\mathbf{J}^{\mathrm{op}}$. However, let B be an arbitrary object of $\mathcal{F}$, and form the pullback

$$\begin{array}{ccc} C & \longrightarrow & \hat{B} \\ \Big\downarrow & & \Big\downarrow \\ f^*J & \longrightarrow & J' \end{array} \qquad ;$$

we may regard C as a diagram of shape $f^*\mathbf{J}^{\mathrm{op}}$, and $[\mathbf{J}^{\mathrm{op}}, f]_*$ sends it to the pullback of f_*C along the unit map $J \to f_*f^*J$, or equivalently to the pullback of $f_*(\hat{B})$ along the transpose of $f^*J \to J'$. Now it is easily verified from the description of $\hat{A}$ given in section 1 that

$$\begin{array}{ccc} \widehat{f_*B} & \longrightarrow & f_*(\hat{B}) \\ \Big\downarrow & & \Big\downarrow \\ J & \longrightarrow & f_*J' \end{array}$$

is a pullback, since morphisms from an arbitrary object A to the pullback correspond to morphisms $f^*A \to \hat{B}$, together with a representation of the j'-dense domain of the corresponding partial map $f^*A \rightharpoondown B$ in the form f^*A' æ f^*A, where A' æ A is j-dense. So we have

$$f_*(B)^+ = \varinjlim(\widehat{f_*B}) \cong f_*(\varinjlim(C)) \cong f_*(\varinjlim(\hat{B})) = f_*(B^+)$$

using the facts that f is tidy and $f^*\mathbf{J} \to \mathbf{J}'$ is initial. Λ

We may now prove our main result. We say a geometric morphism f is *stably good* if the pullback of f along any bounded morphism is good; of course, since pullbacks of bounded morphisms are bounded, it follows from 3.4 that bounded morphisms are stably good.

Theorem 4.5 *Suppose given a pullback square*

$$\begin{array}{ccc} \mathcal{H} & \xrightarrow{k} & \mathcal{G} \\ \Big\downarrow h & & \Big\downarrow g \\ \mathcal{F} & \xrightarrow{f} & \mathcal{E} \end{array}$$

in Top, *where f is tidy and stably good, and g is bounded. Then k is tidy, and the Beck–Chevalley natural transformation $g^*f_* \to k_*h^*$ is an isomorphism.*

Proof By 4.3, we may reduce to the case when g is an inclusion; and the Beck–Chevalley condition for this case is proved in 4.4. So it remains to prove that k is tidy, which we shall do by showing that it is a stable right Beck–Chevalley morphism.

Suppose given a bounded morphism $r\colon \mathcal{G}' \to \mathcal{G}$. Then the composite gr is also bounded, so we may factor it as

$$\mathcal{G}' \xrightarrow{\;g'\;} \mathcal{E}' \xrightarrow{\;p\;} \mathcal{E}$$

where g' is an inclusion and $\mathcal{E}' \simeq [\mathrm{C}, \mathcal{E}]$ for some internal category C in $\mathcal{E}$. Now we may form the cube

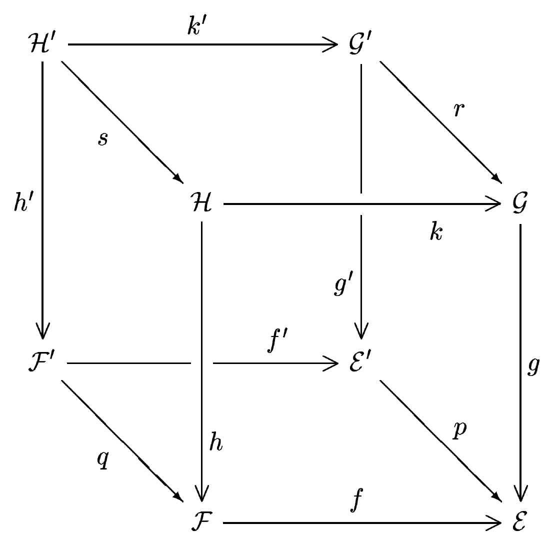

in which all faces with horizontal edges are pullbacks, and the other two faces commute. Moreover, the Beck–Chevalley condition holds for the bottom face by 4.3, and for the front and back faces by 4.4, so an easy diagram-chase shows that the canonical natural transformation $r^*k_*h^* \to k'_*s^*h^*$ is an isomorphism. But h is an inclusion, so h^* is essentially surjective on objects; hence $r^*k_* \to k'_*s^*$ is an isomorphism.

Thus we have shown that k is a right Beck–Chevalley morphism. But the same diagram shows that any pullback of k along a bounded morphism is expressible as a pullback of a (stably good) tidy morphism along an inclusion; so it too is a right Beck–Chevalley morphism. Hence k is a stable right Beck–Chevalley morphism, and so by 4.2 it is tidy. Λ

Corollary 4.6 *A stably good geometric morphism is tidy iff it is a stable right Beck–Chevalley morphism.*

Proof One direction is 4.2; the converse follows from 4.5. Λ

References

[1] M.C. Bunge and J. Funk, Spreads and the symmetric topos, *J. Pure Appl. Alg.* 113 (1996), 1–38.

[2] K.R. Edwards, *Relative finiteness and the preservation of filtered colimits*, Ph.D. thesis, University of Chicago (1980).

[3] P.T. Johnstone, The associated sheaf functor in an elementary topos, *J. Pure Appl. Alg.* 4 (1974), 231–242.

[4] P.T. Johnstone, *Topos Theory*, L.M.S. Monographs no. 10 (Academic Press, 1977).

[5] P.T. Johnstone, Open maps of toposes, *Manuscripta Math.* 31 (1980), 217–247.

[6] P.T. Johnstone, Factorization theorems for geometric morphisms, I, *Cahiers Top. Géom. Diff.* 22 (1981), 3–17.

[7] P.T. Johnstone and I. Moerdijk, Local maps of toposes, *Proc. London Math. Soc.* (3) 58 (1989), 281–305.

[8] A. Joyal and M. Tierney, *An Extension of the Galois Theory of Grothendieck*, Memoirs Amer. Math. Soc. 309 (1984).

[9] T. Lindgren, Ph.D. thesis, Rutgers University (1986).

[10] I. Moerdijk and J.J.C. Vermeulen, *Proper Maps of Toposes*, Memoirs Amer. Math. Soc. 705 (2000).

[11] J.J.C. Vermeulen, Proper maps of locales, *J. Pure Appl. Alg.* 92 (1994), 79–107.