

MACHINE SPACE I: WEAK EXPONENTIALS AND QUANTIFICATION OVER COMPACT SPACES

PETER F. FAUL AND GRAHAM MANUELL

ABSTRACT: Topology may be interpreted as the study of verifiability, where opens correspond to semi-decidable properties. In this paper we make a distinction between verifiable properties themselves and processes which carry out the verification procedure. The former are simply opens, while we call the latter *machines*. Given a frame presentation $\mathcal{O}X = \langle G \mid R \rangle$ we construct a space of machines Σ^{Σ^G} whose points are given by formal combinations of basic machines corresponding to generators in G . This comes equipped with an ‘evaluation’ map making it a weak exponential for Σ^X .

When it exists, the true exponential Σ^X occurs as a retract of machine space. We argue this helps explain why some spaces are exponentiable and others not. We then use machine space to study compactness by giving a purely topological version of Escardó’s algorithm for universal quantification over compact spaces in finite time. Finally, we relate our study of machine space to domain theory and domain embeddings.

MATH. SUBJECT CLASSIFICATION (2010): 54B20, 54C35, 06D22, 06B35, 54D30.

1. Introduction

Topology may be viewed as the logic governing *verifiability*. A property P is said to be *verifiable* if whenever an element satisfies P , this can be established by finite means. Notably absent are any constraints on elements which do not satisfy P — that is, you need not be able to verify this fact. This perspective is due to Smythe in [21, 22] and has been expanded on subsequently in [28] and various other papers discussed below. In this view *open sets* are interpreted as collections of points satisfying verifiable properties — with intersection and union corresponding to logical conjunction and disjunction, respectively. This viewpoint motivates the axioms of a topological space and continuous functions. It is possible to then derive interpretations of various other topological concepts in terms of verifiability. In this paper we make use of the *pointfree* approach which emphasises the verifiable properties P over the collections of points which

Received September 26, 2022.

The author acknowledges financial support from the Centre for Mathematics of the University of Coimbra (UIDB/00324/2020, funded by the Portuguese Government through FCT/MCTES).

satisfy them. The analogue of topological spaces from this perspective are called *locales*, though we will often simply call these ‘spaces’ when speaking intuitively.

It can be instructive to imagine that for each open in a space, there exists some machine or program which carries out the verification procedure. Each machine takes points of the space as input and will either halt (in finite time) or run forever, depending on whether the point belongs to the associated open or not. Here we see an example of the analogy between verifiability in topology and semi-decidability in computability theory [21, 5]. This link is also exhibited by the compactness algorithm, which we discuss later.

We might be tempted to identify machines with opens, which can in turn be seen to correspond to continuous maps from X into the Sierpiński space Σ . Since we are thinking of these machines as ‘real things’ we can interact with, we can expect there to be natural verifiable properties concerning the machines themselves (for example, does the machine halt on some given point) yielding a *space of machines* for each space X . However, if X is not locally compact, the exponential object Σ^X , which we would think of as the space of opens, does not exist and so we are forced to distinguish between machines and opens — unlike functions from X to Σ , machines are not extensional. Moreover, since spaces of machines should always exist, machines are better than opens for some purposes.

In order to formalise these ideas, we describe an explicit construction of a space of machines. We replace the (nonexistent) space of opens Σ^X with a certain *weak exponential*. Our starting point is to fix a presentation for X with generators G . (From the classical perspective these can be viewed as a set of subbasic opens G for the space X .) Since G is a set, the space Σ^G exists. The space X embeds naturally into Σ^G and so the opens of X are restrictions of opens in Σ^G . We will view distinct opens of Σ^G which restrict to the same open U in X as distinct machines which accept (i.e. halt on) precisely the elements of U . A crucial point is that the space Σ^G is locally compact and we can take Σ^{Σ^G} to be the space of machines of X . This space may be thought of as a reasonably canonical weak exponential associated to X with base Σ . More concretely, the points of Σ^{Σ^G} can be thought of as formal joins of formal meets of the generators, which are represented by certain programs that run the ‘basic machines’ from G in parallel.

We then relate machine space to Σ^X , when the latter exists. Under the interpretation so far described an open merely represents a verifiable property

in an abstract sense, whereas a machine is some process that concretely semi-decides memberships of elements. Given an open and a point in a general space X there is *not* an obvious way to verify that the point lies in the open. This helps explain the fact that for general spaces the collection of opens equipped with the Scott topology does not have a continuous evaluation map. We could however expect the evaluation map to be continuous if there were some way to associate a machine to each open. Indeed, we show that when X is locally compact the canonical quotient map of machine space onto Σ^X always has a section, allowing a continuous assignment of opens to machines which represent them.

One way in which machine space is useful is in understanding compactness. From the perspective of verifiability a *compact space* is intuitively a space that can be universally quantified over in finite time [25, 6]. More explicitly, if P is some verifiable property and K is some compact space then the question of whether all the members of K satisfy P is verifiable. That it is sometimes possible to universally quantify over infinite spaces is especially interesting and captures the intuitive idea that ‘compact spaces behave like finite sets’. This idea can be formalised using hyperdoctrines [14, 17, 15], but one might also consider an alternative approach involving the space of opens to have some appeal.

If K is a compact, locally compact space, then we can check if a verifiable property holds on all of K and so we should be able to verify whether the corresponding open in Σ^K is equal to the largest element of Σ^K , namely K itself. In other words, K is compact if and only if the singleton $\{K\}$ is open in Σ^K . See [6]. Of course, this viewpoint is inapplicable outside of the locally compact case, since the exponential Σ^K will not exist.

As might be expected, this can be extended to the general setting by replacing Σ^K with machine space. This perspective on the universal quantification is essentially due to Escardó [5]. Instead of asking if an open of K equals K , we ask if an open in Σ^G *contains* K — that is, if the machine accepts all points of K . By the Hofmann–Mislove theorem [11], there is an associated open corresponding to all of the machines which cover K , which plays the same role as the singleton $\{K\}$ in the previous approach.

Our contribution is obtaining this open via an explicit algorithm for universally quantifying over a compact space. This is a very general and purely topological version of Escardó’s algorithm [5, 7] for universal quantification over Cantor

space. (See also [8] which discusses universal quantification over data types in the setting of higher type computation.)

Finally, we discuss some links with domain theory, give some concrete examples and show how the original algorithm for quantification over Cantor space can be recovered from ours.

In a later paper we intend to explore the links between machine space and powerlocales.

2. Background

2.1. The verifiable interpretation of topology. We begin by discussing the interpretation of topology in terms of verifiable properties. For more details, see [22, 28].

Let X be some collection of points and let us consider what properties the set $\mathcal{O}X$ of verifiable propositions about X should have. Given two proposition U and V we can verify their conjunction $U \wedge V$ by simply verifying U and V in turn. Furthermore, the constantly true proposition is trivially verifiable and so $\mathcal{O}X$ is closed under finite conjunctions. Note that this argument does not extend to infinite conjunctions, since checking infinitely many verifiable propositions in turn is not achievable via finite means. However, verifiable properties are closed under disjunctions of arbitrary cardinality: to verify that a disjunction holds at some point, we need only verify that one disjunct holds, which does not require any infinite processes. Note in particular that the constantly false position is verifiable, since we need only verify propositions when they hold.

In summary, the logic of verifiability admits finite conjunctions and arbitrary disjunctions. Identifying each proposition U with the set of points which satisfy it recovers the usual notion of a *topological space* with conjunctions interpreted as set intersection and disjunctions as set union. Thus, point-set topology provides one formalisation of the theory of verifiability. Indeed, this perspective even extends to *continuous functions*.

Let $f: X \rightarrow Y$ be some function between spaces, and let U be some verifiable property on Y . Supposing f is physically realisable, we can verify if $f(x) \in U$ or equivalently if $x \in f^{-1}(U)$. Thus, we recover that preimages of open sets must be open. Note that from this point of view, all physically realisable functions must necessarily be continuous.

It is instructive to consider why such reasoning *cannot* be applied to prove that every physically realisable function $f: X \rightarrow Y$ is open — that is, that the image of every open V in X is open in Y . How would we verify that that

$y \in f(V)$? We would need to semi-decide that $y = f(x)$ for some $x \in V$. But equality is not verifiable in general and so we are not always able to do this. If we could semi-decide equality, then each singleton $\{y\}$ would necessarily be open, which would make Y a discrete space, and indeed, all functions into discrete spaces *are* open.

Concrete topological spaces can be fruitfully interpreted from the perspective of verifiability.

Example 2.1. Consider the case of the real numbers \mathbb{R} . We might imagine having devices that can measure a quantity x , each to some fixed precision giving us a rational number. Moreover, let us assume that for any precision there exists a device that is at least that precise.

We can use these to semi-decide if $x > a$ for some given $a \in \mathbb{R}$, since if $x > a$ then there is always some ε such that $x - \varepsilon > a$ and so using a device with such a precision will verify that x is indeed greater than a . By a similar argument, we can semi-decide if $x < b$ and hence if $x \in (a, b)$. On the other hand, it is not always possible to check if $x \geq a$, as in the scenario where $x = a$ we need infinite precision to be sure that x is not actually very slightly smaller than a and any given device has only finite precision. This agrees with the familiar topology on \mathbb{R} . \triangle

2.2. Pointfree topology. The logic of verifiability can also be studied purely through the verifiable propositions themselves, without any regard to the points which may or may not satisfy these propositions. This is the perspective of *pointfree topology*, the study of topology through the lattice of open sets. In this section we introduce some of the basic notation and concepts which will be used throughout the rest of the paper. See [16, 28] for further details.

We have seen that verifiable properties are closed under finite conjunctions and arbitrary disjunctions. Identifying logically equivalent propositions and ordering the equivalence classes by logical entailment we obtain a lattice. This is the Lindenbaum–Tarski algebra for the logic. For the conjunctions and disjunctions to behave as expected, we require that this lattice is distributive. This motivates the following definition.

Definition 2.2. A *frame* is a complete lattice which satisfies the distributivity law

$$u \wedge \bigvee_{i \in I} v_i = \bigvee_{i \in I} u \wedge v_i.$$

A *locale* X is formally the same thing as a frame, but is thought of as being the abstract space that has the corresponding frame $\mathcal{O}X$ as its lattice of opens. \triangle

Of course, the open sets of any topological space form a frame under intersection and union, though there are also non-spatial examples.

We can also describe continuous functions from this perspective. Since the preimage of an open is an open, every continuous function induces a function between the lattices of open sets but in the opposite direction. From set-theoretic properties of preimages, we see that this map preserves finite meets and arbitrary joins. Thus, we define a locale morphism as follows.

Definition 2.3. A *frame homomorphism* $f: L \rightarrow M$ is a function $f: L \rightarrow M$ between frames that preserves finite meets and arbitrary joins. A *locale morphism* $f: X \rightarrow Y$ is a frame homomorphism $f^*: \mathcal{O}Y \rightarrow \mathcal{O}X$ between the corresponding frames of opens. \triangle

In good situations, a topological space can be recovered from its frame of open sets. (See [16] for more details.) In particular, we can talk about *points* of a locale. As in \mathbf{Top} , we can identify points of a locale with maps from the terminal object 1 . These correspond to frame homomorphisms from $\mathcal{O}X$ to $\mathcal{O}1 \cong \{0, 1\}$, which can be understood logically as assigning truth values to each verifiable proposition P corresponding to whether P holds at x .

Since frames are algebraic structures, they can be *presented* by generators and relations. In particular, free frames exist. The free frame on G can be described explicitly as the frame of downsets on the meet-semilattice of finite subsets of G ordered by reverse inclusion. The finite subsets of G are interpreted as a formal meets of generators, while the downsets are viewed as formal joins of these. Note that every formal expression involving frame operators on the generators can be brought into the form of a join of finite meets by distributivity.

Definition 2.4. A *presentation* for a frame consists of a set of generators G and a set of relations R consisting of formal equalities (or inequalities) between formal combinations of generators — for example, $g_1 \wedge g_2 = \bigvee_{i=3}^{\infty} g_i \wedge g_2$. Explicitly, R can be defined to be a subset of $F(G) \times F(G)$ where $F(G)$ is the free frame on the set G .

The frame $\langle G \mid R \rangle$ defined by such a presentation contains elements corresponding to the generators $g \in G$ and which satisfy the relations given in R . Moreover, it is the initial frame satisfying this property: for any frame M and function $f: G \rightarrow M$ for which the images of the generators under f satisfy the

relations from R , there is a unique frame homomorphism $\bar{f}: \langle G \mid R \rangle \rightarrow M$ making the triangle commute.

$$\begin{array}{ccc}
 & \langle G \mid R \rangle & \\
 & \uparrow & \searrow \bar{f} \\
 G & \xrightarrow{f} & M
 \end{array}$$

Explicitly, the frame $\langle G \mid R \rangle$ can be constructed as the quotient of the free frame $F(G)$ by the congruence relation generated by R .

Note that taking M in the diagram above to be the initial frame $\{0, 1\}$, we see that the points of $\langle G \mid R \rangle$ are given by specifying the subset of the generators that are viewed as ‘true’ such that relations become logical formulae. For example, our relation above becomes $g_1 \wedge g_2 \iff \exists i \geq 3. g_i \wedge g_2$, so g_1 and g_2 should be true if and only if g_i and g_2 are true for some $g_i \geq 3$. \triangle

One particularly important locale is Σ corresponding to *Sierpiński space*. As a space this is $\{\perp, \top\}$ with the topology generated by the single open $\{\top\}$. It can be formally defined as the locale corresponding to the free frame on a single generator (corresponding to $\{\top\}$). Locale maps $u: X \rightarrow \Sigma$ correspond to opens of X by the universal property of the free frame. Spatially, this can be understood as taking the preimage of $\{\top\}$. From the verifiability perspective this allows us to interpret the elements of Σ as corresponding to whether the verification process halts (\top), or runs forever (\perp).

Note that if G is a set then Σ^G is the locale corresponding to the free frame on G and so a presentation of X can be understood as an equaliser of locales $X \hookrightarrow \Sigma^G \rightrightarrows \Sigma^R$. (Recall that locales are dual to frames and so a coequaliser of frames becomes an equaliser of locales.)

The equaliser morphism from $X \hookrightarrow \Sigma^G$ can be viewed as analogous to an embedding of topological spaces. Regular subobjects of locales are called *sublocales* and correspond to frame quotients.

Another concept we recall here is that of compactness. The usual topological definition in terms of finite subcovers works equally well in this setting.

Definition 2.5. A locale X is compact if whenever $\bigvee_{i \in I} u_i = 1$ there exists a finite set $F \subseteq I$ such that $\bigvee_{i \in F} u_i = 1$. \triangle

As mentioned before, compactness has a very nice interpretation in terms of universal quantification of verifiable properties. We will see some instances of this in later sections.

2.3. The Scott topology and exponentials. The set of opens $\mathcal{O}X$ of a locale X is itself often endowed with the *Scott topology*.

Definition 2.6. A subset $V \subseteq \mathcal{O}X$ is *Scott-open* if it is upward closed and if whenever a join $\bigvee D \in V$ for a directed set D , then some $d \in D$ lies in V . Recall that a set D is *directed* if every finite subset of D has an upper bound in D . \triangle

If X is locally compact, then X is exponentiable. In particular, the exponential locale Σ^X exists and coincides with the Scott topology on $\mathcal{O}X \cong \text{Hom}(X, \Sigma)$. Furthermore, Σ^X is itself locally compact. However, if X is not locally compact the evaluation map $(f, x) \mapsto f(x)$ is not continuous with respect to the Scott topology and the exponential Σ^X does not exist.

The Scott topology can actually be defined more generally on *dcpos*, which are often used to model data types in programming languages.

Definition 2.7. A poset which admits joins of all directed subsets is called a *dcpo*. We will write $\bigvee^\uparrow S$ for the join of a directed set S . A morphism of dcpos or a *Scott-continuous function* is a monotone map which preserves directed suprema. \triangle

The order on a dcpo can be thought of as a definability order: the least element (if it exists) corresponds to a purely divergent computation, while maximal elements are completely defined. For more on this topic see [1, 2, 10].

3. Machine space

In this section we make certain aspects of the verifiable interpretation of topology more precise, specifically the relationship between opens and the processes that semi-decide some fact. Our approach makes use of presentations of frames and inspires a number of purely topological results — many with interesting interpretations from the point of view of verifiability.

Given a space X with $\mathcal{O}X = \langle G \mid R \rangle$ we interpret each generator $g \in G$ as a basic black-box machine which takes as input the points of X and either halts after some finite amount of time, or runs forever. More complicated machines can be constructed from these basic machines. For instance, we may construct a composite machine by taking a finite collection of machines and insist that each

machine halts on the given point before the composite machine is considered to have halted. We can also construct a further machine by running any number of these composite machines in parallel and halting on a point if at least one branch halts. Topologically this corresponds to taking formal joins of finite formal meets of generators.

Definition 3.1. Given a space X with $\mathcal{O}X = \langle G \mid R \rangle$, an (*idealised*) machine m over X is a formal join of finite formal meets of generators written $m = \bigvee_{i \in I} \bigwedge_{j \in J_i} g_j$ with each J_i finite. \triangle

A point $x \in X$ is accepted by a machine m if there exists some J_k such that x is accepted by each ‘basic machine’ g_j for all $j \in J_k$. In this way each machine m will define an open on which it halts. Many of these machines will correspond to the same opens, as described by the relation R , so that each open in X corresponds to an equivalence class of machines (see Definition 3.5 and Proposition 3.6).

Note that such relations must be taken as given axiomatically and cannot be verified to hold. (After all, we cannot show in finite time that two machines halt on precisely the same points.) It is important to maintain a sharp distinction between the properties we assume as axioms, which define the spaces under consideration and cannot be proven empirically, and the properties we observe about the spaces we describe, which concern unspecified points and must be verifiable.

In what follows we will formally define the space of machines and its relation to the exponential Σ^X , when it exists, and to the Scott topology on $\mathcal{O}X$ more generally.

3.1. Presentations and weak exponentials. For a general space X the space of opens Σ^X need not exist. However, given a presentation $\mathcal{O}X = \langle G \mid R \rangle$ there is however always a natural space of machines, whose points are precisely the machines described above and whose opens have a natural interpretation in terms of observable properties of machines. We claim that this space is Σ^{Σ^G} and will justify this assertion in the paragraphs below.

Recall that $\langle G \mid R \rangle$ is a quotient of the free frame on G , which in turn is the frame of opens of Σ^G . The opens of Σ^G are precisely the machines mentioned in Definition 3.1.

The frame quotient $\Sigma^G \twoheadrightarrow \langle G \mid R \rangle$ exhibits X as a sublocale of Σ^G . The points of Σ^G correspond to subsets of G and we can then interpret this inclusion as sending points of X to the set of generating opens in which they lie. The

points in the image of this inclusion correspond to the frame homomorphisms $\mathcal{O}(\Sigma^G) \rightarrow \mathcal{O}1$ which respect the relations. General points of Σ^G may then be interpreted as ‘generalised points’ of X and can be given as input to machines on X . This perspective will be useful when we study compactness in Section 4.

So machines on X correspond to opens of Σ^G . But Σ^G is always locally compact and so we can consider Σ^{Σ^G} — the space of opens of Σ^G . This will be our space of machines for X . The topology on Σ^{Σ^G} is that of the double powerlocale on G (see [29, 31]) and is generated by opens $\boxtimes U$ for each $U \subseteq G$ where a machine $m = \bigvee_{i \in I} \bigwedge_{j \in J_i} g_j$ lies in $\boxtimes U$ if and only if there exists some J_k such that $g_j \in U$ for each $j \in J_k$. Intuitively this means we can semi-decide if there exists some parallel branch of a machine m such that each of the generators that run in that branch are contained in U .

The fact that we have nontrivial verifiable properties indicates that these machines are not complete black boxes and that there is some information that leaks while they run. In order to understand this better, let us consider the following compelling but flawed model of the situation, before showing how to remedy it. Note that with all such “real-world” models the set of generators is assumed to be countable.

Model 3.2. Let X be a locale with $\mathcal{O}X = \langle G \mid R \rangle$. We take each generator $g \in G$ to be a black box which accepts points of X as input and either halts after some finite time or runs forever. We imagine these generators arranged in a line on some suitably long table. Also on the table are a number of mobile robots which themselves accept points of X as input and will roll from generator to generator testing the point in the following manner.

The robot progresses through different stages. On the i^{th} stage it selects a group of finitely many generators, writes the number i onto them and (if they are not already running) runs them with the point as input. The generators are left to run as the next stage is started. If at any point every generator in a group has halted, the a light on the robot’s head flashes green indicating that it accepts the point.

Such a robot instantiates a machine $\bigvee_{i \in I} \bigwedge_{j \in J_i} g_j$ where each J_i is the finite set of generators in group i . △

If we take these robots to be our machines, then we can indeed verify $\boxtimes U$ by simply checking that for every machine in some group is contained in U . However, there are additional questions that we can semi-decide which do *not* occur as opens in Σ^{Σ^G} , such as:

- (1) “is the generator g_1 visited before the generator g_2 ?”
- (2) “is the generator g ever visited?”
- (3) “is the group $\{g_1, g_2\}$ selected at least twice?”

The problem is that the robots, as described, do not actually correspond to points of Σ^{Σ^G} , rather they are points in a ‘freer’ structure in which \wedge and \vee are not assumed to be commutative, idempotent or absorbent. We can eliminate questions of the first kind by making the robot choose the order of the generators visited (both within and between groups) nondeterministically. For questions of the second kind, note that the machine g should be identified with the machine $g \vee (g \wedge h)$. So if the robot visits the generators $S \subseteq G$ in a particular branch, then we should require it to visit a branch with generators T for all $T \supseteq S$. This will ensure that questions of the second form will always be true and hence trivially verifiable. Finally, we ensure that no group of machines is selected more than once to render third question trivial in a similar way.

Note that original model already captured associativity and distributivity of the operations. Distributivity follows the fact that the formal operations are already reduced the canonical form of a join of meets. These axioms together are then enough to give a frame, which ensures that no extraneous verifiable properties of this sort exist.

Of course, if actually implemented in the real world there would almost certainly be additional questions one could ask pertaining to the specific engineering of the robots, manufacturing defects, and so on. However, we consider an idealised setting where the only information that can be gleaned is related to the robot’s interaction with the generators.

We believe this idealised model is still compelling evidence that such a space of machines is realisable. Thus, the space of machines is a sense more concrete than the space of opens, even when the latter exists, and it is interesting to see how we might obtain the space of opens from machine space. We will do this in Section 3.2.

Notice that Σ^{Σ^G} by itself contains no information about the behaviour of the machines on inputs from the locale X with presentation $\langle G \mid R \rangle$, as this data is contained in the relation R . We can represent this data if we equip Σ^{Σ^G} with an evaluation map $\tilde{ev}: \Sigma^{\Sigma^G} \times X \rightarrow \Sigma$ which records whether a point lies in a machine. This makes Σ^{Σ^G} a *weak exponential* as we show below. (Compare the construction of weak exponentials in Top given in [20].)

Proposition 3.3. *The space Σ^{Σ^G} together with the evaluation map $\tilde{ev}: \Sigma^{\Sigma^G} \times X \rightarrow \Sigma$ given by the composite $\Sigma^{\Sigma^G} \times X \hookrightarrow \Sigma^{\Sigma^G} \times \Sigma^G \xrightarrow{ev} \Sigma$ is a weak exponential with base Σ and exponent X .*

Proof: Since X embeds into Σ^G via $i_X: X \hookrightarrow \Sigma^G$, the product $\Sigma^{\Sigma^G} \times X$ embeds into $\Sigma^{\Sigma^G} \times \Sigma^G$. But Σ^G is locally compact and so we have an evaluation map $ev: \Sigma^{\Sigma^G} \times \Sigma^G \rightarrow \Sigma$. To see that $\tilde{ev} = ev \circ (\Sigma^{\Sigma^G} \times i_X)$ makes Σ^{Σ^G} a weak exponential we consider a morphism $u: A \times X \rightarrow \Sigma$ and must construct a morphism v such that $u = \tilde{ev} \circ (v \times X)$. Consider the following diagram.

$$\begin{array}{ccccc}
 A \times X & & & & \\
 \downarrow \bar{u}' \times X & \searrow A \times i_X & & \searrow u & \\
 & A \times \Sigma^G & & & \\
 & \downarrow \bar{u}' \times \Sigma^G & \dashrightarrow u' & & \\
 \Sigma^{\Sigma^G} \times X & \xrightarrow{\Sigma^{\Sigma^G} \times i_X} & \Sigma^{\Sigma^G} \times \Sigma^G & \xrightarrow{ev} & \Sigma
 \end{array}$$

Observe that u defines an open in $A \times X$. Since $A \times X$ is a subspace of $A \times \Sigma^G$, there exists a (not necessarily unique) open u' in $A \times \Sigma^G$ which restricts to give u . By the universal property of the exponential Σ^{Σ^G} , we obtain a morphism \bar{u}' making the bottom right triangle commute. It is not hard to see that the left-hand trapezium commutes and then a simple diagram chase confirms that $u = \tilde{ev} \circ (\bar{u}' \times X)$. \blacksquare

We arrive at the following definition.

Definition 3.4. Let X be a locale with $\mathcal{O}X = \langle G \mid R \rangle$. We define the associated (*idealised*) machine space to be Σ^{Σ^G} together with the evaluation map $\tilde{ev}: \Sigma^{\Sigma^G} \times X \rightarrow \Sigma$ defined in Proposition 3.3. \triangle

3.2. Relationship between machine space and the space of opens. It is enlightening to consider how Σ^{Σ^G} relates to Σ^X in the case that the latter does exist. Intuitively, we expect there to be a quotient map $q: \Sigma^{\Sigma^G} \rightarrow \Sigma^X$ which sends a machine to its corresponding open.

Definition 3.5. Let X be a locally compact locale with $\mathcal{O}X = \langle G \mid R \rangle$. We may define the map q of machine space via the universal property of the exponential Σ^X applied to $\tilde{e}v$.

$$\begin{array}{ccc}
 \Sigma^{\Sigma^G} & & \Sigma^{\Sigma^G} \times X \\
 \downarrow q & & \downarrow q \times X \quad \searrow \tilde{e}v \\
 \Sigma^X & & \Sigma^X \times X \xrightarrow{\text{ev}} \Sigma^{\Sigma^G} \times \Sigma^G
 \end{array}$$

Equivalently, q is given by Σ^{i_X} where i_X is the inclusion of X into Σ^G . \triangle

This map behaves in accordance with our intuition. The diagram can be interpreted as meaning $\tilde{e}v(m, x) = \text{ev}(q(m), x)$ which is to say that m halts on x if and only if x lies in $q(m)$. Moreover, it can be shown to be a quotient map. In fact, it necessarily has a section. (Compare the Σ -split inclusions of Taylor [24].)

Proposition 3.6. *Let X be a locally compact locale with $\mathcal{O}X = \langle G \mid R \rangle$. The map $q: \Sigma^{\Sigma^G} \rightarrow \Sigma^X$ as defined in Definition 3.5 has a section s satisfying $\tilde{e}v \circ (s \times X) = \text{ev}$.*

Proof: We define s from the ‘weak’ universal property of the weak exponential applied to ev .

$$\begin{array}{ccc}
 \Sigma^{\Sigma^G} & & \Sigma^{\Sigma^G} \times X \\
 \downarrow q \quad \uparrow s & & \downarrow q \times X \quad \uparrow s \times X \quad \searrow \tilde{e}v \\
 \Sigma^X & & \Sigma^X \times X \xrightarrow{\text{ev}} \Sigma^{\Sigma^G} \times \Sigma^G
 \end{array}$$

It is immediate that $\tilde{e}v \circ (s \times X) = \text{ev}$ and so all that remains is to prove that s is a section of q .

By the diagram we have $\text{ev} \circ (q \times X) \circ (s \times X) = \tilde{e}v \circ (s \times X) = \text{ev} = \text{ev} \circ (\text{id}_{\Sigma^X} \times X)$. Hence by the uniqueness condition of the universal property $qs = \text{id}_{\Sigma^X}$, as required. \blacksquare

The section s continuously ‘picks out representatives of the equivalence classes defined by the quotient’. It can be viewed as sending an open u to a machine which carries out the procedure of verification for u .

The relationship between machine space and the exponential Σ^X may be fruitfully interpreted from the perspective of verifiability. If one is not careful, one can be led to believe that the verifiable interpretation of topology implies that the space of opens Σ^X should exist for all spaces X . This deficient reasoning would go as follows: the collection of opens should have some topological structure corresponding to the verifiable properties on it. Let us call this space $\mathcal{O}X$. Given an open $u \in \mathcal{O}X$ and a point $x \in X$, since u is a verifiable property, we can verify whether x lies in u . Carrying out this procedure should take a value in Σ corresponding to whether the process ran forever or halted. Thus, this procedure represents a physically realisable evaluation map from $\mathcal{O}X \times X$ to Σ which implies that this evaluation map is continuous. It is then not hard to check this would satisfy the properties of an exponential object, proving that $\mathcal{O}X = \Sigma^X$.

From the perspective of machine space the error lies in assuming that given an open and a point, it is possible to semi-decide membership of the point in the open. An open merely *describes* the verifiable property under consideration — it is a *machine* that actually physically carries out this verification. From this point of view, the fact that a section of the quotient map exists whenever X is locally compact perfectly explains why a continuous evaluation map exists for $\mathcal{O}X$. Given an open u and a point x , one can apply the section to u to acquire a machine $s(u)$ which can then verify membership of x in u .

On the other hand, when X is not locally compact the argument breaks down and we will see that no way to continuously assign opens to machines in this case. Formulating the quotient of machine space for non-locally-compact locales presents some difficulties in the pointfree setting, so in the following section we will examine this from the point of view of point-set topology. In Section 3.2.2 we formalise what we can in the pointfree setting. The first section is sufficient for understanding the intuition, while the second is more technical.

3.2.1. The space of opens as a topological space. Recall that if a locale X is given by $\mathcal{O}X = \langle G \mid R \rangle$ then the frame $\mathcal{O}X$ is a quotient of $\mathcal{O}(\Sigma^G)$ by the congruence generated from the relations R . Let us call this frame quotient map \bar{q} . Note that $\mathcal{O}(\Sigma^G)$ equipped with the Scott topology is precisely the machine space Σ^{Σ^G} . We can now use \bar{q} to induce a topology on $\mathcal{O}X$.

Proposition 3.7. *The quotient topology induced on $\mathcal{O}X$ by the function $\bar{q}: \Sigma^{\Sigma^G} \rightarrow \mathcal{O}X$ is the Scott topology on $\mathcal{O}X$.*

Proof: We need to show that $U \subseteq \mathcal{O}X$ is Scott-open if and only if $\bar{q}^{-1}(U)$ is Scott-open. First suppose U is Scott-open. We start by showing that $\bar{q}^{-1}(U)$ is an upset. Take $a \in \bar{q}^{-1}(U)$ and suppose $a \leq b$. Then $\bar{q}(a) \leq \bar{q}(b)$ and $\bar{q}(a) \in U$, giving $\bar{q}(b) \in U$. Hence, $b \in \bar{q}^{-1}(U)$ and $\bar{q}^{-1}(U)$ is an upset. Now suppose $\bigvee^\uparrow D \in \bar{q}^{-1}(U)$. This means $q(D) \subseteq U$ and $\bigvee^\uparrow \bar{q}(D) \in U$. But $\bar{q}(D)$ is also directed and so, by Scott-openness of U , there is a $d \in D$ such that $\bar{q}(d) \in U$. Hence, $d \in \bar{q}^{-1}(U)$ and $\bar{q}^{-1}(U)$ is Scott-open.

Now take $U \subseteq \mathcal{O}X$ and suppose $\bar{q}^{-1}(U)$ is Scott-open. Say $a' \in U$ and $a' \leq b'$. Now take a to be any element of $\bar{q}^{-1}(a')$ and b to be an element of $\bar{q}^{-1}(b')$ lying above a . This is always possible by joining with a as necessary. Then $a \in \bar{q}^{-1}(U)$, as $a' = \bar{q}(a) \in U$. Hence, $b \in \bar{q}^{-1}(U)$, since $\bar{q}^{-1}(U)$ is upward closed. Thus, $b' = \bar{q}(b) \in U$ and so U is an upset. Finally, suppose $\bigvee^\uparrow D' \in U$. Set $D = \bar{q}^{-1}(D')$. Note that D is directed, that $\bar{q}(D) = D'$ and that $\bar{q}(\bigvee^\uparrow D) = \bigvee^\uparrow \bar{q}(D) = \bigvee^\uparrow D' \in U$. Thus, $\bigvee^\uparrow D \in \bar{q}^{-1}(U)$ and so there is a $d \in D$ such that $d \in \bar{q}^{-1}(U)$. Then $\bar{q}(d)$ is the desired element $d' \in D'$ with $d' \in U$ and so U is Scott-open. \blacksquare

Remark 3.8. Note that a similar proof shows this for any Scott-continuous order quotient between dcpos.

In the case that X is locally compact, this map \bar{q} agrees with the map q we previously defined by the universal property.

Lemma 3.9. *The map $\bar{q}: \Sigma^{\Sigma^G} \rightarrow \mathcal{O}X$ is equal to $q: \Sigma^{\Sigma^G} \rightarrow \Sigma^X$ as defined in Definition 3.5 whenever X is locally compact.*

Proof: We first note that the codomains agree since $\mathcal{O}X = \Sigma^X$ for locally compact X . Now recall that $q = \Sigma^{i_X}$. The natural transformation $(-)^{i_X}$ is the mate of $(-) \times i_X$ with respect to the adjunction $(-) \times X \dashv (-)^X$ and so we have the following commutative diagram.

$$\begin{array}{ccc}
 \text{Hom}(1, \Sigma^X) = \text{Hom}(1 \times X, \Sigma) & & \\
 \uparrow & & \uparrow \\
 \text{Hom}(1, \Sigma^{i_X}) & & \text{Hom}(1 \times i_X, \Sigma) \\
 \uparrow & & \uparrow \\
 \text{Hom}(1, \Sigma^{\Sigma^G}) & = & \text{Hom}(1 \times \Sigma^G, \Sigma)
 \end{array}$$

The map $\text{Hom}(i_X, \Sigma)$ corresponds to the action of the frame map $i_X^* = \bar{q}$, while $\text{Hom}(1, \Sigma^{i_X})$ gives the action of Σ^{i_X} on points. Thus, we have that Σ^{i_X} and \bar{q} agree on points and hence they coincide. ■

The failure of $\mathcal{O}X$ with the Scott topology to be an exponential Σ^X when X is not locally compact is due to this quotient map being badly behaved in general. Indeed, if there were a section to q as in Proposition 3.6 we could show X to be locally compact after all. (The proofs of the following two results require us to assume X itself is spatial, but are useful for intuition. We give pointfree results in Section 3.2.2.)

Proposition 3.10. *If $\bar{q}: \Sigma^{\Sigma^G} \rightarrow \mathcal{O}X$ has a continuous section $s: \mathcal{O}X \rightarrow \Sigma^{\Sigma^G}$, then X is locally compact. Moreover, the evaluation map for $\mathcal{O}X$ is $\tilde{\text{ev}} \circ (s \times X)$.*

Proof: Let $\text{ev} = \tilde{\text{ev}} \circ (s \times X)$. On points we have $\text{ev}(U, x) = \tilde{\text{ev}}(s(U), x)$, which is equal to \top precisely when $i_X(x) \in s(U)$ in Σ^G and hence when $x \in U$ in X . Thus, ev is the evaluation map for exponential of the underlying sets. Moreover, ev is continuous, since it is a composition of continuous maps. Every continuous map from $Y \times X \rightarrow \Sigma$ factors through ev to give a set-theoretic function from Y to Σ^X . But this map is continuous since $\mathcal{O}X$ has the final topology with respect to \bar{q} . Thus, $\mathcal{O}X$ is indeed the exponential in Top and hence X is locally compact. ■

In fact, we can say more. It is known that lack of exponential objects is related to the fact that quotients are not stable under pullback (see [4]). We have the following following result.

Proposition 3.11. *If $\bar{q} \times X: \Sigma^{\Sigma^G} \times X \rightarrow \mathcal{O}X \times X$ is a quotient map, then $\mathcal{O}X$ (with the Scott topology) is the exponential Σ^X .*

Proof: Consider the following diagram.

$$\begin{array}{ccccc}
 & & h'' \times X & & \\
 & & \curvearrowright & & \\
 A \times X & \xrightarrow{h' \times X} & \Sigma^{\Sigma^G} \times X & \xrightarrow{\bar{q} \times X} & \mathcal{O}X \times X \\
 & \searrow h & \downarrow \tilde{\text{ev}} & \swarrow \text{ev} & \\
 & & \Sigma & &
 \end{array}$$

Here ev is defined as the evaluation map from the exponential of the underlying sets. It is continuous since the composite $\tilde{\text{ev}} = \text{ev} \circ (q \times X)$ is continuous and $\bar{q} \times X$ is a quotient map.

If $h: A \times X \rightarrow \Sigma$ then by the weak universal property of $\tilde{\text{ev}}$ we have a map $h': A \rightarrow \Sigma^{\Sigma^G}$ making the diagram commute. Then $h'' = \bar{q}h'$ satisfies the condition needed for the weak universal property of $\text{ev}: \mathcal{O}X \times X \rightarrow \Sigma$ to hold. Moreover, it is unique since we have uniqueness on the underlying sets by the universal property of the exponential in Set . Thus, $\mathcal{O}X$ satisfies the universal property of the exponential Σ^X , as required. ■

3.2.2. Pointfree quotients of machine space. For non-locally-compact locales, it is not obvious how to define the Scott topology on $\mathcal{O}X$ and the corresponding quotient machine space in a pointfree way. Nonetheless we can still prove some results that relate retracts of machine space to local compactness and existence of the exponential Σ^X . This section is more technical than previous ones and can be skipped by readers who are mainly interested in intuition.

We have seen that a section to the quotient $q: \Sigma^{\Sigma^G} \rightarrow \Sigma^X$ can be used to recover the evaluation map on Σ^X . In fact, even we can say even more. Suplattice and preframe homomorphisms between frames are known to correspond to angelically and demonically nondeterministic maps respectively [23, 32, 31] and it would seem that a nondeterministic choice of representative for each open would be sufficient to construct the evaluation map.

Quotient maps with nondeterministic sections correspond to certain *triquotient* maps. For example, see [27] for how open and proper quotients can be understood in these terms.

Definition 3.12. A locale map $f: X \rightarrow Y$ is a *triquotient* if there exists a dcpo morphism $f_{\#}: \mathcal{O}X \rightarrow \mathcal{O}Y$, called a *triquotient assignment*, satisfying $f_{\#}(a \wedge f^*(b)) = f_{\#}(a) \wedge b$ and $f_{\#}(a \vee f^*(b)) = f_{\#}(a) \vee b$. (The map $f_{\#}$ is automatically a dcpo retraction of f^* .) \triangle

Triquotients generalise both open and proper quotients and every retraction is a triquotient where we take the section as the triquotient assignment. Plewe has shown in [18] that triquotients are stable under pullback and hence $q \times X$ is a triquotient whenever q is. We now state a very general result about how to define the evaluation map (as long as X is locally compact).

Proposition 3.13. *Let X be a locally compact locale and suppose $q: \Sigma^{\Sigma^G} \rightarrow \Sigma^X$ is a triquotient. Then $\text{ev}^* = t \circ \tilde{\text{ev}}^*$ where t is any triquotiency assignment of $q \times X$.*

Proof: From the definition of q we have the following diagram.

$$\begin{array}{ccc}
 \mathcal{O}(\Sigma^{\Sigma^G}) \oplus \mathcal{O}X & & \\
 \uparrow & \swarrow \tilde{\text{ev}}^* & \\
 q^* \oplus \mathcal{O}X & & \mathcal{O}(\Sigma^{\Sigma^G}) \oplus \mathcal{O}(\Sigma^G) \\
 \downarrow t & \swarrow \text{ev}^* & \\
 \mathcal{O}(\Sigma^X) \oplus \mathcal{O}X & &
 \end{array}$$

Then $t \circ \tilde{\text{ev}}^* = t \circ (q^* \oplus \mathcal{O}X) \circ \text{ev}^* = \text{ev}^*$, as required. \blacksquare

We now consider the situation when X is not necessarily locally compact. Even without being able to easily define the space of opens in a pointfree way, we can still provide a kind of converse to Proposition 3.6.

For this we will need the definition of a continuous dcpo.

Definition 3.14. The *way-below relation* \ll on a dcpo L is defined so that $a \ll b$ if and only if $b \leq \bigvee^\uparrow D$ implies there is some $d \in D$ with $a \leq d$. We say L is *continuous* if for every $b \in L$ we have $b = \bigvee^\uparrow \{a \in L \mid a \ll b\}$. \triangle

A locale X is locally compact if and only if its frame of opens $\mathcal{O}X$ is continuous.

The result involves a general kind of map that generalises both retractions and open quotients and is incomparable with triquotients.

Definition 3.15. We say a locale map $e: X \rightarrow Y$ is a *semi-open quotient* if e^* has a suplattice retraction. \triangle

Proposition 3.16. *If there is a semi-open quotient map $e: \Sigma^{\Sigma^G} \rightarrow \mathcal{O}X$ (where $\mathcal{O}X$ has the Scott topology) then X is locally compact.*

Proof: Every locale of the form Σ^Y is injective (see [12, Chapter VII]). Thus, in particular $\mathcal{O}(\Sigma^{\Sigma^G})$ is a projective frame and hence a retract of a free frame. But since e is a semi-open quotient, $\mathcal{O}(\mathcal{O}X)$ is a suplattice retraction of $\mathcal{O}(\Sigma^{\Sigma^G})$. Thus, $\mathcal{O}(\mathcal{O}X)$ is a suplattice retraction of a free frame.

The free frame is constructed as a frame of downsets and is thus constructively completely distributive [9]. But a suplattice retract of a constructively

completely distributive lattice is constructively completely distributive (see [19]) and hence so is $\mathcal{O}(\mathcal{O}X)$.

But it is shown in [26] that the points of a completely distributive frame are a continuous dcpo. Thus, $\mathcal{O}X$ is continuous and so X is locally compact. ■

If we could define the map $\bar{q}: \Sigma^{\Sigma^G} \rightarrow \mathcal{O}X$ and it were semi-open then this result would imply that X is locally compact as in Proposition 3.10. It would be interesting to know if Proposition 3.16 also holds for triquotient maps.

4. Compactness and universal quantification

We believe that machine space also has the potential to be useful in clarifying other aspects of topology. In this section we use it to study compactness. Escardó has explained compactness in terms of an algorithm for universal quantification in the setting of semantics of programming languages [5, 7, 8]. Our approach yields a purely topological version of this algorithm.

4.1. An algorithm for universal quantification. Recall that a locally compact locale X is compact precisely when $\{1\}$ is open in Σ^X . More generally, if a X is embedded into a locally compact space Y , then X is compact if and only if there is an open of Σ^Y consisting of the opens that cover X (see [11]). In terms of machine space this says that if $\mathcal{O}X = \langle G \mid R \rangle$ then X is compact if and only if there is an open in Σ^{Σ^G} consisting of the machines which halt on all of X . From the perspective of verifiability, this means we can semi-decide if a given machine always halts on X . Indeed, we will provide an algorithm to carry out this very procedure.

It is important to understand exactly what it is such an algorithm needs to do. Of course, to give such a procedure we need to know a precise description of the space X by generators and relations. From this we can mathematically derive whether a particular formal combination of generators is equal to 1. On the other hand, recall that ‘the exact composition’ of the machines in machine space is opaque to us. Indeed, the opens of machine space only allow us to test very particular properties of the machines. In the first case we are dealing with discrete syntax, while in the second we are working with abstract machines which we do not have knowledge about a priori.

We are now ready to consider the algorithm for universal quantification over a compact locale X with presentation $\langle G \mid R \rangle$. (To actually run such an algorithm we should restrict G to be countable, but we can still imagine unbounded parallelism in theory.)

Algorithm 1 Semi-decision procedure for universal quantification over X

$\forall_X: \Sigma^{\Sigma^G} \rightarrow \Sigma$
function $\forall_X(m)$
 for each $S \in \mathcal{P}_{\text{fin}}(\mathcal{P}_{\text{fin}}(G))$ **do** ▷ performed in parallel
 if $\bigvee_{F \in S} \bigwedge_{g \in F} g \sim 1$ with respect to R **then**
 for each $F \in S$ **do**
 test($m \in \boxtimes F$) ▷ semi-decides if a branch is contained in F
 HALT

Here $\mathcal{P}_{\text{fin}}(T)$ denotes the set of finite subsets of T and **test**($m \in \boxtimes F$) halts precisely when $m \in \boxtimes F$. Recall that a machine $m = \bigvee_{i \in I} \bigwedge_{j \in J_i} g_j$ lies in $\boxtimes F$ if and only if there exists some J_k such that $g_j \in F$ for each $j \in J_k$ — that is, if every generator in some branch of m lies in F .

Note that the algorithm itself involves testing whether $\bigvee_{F \in S} \bigwedge_{g \in F} g$ is a cover of X . This is on the level of syntax and how to test this is determined mathematically before running the algorithm. The role of the algorithm is to translate from the level of syntax to the level of opaque machines.

Theorem 4.1. *Let X be a compact locale with $\mathcal{O}X = \langle G \mid R \rangle$. Algorithm 1 semi-decides if the given machine m halts on all of X .*

Proof: Let $m = \bigvee_{i \in I} \bigwedge_{j \in J_i} g_j$. First suppose that this covers X so that $\bigvee_{i \in I} \bigwedge_{j \in J_i} g_j = 1$ in X . Since X is compact there is a finite $I' \subseteq I$ such that $\bigvee_{i \in I'} \bigwedge_{j \in J_i} g_j = 1$. Now set $S = \{J_i \mid i \in I'\} \in \mathcal{P}_{\text{fin}}(\mathcal{P}_{\text{fin}}(G))$. This set S will be considered in some parallel branch of the algorithm and we have $\bigvee_{F \in S} \bigwedge_{g \in F} g \sim 1$ in R by construction. Note that **test**($m \in \boxtimes J_i$) halts for each $J_i \in S$, since $\bigwedge J_i \leq m$ and hence $m \in \boxtimes J_i$. Therefore, this branch will reach **HALT** and so the entire computation halts.

Conversely, if the computation halts on m . Then there is some S which provides a finite refinement of the cover given by m which covers X . Hence $\bigvee_{i \in I} \bigwedge_{j \in J_i} g_j$ covers X and m halts on all of X , as required. ■

Remark 4.2. The usual way to link compactness to universal quantification is via the characterisation in terms of closed product projections. Thus, it is perhaps remarkable that this algorithm instead makes use of the standard open cover definition.

It is worth discussing **test**($m \in \boxtimes F$) in more detail. Recall that in general an open merely suggests the existence of a procedure to semi-decide some

membership and is not the procedure itself. Hence it is not a priori clear that there is some uniform way to compute $\mathbf{test}(m \in \boxtimes F)$. However, note that Σ^{Σ^G} is locally compact and hence we expect there to be continuous assignment of opens to machines.

In Model 3.2 the procedure associated to $\boxtimes F$ is simply the act of observing the robot and comparing the generators it visits to the finite collection F . It is clear that this can be done uniformly in F . Alternatively, we might regard a machine as a program $m: \Sigma^G \rightarrow \Sigma$. To test whether a branch is contained in a subset F , we simply supply as input the generalised point $p_F: G \rightarrow \Sigma$ which halts precisely on the elements of F .

Example 4.3 (Cantor space). Cantor space is the space $2^{\mathbb{N}}$ of infinite binary sequences. A presentation is given by $\langle z_n, u_n, n \in \mathbb{N} \mid z_n \wedge u_n = 0, z_n \vee u_n = 1 \rangle$. Intuitively, z_n should be thought of as the open consisting of the sequences whose n^{th} digit is 0 and u_n the open of sequences with 1 in the n^{th} position. A machine will be given by a formal expression of the form $\bigvee_{i \in I} \left(\bigwedge_{j \in J_i} z_j \wedge \bigwedge_{k \in K_i} u_k \right)$.

Note that a we can easily check if a *finite* join $\bigvee_{i \in I} \left(\bigwedge_{j \in J_i} z_j \wedge \bigwedge_{k \in K_i} u_k \right)$ is a cover of $2^{\mathbb{N}}$ by distributing the joins over the meets and checking each conjunct a cover. A conjunct will be a cover if and only if it contains both z_i and u_i as disjuncts for some $i \in \mathbb{N}$.

The implementation of $\mathbf{test}(m \in \boxtimes F)$ depends on the precise model. In any case, we arrive at an algorithm which achieves the same ends as that given by Escardó in [5, 8]. \triangle

We will explore how our approach relates to other representations of spaces and the link to Escardó's setting in Section 4.3.

Example 4.4 (The closed interval). Consider the case of the closed real interval $X = [0, 1]$. We can take a generating set consisting of the open intervals with rational end points, $G = \{(a, b) \mid a, b \in \mathbb{Q}, 0 \leq a < b \leq 1\}$. Then in the computational model, $[0, 1]$ embeds as a subset of partial functions on G and so a real r corresponds to a function that halts on (a, b) if and only if $r \in (a, b)$. Then for a 'machine' $m: \Sigma^G \rightarrow \Sigma$ the algorithm halts if m halts on all of the (functions corresponding to) reals in $[0, 1]$. It is elementary to compute when a formal join of rational intervals covers $[0, 1]$. Finally, $\mathbf{test}(m \in \boxtimes F)$ is given by simply calling $m(p_F)$ where $p_F \in \Sigma^G$ is given by $p_F(g) = \top \iff g \in F$. \triangle

Another element to discuss further is the process of determining whether $\bigvee_{F \in S} \bigwedge_{g \in F} g \sim 1$ with respect to R . We must do this by hand in order to construct the explicit algorithm for a given locale. In concrete examples this is typically straightforward, but ideally we might hope to have a general algorithm that does this for us. However, generating a congruence from given relations involves a transfinite procedure and is hence a computationally nontrivial process.

Here it is natural to consider links to *formal topology*, which is a predicative approach to topology that considers as its primary objects of study structures which correspond to presentations of locales. Of relevance here are the paper [3], which discusses inductive generation of topologies, and [30], which covers compactness and the relation to locale theory. Theorem 15 of [30] gives conditions in terms of the generators for a set of formal finite joins of opens to be the set of all finite covers of a locale. Moreover, Proposition 11 of the same paper implies that the presentation of a compact locale may always be given in a way such that compactness is manifest and in which case it is easy to find the finite covers algorithmically. In general, a proof of compactness is likely to lead to a description of the finite covers.

Remark 4.5. Constructively, there is a ‘dual’ notion to compactness called *overtness* (see [25, 13]) which has the same relation to existential quantification as compactness has to universal quantification. One can give an algorithm for existential quantification that is very similar to Algorithm 1 except that instead of iterating over covers we iterate over ‘positive’ (classically: non-zero) elements. Classically, every locale is overt, though the resulting algorithm is still more subtle than naively searching through every real number in turn. Instead, it would use the fact that non-trivial open intervals with rational endpoints form a countable base for \mathbb{R} , which has the advantage of avoiding unbounded parallelism.

4.2. Topological consequences. It is possible to unwind Algorithm 1 to give an open of Σ^{Σ^G} . Each expression $\mathbf{test}(m \in \boxtimes F)$ is replaced with the corresponding open $\boxtimes F$. The second **for** loop takes the meet of these opens, $\bigwedge_{F \in S} \boxtimes F$. The first **for** (together with the **if**) corresponds to taking the join over all $S \in \mathcal{P}_{\text{fin}}(\mathcal{P}_{\text{fin}}(G))$ such that $\bigvee_{F \in S} \bigwedge_{g \in F}$ covers X . Let us denote the set of such S by \mathcal{C} . Thus, the resulting open is $\bigvee_{S \in \mathcal{C}} \bigwedge_{F \in S} \boxtimes F$. When X is compact, this is the open of machines which cover X . The following result is then immediate from Theorem 4.1.

Corollary 4.6. *Let X be a compact locale with $\mathcal{O}X = \langle G \mid R \rangle$ and let \mathcal{C} the set $S \in \mathcal{P}_{\text{fin}}(\mathcal{P}_{\text{fin}}(G))$ satisfying $\bigvee_{F \in S} \bigwedge_{g \in F} g = 1$ in X . Then the open $\bigvee_{S \in \mathcal{C}} \bigwedge_{F \in S} \boxtimes F$ contains precisely the machines which cover X .*

Note that we can still define this open for non-compact X ; however, in general it will only contain machines which correspond to covers which have a finite subcover.

Also note Corollary 4.6 easily generalises to arbitrary compact sublocales K of a general locale X by simply viewing K as direct sublocale of Σ^G .

Remark 4.7. Corollary 4.6 also follows from the Hofmann–Mislove theorem [11]. Moreover, the Hofmann–Mislove theorem also provides a converse giving that the set of all covering machine of X is Scott-open if and only if X is compact. (See also [5, Lemma 7.4].) However, we believe restriction to machine space and the explicit algorithm provide an interesting new perspective on the situation.

Remark 4.8. Of course, in the overt case we can similarly extract an open of all ‘positive’ machines for X from the algorithm mentioned in Remark 4.5.

4.3. Link to domain-theoretic approaches. Escardó approaches his compactness algorithm from the point of view of programming language semantics and domain theory. It is interesting to consider how our more topological approach compares to this one. (Until now our pointfree results have been constructively valid. For simplicity, we will work classically in this section.)

Recall that data types in programming languages can be modelled as certain kinds of dcpos. The dcpos that arise in denotational semantics are usually *continuous*, in which case the Scott topology is particularly well-behaved. We also note that resulting topological spaces are locally compact.

Escardó considers the data type of infinite binary sequences. We imagine these are implemented as streams where the terms are computed in order one by one on demand. Such a data consists not only of infinite sequences, but also additional elements corresponding to sequences whose first n terms are defined, but for which accessing later terms in the sequence causes the computation to hang.

The elements of this data type form a dcpo ordered by definability so that, for instance, a sequence which hangs after producing $0, 1, 0$ is smaller than one which hangs after $0, 1, 0, 0$. Equipping this with the Scott topology we obtain a space containing Cantor space as the subspace of completely defined sequences.

Our embedding of a locale X into the space Σ^G for a set of generators G can be thought of as a reasonably canonical way to obtain additional ‘partially defined’ elements in a similar way. In fact, since the space Σ^G is simply the power set $\mathcal{P}(G)$ (which is a continuous dcpo) equipped with the Scott topology, it can be understood as corresponding to the data type of partial functions from a discrete data type with $|G|$ elements to the unit type.

In this way our construction is an example of a *domain embedding* where a topological space is represented as a subquotient of a dcpo. In the domain theory literature, there is some desire for spaces to occur as the subspace of maximal elements of a dcpo, but this not the case for our embedding, and indeed, this impossible if the space is not T_1 . However, we can get closer to this ideal if we take the closure of X as a subspace in $\mathcal{P}(G)$. A closed subset of a dcpo is a downward closed set which is closed under directed joins. Moreover, a closed subspace of a continuous dcpo is a continuous dcpo and the Scott topology on the subset agrees with the subspace topology. Thus, the closure of X in $\mathcal{P}(G)$ is a dcpo which contains the points of X and possibly some ‘more undefined’ approximations to them and since this closure is the smallest Scott-closed set containing X , the points of X lie as close to the top of the dcpo as possible in some sense.

Example 4.9. Let us consider the case of Cantor space from Example 4.3. Recall that $\mathcal{O}(2^{\mathbb{N}})$ can be expressed as a quotient of $\mathcal{O}(\Sigma^G)$ where $G = \{z_n \mid n \in \mathbb{N}\} \sqcup \{u_n \mid n \in \mathbb{N}\}$ by the congruence $C = \langle z_n \wedge u_n = 0, z_n \vee u_n = 1, n \in \mathbb{N} \rangle$. The closure of such a sublocale corresponds to the congruence generated by the pairs $(0, c) \in C$. In this case we obtain a the following presentation for the closure: $\langle z_n, u_n, n \in \mathbb{N} \mid z_n \wedge u_n = 0 \rangle$. This describes a dcpo the elements of which can be thought of as *partial* functions from \mathbb{N} to $2 = \{0, 1\}$. The points of Cantor space correspond to the total functions, which do appear the maximal elements of this dcpo. \triangle

In general, the compactness algorithm works equally well if we replace $\mathcal{P}(G)$ with such a continuous dcpo D , or indeed any locally compact sublocale of machine space which contains X . The algorithm takes an element m of Σ^D as input and simply applies Algorithm 1 to $s(m)$ where s is a section of the quotient $\Sigma^{\Sigma^G} \rightarrow \Sigma^D$ from Proposition 3.6.

Note that the only place $s(m)$ appears in the algorithm is at the step where we run $\mathbf{test}(s(m) \in \boxtimes F)$. Thus, we may replace this step with any equivalent computation involving m .

In this setting it is natural to consider an intuitive model of such a dcpo as a data type in an idealised programming language. If we view Σ^{Σ^G} from this perspective recall that $\mathbf{test}(m \in \boxtimes F)$ is given by $m(p_F)$ where $p_F: G \rightarrow \Sigma$ halts precisely on the elements of F . Other examples can be handled in a similar way.

Example 4.9 (Continued). Let us return to the example of Cantor space. Using the dcpo A of partial functions from \mathbb{N} to 2 instead of machine space we have a compactness algorithm where $\mathbf{test}(s(m) \in \boxtimes F)$ is implemented as follows: for each $E \in \mathcal{P}_{\text{fin}}(F)$ which does not contain both u_n and z_n for any n , we run $m(f_E)$ where

$$f_E(n) = \begin{cases} 0 & \text{if } z_n \in F, \\ 1 & \text{if } u_n \in F, \\ \perp & \text{otherwise} \end{cases}$$

and halt if any of these do.

The inclusion $i: A \rightarrow \Sigma^G$ maps f to a function which halts on z_n if $f(n) = 0$ and halts on u_n if $f(n) = 1$. The quotient map $q: \Sigma^{\Sigma^G} \rightarrow \Sigma^A$ sends m to a map $x \mapsto m(i(x))$. Note that for a given $m \in \Sigma^A$, the above procedure $\mathbf{test}(s(m) \in \boxtimes F)$ for $F \in \mathcal{P}_{\text{fin}}(G)$ can be easily extended to infinite $F \in \Sigma^G$. This gives the associated section $s: \Sigma^A \rightarrow \Sigma^{\Sigma^G}$ of q . \triangle

The above representation of Cantor space is not quite the one used by Escardó, since he views Cantor space not as a function, but as a *stream* — that is, there is an ordering on indices so that if the value at i is defined, so are all previous values. This is related to alternative presentation of Cantor space in terms of prefixes.

Example 4.10. A presentation for $2^{\mathbb{N}}$ can be given with generators ℓ_p corresponding to sequences with finite prefix $p \in 2^*$. We have

$$\begin{aligned} \mathcal{O}2^{\mathbb{N}} = \langle & \ell_p, p \in 2^* \mid \ell_p \wedge \ell_q = \ell_q \text{ for } p \prec q, \\ & \ell_p \wedge \ell_q = 0 \text{ if } q \not\prec p \text{ and } p \not\prec q, \\ & \ell_{p\#0} \vee \ell_{p\#1} = \ell_p \text{ for all } p \rangle \end{aligned}$$

where $p \prec q$ means p is a prefix of q and $\#$ denotes concatenation.

However, taking the closure of this sublocale of Σ^{2^*} does not yield the correct the dcpo. This is because it ignores the natural order structure on the prefixes. The solution is to modify the parent space Σ^G , so that instead of taking the free frame on a *set* G , we take the free frame on a poset. In this case, we

equip 2^* with the reverse of the prefix order. Then $\mathcal{O}2^{\mathbb{N}}$ is the quotient of this frame by the congruence generated by the second and third relations above (since the first relation is already handled by the order structure on G). The closure of this sublocale then has the presentation $\langle \{\ell_p : p \in (2^*)^{\text{op}}\} \in \text{Pos} \mid \ell_p \wedge \ell_q = 0 \text{ if } q \not\prec p \text{ and } p \not\prec q \rangle$. It is not hard to see that this is the locale of finite and infinite sequences binary sequences. (Indeed, it is the third relation in the presentation for $2^{\mathbb{N}}$ above that forces all sequences to be infinite.) This is precisely the dcpo considered by Escardó. (Finite sequences are interpreted as sequences which hang after some point.)

We can now see how our algorithm reduces to a (non-optimised) version of Escardó's algorithm.

Let us first consider how to check if a finite join $\bigvee_{i \in I} \left(\bigwedge_{p \in J_i} \ell_p \right)$ is a cover of $2^{\mathbb{N}}$. We first compute the finite meets using the first two relations to obtain a join of basic generators (omitting zeros from the join). Now let N be the length of the longest prefix. We pad each p in the join with 0s and 1s to obtain all possible strings of length N with prefix p . The join is a cover if and only if the resulting set contains every string of length N .

Finally, $\text{test}(s(m) \in \boxtimes F)$ is computed as follows. If $\bigwedge F = 0$ there is nothing to do. Otherwise, $\bigwedge F = \ell_p$ and we run m on the finite sequence p . \triangle

Remark 4.11. The locale corresponding to free frame on a poset P is $\Sigma^{\mathcal{I}P}$ where $\mathcal{I}P$ is the *ideal completion* of P . In the case above, the ideal completion of $(2^*)^{\text{op}}$ is actually already the desired dcpo. Nonetheless, it is still interesting to see that this fits into our general approach.

References

- [1] S. Abramsky. Domain theory in logical form. *Ann. Pure Appl. Logic*, 51(1-2):1–77, 1991.
- [2] R. Cartwright, R. Parsons, and M. AbdelGawad. Domain theory: An introduction, 2016. arXiv preprint, [arXiv:1605.05858](https://arxiv.org/abs/1605.05858).
- [3] T. Coquand, G. Sambin, J. Smith, and S. Valentini. Inductively generated formal topologies. *Ann. Pure Appl. Logic*, 124(1-3):71–106, 2003.
- [4] B. J. Day and G. M. Kelly. On topological quotient maps preserved by pullbacks or products. *Math. Proc. Cambridge Philos. Soc.*, 67(3):553–558, 1970.
- [5] M. Escardó. Synthetic topology: of data types and classical spaces. *Electron. Notes Theor. Comput. Sci.*, 87:21–156, 2004.
- [6] M. Escardó. Notes on compactness, 2005. Unpublished notes, <https://www.cs.bham.ac.uk/~mhe/papers/compactness.pdf>.
- [7] M. Escardó. Infinite sets that admit fast exhaustive search. In *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 443–452. IEEE, 2007.
- [8] M. Escardó. Exhaustible sets in higher-type computation. *Log. Methods Comput. Sci.*, 4, 2008.

- [9] B. Fawcett and R. J. Wood. Constructive complete distributivity I. *Math. Proc. Cambridge Philos. Soc.*, 107(1):81–89, 1990.
- [10] J. Goubault-Larrecq. *Non-Hausdorff topology and domain theory*, volume 22 of *New Mathematical Monographs*. Cambridge University Press, 2013.
- [11] K. H. Hofmann and M. W. Mislove. Local compactness and continuous lattices. In B. Banaschewski and R.-E. Hoffmann, editors, *Continuous lattices*, pages 209–248. Springer, 1981.
- [12] P. T. Johnstone. *Stone spaces*, volume 3 of *Cambridge studies in advanced mathematics*. Cambridge University Press, Cambridge, 1982.
- [13] A. Joyal and M. Tierney. An extension of the Galois theory of Grothendieck. *Mem. Amer. Math. Soc.*, 51(309), 1984.
- [14] F. W. Lawvere. Adjointness in foundations. *Dialectica*, pages 281–296, 1969.
- [15] G. Manuell. *Quantalic spectra of semirings*. PhD thesis, University of Edinburgh, 2019.
- [16] J. Picado and A. Pultr. *Frames and Locales: topology without points*. Springer Science & Business Media, 2011.
- [17] A. M. Pitts. Notes on categorical logic, 1991. Unpublished lecture notes, <https://www.cl.cam.ac.uk/~amp12/papers/notcl/notcl.pdf>.
- [18] T. Plewe. Localic triquotient maps are effective descent maps. *Math. Proc. Cambridge Philos. Soc.*, 122(1):17–43, 1997.
- [19] R. Rosebrugh and R. J. Wood. Constructive complete distributivity IV. *Appl. Categ. Structures*, 2(2):119–144, 1994.
- [20] J. Rosický. Cartesian closed exact completions. *J. Pure Appl. Algebra*, 142(3):261–270, 1999.
- [21] M. B. Smyth. Power domains and predicate transformers: A topological view. In *International Colloquium on Automata, Languages, and Programming*, pages 662–675. Springer, 1983.
- [22] M. B. Smyth. Topology. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 1, pages 641–761. Clarendon Press, Oxford, 1992.
- [23] H. Søndergaard and P. Sestoft. Non-determinism in functional languages. *The Computer Journal*, 35(5):514–523, 1992.
- [24] P. Taylor. Subspaces in abstract stone duality. *Theory Appl. Categ.*, 10(13):301–368, 2002.
- [25] P. Taylor. Foundations for computable topology. In G. Sommaruga, editor, *Foundational theories of classical and constructive mathematics*, pages 265–310. Springer, 2011.
- [26] S. Vickers. Information systems for continuous posets. *Theoret. Comput. Sci.*, 114(2):201–229, 1993.
- [27] S. Vickers. Locales are not pointless. In C. Hankin, I. Mackie, and R. Nagarajan, editors, *Theory and Formal Methods of Computing 94: Proceedings of the Second Imperial College Workshop*, pages 199–216, London, 1995. Imperial College Press.
- [28] S. Vickers. *Topology via logic*. Cambridge University Press, 1996.
- [29] S. Vickers. The double powerlocale and exponentiation: a case study in geometric logic. *Theory Appl. Categ.*, 12(13):372–422, 2004.
- [30] S. Vickers. Compactness in locales and in formal topology. *Ann. Pure Appl. Logic*, 137(1-3):413–438, 2006.
- [31] S. J. Vickers and C. F. Townsend. A universal characterization of the double powerlocale. *Theoretical Computer Science*, 316(1-3):297–321, 2004.
- [32] G. Winskel. A note on powerdomains and modality. In *International Conference on Fundamentals of Computation Theory*, pages 505–514. Springer, 1983.

PETER F. FAUL
INDEPENDENT RESEARCHER
E-mail address: `peter@faul.io`

GRAHAM MANUELL
CMUC, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF COIMBRA, COIMBRA, PORTUGAL
E-mail address: `graham@manuell.me`