

Continuous Domains in Logical Form

Mathias Keigelmann

A thesis submitted
to The University of Birmingham
for the degree of Doctor of Philosophy

School of Computer Science
The University of Birmingham
Birmingham B15 2TT
United Kingdom

June 1999

Abstract

This thesis investigates the mathematical foundations that are necessary for an extension of Abramsky's domain theory in logical form to continuous domains.

We present a multi-lingual sequent calculus, that is a positive logic allowing sequents that relate propositions from different languages. This setup necessitates a number of syntactic adjustments. In particular, we discuss different reformulations of the cut rule and how they can be used as a basis for a category MLS of logical systems. Then we investigate cut elimination in this logic. From a semantic point of view this can be seen as enabling us to perform domain constructions in purely syntactic form.

The category MLS has a number of different manifestations, and we study it with logical, localic, topological and categorical methods. From a topological point of view, we show that MLS is equivalent to the category of stably compact spaces with certain closed relations. By putting together cut elimination and representation theorems for these spaces we get a continuous domain theory in logical form.

Contents

Introduction	7
1 The Basics	15
1.1 Topology and Order	15
1.1.1 Topology	15
1.1.2 Domains	18
1.2 Stone Duality	24
1.2.1 The Adjunction $\Omega \dashv \text{pt}$	24
1.2.2 Sobriety and Spatiality	26
1.2.3 Hierarchy of Stone Dualities I	32
1.3 Stably Compact Spaces	36
1.3.1 The Definition	36
1.3.2 Compact Pospaces	39
1.3.3 Stably Compact Domains	44
1.3.4 Hierarchy of Stone Dualities II	48
1.4 Abramsky's Domain Theory in Logical Form	51
1.4.1 Prelocales	51
1.4.2 Prelocalic Description of Domains	52
1.4.3 Domain Constructions	53
1.4.4 Logic	55
2 Syntax	57
2.1 Multi-Lingual Sequent Calculus	57
2.1.1 The Logic	59
2.1.2 A Category of Consequence Relations	64
2.2 Cut Elimination	70
2.2.1 Simple Elimination	70
2.2.2 Construction of Consequence Relations	73
2.2.3 Coproducts and Products	81
3 Semantics	91
3.1 Logic and Topological Spaces	91
3.1.1 Theories and Models	92

3.1.2	Algebraisation of the Logic	98
3.1.3	Topological Semantics	104
3.1.4	Semantics of Morphisms	116
3.2	Domain Constructions	126
3.2.1	Representing Stably Compact Spaces	126
3.2.2	Individual Constructions	132
3.2.3	Function Spaces	137
3.2.4	Relation Spaces	145
3.3	Relations and Functions	154
3.3.1	Diagonals	154
3.3.2	Regular Categories	159
3.3.3	Closed Relations	165
	Conclusion	167
	Index	169
	Categories	169
	Notation	171
	Bibliography	173

Introduction

Domains are used in the theory of semantics of programming languages to give *denotational* models for programs. In this approach a data type corresponds to a structured set, called a *domain*, and a program is interpreted as a function between the domains corresponding to the input and the output types of the program. One can then employ mathematical reasoning on the domain theoretic side to get insights into the behaviour of the program. This approach goes back to the early work of Scott [Sco70] and [Sco76]. Domain theory also has strong links to other areas of mathematics, most notably T_0 topology [GHK⁺80]. One reason for this is that the specialisation order on every sober space is a dcpo. Moreover, the lattice of open sets in a topological space is a particular dcpo and so domain theoretic methods are often useful for their analysis.

Another branch of semantics, *program logics*, studies logical systems that either describe properties of programs or how fragments of programs transform them. The best known example is probably Hoare logic [Hoa69] which studies triples

$$\{P\} \ S \ \{Q\}$$

where P and Q are predicates and S is a statement or a fragment of a program. The intended reading of such a triple is that if P holds before the execution of S , then Q holds afterwards. In [Abr87] and [Abr91b] Abramsky explicates the connection between denotational semantics and program logics via his programme of *domain theory in logical form*. The basic idea is to utilise the correspondence between topology and certain logics.

From the computer science perspective, topology enters the theory in several ways. For one, functions between domains arise as the semantics of programs and the continuity of these functions is an abstract way of capturing computability. The Scott topology, the topology usually considered on domains, encodes the idea that to produce any output a computable function can only look at a finite portion of its input. We can also give some meaning to the topology on any space, namely as a *logic of observable properties*. This logic has arbitrary disjunctions but only finite conjunctions, analogous to the definition of a topology; sometimes this is also referred

to as *geometric logic*. The asymmetry is due to the fact that an infinite disjunctions of properties can be observed by witnessing any one of them, testing an infinite conjunction, on the other hand, requires infinitely many ‘experiments’ rendering such a property non-observable in general. So, the open sets of a topological space may be identified with *observable* or *semi-decidable properties* of its points. This idea was introduced by Smyth [Smy83b]. For a much more thorough discussion of these issues see [Vic89, Chapter 2], [Abr87] and [Smy92b]. Coming back to domains we conclude that the propositions of a program logic should correspond to Scott-open sets. Program fragments, in the form of continuous functions, act on these propositions by taking the preimage and can thus be understood as *predicate transformers* in the sense of Dijkstra’s weakest precondition calculus [Dij76].

Topology and the study of the logic of observable properties are essentially the same thing. We have just seen how we can go from a topological space to such a logic. Conversely, for every logic of observable properties we can look at its Lindenbaum algebra, the quotient of the logic under equivalence. This algebra is a lattice with arbitrary suprema and finite infima, which are connected by a distributivity law. Such lattices are called *frames* or *locales*. This is the starting point of *locale theory* [Joh82] which can also be seen as topology without points [Joh83]. Among other things, locales have the advantage that free locales exist and that one can thus construct them from generators and relations. The connection between topological spaces and locales is the subject of *Stone duality* which says that certain categories of spaces are equivalent to certain categories of locales.

Stone duality is the main foundation of domain theory in logical form as pioneered by Abramsky [Abr87]. A domain can be seen as a topological space by endowing it with the Scott topology, thus also turning it into a locale. A logical description of the domain is a logic of observable properties, a so-called *prelocale*, such that its Lindenbaum algebra generates this locale. Now suppose we take the product or the function space of domains, form a power domain or perform any other domain construction. If we already have prelocales corresponding to these domains, then the question is how we can translate such a construction into prelocalic terms. The aim is to come up with *purely syntactic rules* to build a new prelocale from the given ones such that it corresponds to the resulting domain. In [Abr91b] Abramsky gives them for the domain constructions that are commonly considered in domain theory. On top of the logic and the constructions for prelocales one can build different logics to talk about elements in different domains or about functions between them. We discuss some details of Abramsky’s programme briefly in Section 1.4.

There are several benefits of this logical viewpoint. On the one hand, it yields new insights for domain theory because we now have to understand

how a construction acts on properties rather than on points. The localic description is sometimes even smoother than the domain theoretic one, e.g. in the case of the Plotkin power domain and for bilimits. This new angle may also help us understand why and how certain constructions work. Furthermore, there are concrete application of this theory: If we have a good denotational semantics of a phenomenon we want to model, this theory gives us a corresponding logic almost for free. Its advantage over *ad hoc* formalisms is that the close connection to the denotational model guarantees that the logic is equally suitable for reasoning about the particular system. For examples of this technique see [Abr91a] and [Abr90].

The theory as described by Abramsky does not encompass all domains; his notion of prelocale corresponds to a class of certain *algebraic* domains, the so-called bifinite or SFP domains. In fact, most of the classical domain theory in computer science used to be focused on algebraic domains. These are domains where every element can be approximated by compact elements which in turn can be thought of as finite pieces of information. This is good enough for many applications since most programs process discrete data and we are often only interested in discrete properties like termination or correctness. However, apart from their mathematical appeal, it has long been argued that *continuous domains* have a similar importance (see for example [Jun90]). The reasons for this are mainly real number computation, the modelling and computational analysis of other continuous mathematical structures, and probabilistic power domains which are used to study non-deterministic and stochastic phenomena.

We consider some examples of activities in these fields. In [Esc96] Escardo studies PCF with an additional ground type for real numbers. In another paper he and Edalat look at integration in this extension of PCF [EE96]. This is based on Edalat's research of computational measure theory [Eda95] and integration [Eda94] using continuous domains. Apart from real number computation, which we have already mentioned, this theory has applications in theoretical physics, neural networks and fractal image compression. Edalat's approach uses the probabilistic power domain which was introduced by Plotkin and Jones [JP89]. It is a space of valuations which are the domain theoretic analogue of a measure. The original purpose of this power domain was to give a model for a non-deterministic programming language [Jon90]. It can also be used to model other types of stochastic behaviour. Sünderhauf studies computational models for uniform totally bounded spaces [Sün94]. His models are dcpo's which are unfortunately not necessarily continuous. Edalat and Heckmann give a domain theoretic description of metric spaces [EH95], and Lawson shows in [Law97] that every Polish space is a *maximal point space*, i.e. homeomorphic to the maximal points of a continuous domain. Hence, continuous domains can be used as computational models for classical spaces.

Continuous domain theory is an active field, as the previous paragraph illustrates, and its importance is likely to grow as its applications become more widespread. This is the motivation for trying to extend the domain theory in logical form to continuous domains. In a nutshell, this is the main aim of this thesis. It is worth mentioning that also from a mathematical point of view it is interesting to extend the class of spaces considered. The study of continuous domains often unravels certain features of the theory that are not prominent in the algebraic case. A good illustration of this is the paper [JS96] which can be seen as the first steps of *continuous domains in logical form*. There the dual roles played by compact saturated and open subsets become apparent, something which has been obscured in the algebraic case.

There are some other developments starting from Abramsky's theory. Zhang studies the logic of stable domain theory [Zha91], Bonsange gives a Stone-type duality for non-sober spaces and discusses the link to an infinitary logic with arbitrary conjunctions and disjunction [Bon96]. In [RB96] Brink and Rewitzky explore the connection with Priestley duality to study the exact relationship between information systems, power domains, predicate transformers and relational models.

One advantage of restricting domain theory in logical form to algebraic domains is the following: Their Stone duals are algebraic lattices which have the pleasant property that the compact elements form a sub-lattice. This unlocks the door to a finitary, localic description of these spaces. In the continuous case the Stone dual is only a certain kind of continuous lattice and as such no longer has a canonical basis. In their papers [JS96] and [JS98] Jung and Sünderhauf show how one can, nonetheless, get a good description of such spaces via *strong proximity lattices*. These structures have an order of approximation in addition to their (logical) lattice order. While these two papers were motivated by purely topological considerations their findings are the starting point for a continuous domain theory in logical form. The connection to our setup as explained in Section 3.1.2 and the results of Section 2.2 show that they have a logical and proof-theoretic content. In particular, this constitutes an independent justification for the two axioms which distinguish strong proximity lattices from the structures studied in [Smy92a].

The spaces considered in [JS96] are the *stably compact spaces* which are compact, locally compact, sober spaces such that the intersection of two compact saturated subsets is again compact. In the literature they appear in many guises and under a number of names. In [AJ94] and [JS96], for example, they are called 'coherent'. This term, however, might be a possible source of confusion: The spaces corresponding to the 'coherent' locales of [Joh82] and [Vic89] are called 'spectral' in the latter book. Reserving the use of the term 'coherent' for algebraic locales is justified by the link to 'coherent logic' which is a well-established term. As there is also the danger of

mixing them up with Girard's 'coherence spaces' we use the more descriptive term 'stably compact'. The domains that are normally used in semantics like Scott domains, bifinite and FS domains are stably compact with respect to their Scott topology. The stably compact spaces are also closely related to the compact ordered spaces introduced by Nachbin [Nac48], and they can be seen as the T_0 equivalent of compact Hausdorff spaces.

This thesis is a mathematical investigation of different aspects of a logical description for the class of stably compact spaces and extends Abramsky's work in three directions. The first one is syntactic and goes back to a suggestion by Andrew Moshier. The idea is to see the auxiliary relation on strong proximity lattices as the entailment \vdash of a sequent calculus. This allows us to understand the completeness proofs for individual domain constructions in Abramsky's programme as cut elimination. To make this precise we introduce a *multi-lingual sequent calculus*, i.e. is a logic that considers sequents between different languages or worlds. From the semantic point of view we can think of them as logical descriptions of different domains. But the calculus can also be motivated independently as one that allows reasoning about inferences between different areas of reasoning. The logic is very weak; the only logical connectives are conjunction and disjunction, and we exclude the identity axiom, while retaining all structural rules. Given this set-up we consider an alternative to Gentzen's cut rule and explore how it interacts with the logical rules and how it relates to other cut rules. We argue on purely proof theoretic grounds that consequence relations \vdash should be interpolative. This property is well known from semantics, where it is justified by reference to effective computability. These results are then used to define a category MLS of logical systems which we study from a number of different angles. There is also a good theory of cut elimination in the multi-lingual setting. This can be used to perform domain constructions in purely logical form: Rather than showing completeness with respect to a domain which is known to satisfy a certain universal property, we can perform and verify the construction directly in MLS.

The two other directions come out of the semantics of the multi-lingual sequent calculus. We study its model theory in the style of Stone's work on Boolean algebras, which in spite of the ostensible simplicity of our logic is quite intricate. The category MLS turns out to be equivalent to stably compact spaces with certain relations between them. Taking a closer look at this proof, we get a representation theorem that characterises when a sequent calculus in MLS encodes a given stably compact space. It can also be used to construct such calculi: Given a space X and a 'language' together with open and compact interpretations of its formulae in X , the theorem tells us how we have to define entailment to turn it into a logical description of X . The problem is then to come up with inductive rules that gener-

ate this logic. This technique is applied to a number of domain constructions, thus extending domain constructions in logical form to the category of all stably compact spaces which contains a large class of continuous domains. Moreover, we have two methods to perform such constructions: One, purely syntactic, using cut elimination and a second one which relies on the topological semantics. The constructions that we perform as examples for either method show that both options are feasible. Purely logical constructions avoid the heavy topological machinery to some extent, but there is quite some overhead, as the construction of coproducts shows. Doing this example in full detail allows us to see exactly what is involved and how it compares to the semantic approach.

As mentioned before, a morphism in MLS corresponds to a relation—rather than a continuous function—between stably compact spaces, and this is the third direction in which the present work extends Abramsky’s. The relations that arise in this way can be characterised topologically and are in one-to-one correspondence to topological set-valued functions. It is not clear whether it is possible to give a logical function space construction in the current setup, but for the relation space, on the other hand, the situation is very satisfactory. We discuss the corresponding construction semantically and syntactically, and also identify the universal property that defines it. This almost makes MLS a symmetric monoidal category, but the mediating morphism is only unique as a function but not as a relation. To make this precise we need a characterisation of morphisms that are functions in disguise. This can be done in syntactic, localic, topological or categorical terms.

The thesis brings together ideas from proof theory, topology, domain theory and category theory, and some familiarity with the basic concepts from these fields is presupposed. Our principle reference is the handbook article [AJ94] which covers domain theory and some aspects of order theory and T_0 topology. When we need more advanced ideas from topology, locale theory or other areas, references are given in the text. We feel free to use the language of category theory throughout the thesis; a good introduction of the necessary material are the first chapters of [Mac71] or [McL92]. As a convention we compose functions from right to left, i.e.

$$\cdot \xrightarrow{f} \cdot \xrightarrow{g} \cdot$$

is written as $g \circ f$, composition of relations and consequence relations \vdash are from left to right. This may be slightly confusing, but it is in accordance with the usual mathematical practice.

The outline of the thesis is as follows: Chapter 1 provides the technical background for the thesis. It reviews material that can be found in the literature and contains no new results. The only original contribution of

this chapter is the short note on complete distributivity in Section 1.3.3, as well as the organisation of the exposition and maybe some of the proofs. The chapter does not discuss all the basics needed in the rest of the thesis, only those that lead to the logical description of stably compact spaces. Later we also require information about exponentials of topological spaces and regular categories, but as they are not part of the main thrust of this work we discuss them when we actually need them, that is at the beginning of Sections 3.2.3 and 3.3.2, respectively. They can be seen as continuations of Chapter 1, and the same disclaimers apply.

The first two sections cover basic T_0 topology, domain theory and Stone duality, mainly to fix notation. For this reason many proofs are omitted. We only include them if they are not easily found in the literature or if they contain an interesting idea that is useful for the subsequent development of the theory. Many subsections give explicit pointers to the literature where more details can be found. If a theorem is given without proof and no other reference is given, then it can be found there.

Section 1.3 introduces stably compact spaces. As their logical description is at the centre of this thesis and the relevant details are scattered over a number of books and articles in the literature, we discuss them in quite some detail including all the proofs. The section is meant as an exposition of the basic theory of stably compact spaces, their Stone duality, stably compact domains and the link to compact pospaces. The latter is interesting from a mathematical standpoint as it shows the connection to classical spaces, but they are also an important technical tool and we make heavy use of them in Section 3.1.4.

In the last section of the introductory chapter we give a quick overview of Abramsky's work. Its main purpose is to allow the reader who is not familiar with it to see where the approaches differ and in what sense and to what degree Abramsky's programme has been extended successfully to continuous domains.

In Chapter 2 we take a completely different point of view. In the first section we study the multi-lingual sequent calculus and how we can turn it into a category. We can think of the objects and morphisms in this category MLS as sequent calculi for internal reasoning and for reasoning about inferences between different logical domains, respectively. There are a number of interesting properties such sequent calculi can have and we investigate to some extent how they interact.

Section 2.2 discusses cut elimination for this logic and explains how it can be used to construct new objects and morphisms from old ones. As an application we construct products and coproducts in a worked example. In the light of the last chapter this can be understood as performing domain constructions in purely logical terms.

The last chapter ties all these ideas together. Section 3.1 contains the core of this thesis and is technically the most demanding part. It shows

that we can think of MLS as a syntactic description of the category of stably compact spaces and certain closed relations or multi-functions between them. This is done by proving that these categories are equivalent. The category has a number of syntactic, localic and topological manifestations and we discuss their relationships.

In the second section we apply this to do domain constructions. First we characterise when an object from MLS represents a stably compact space. This is the key to doing concrete domain constructions in logical form and we go through a number of them.

The final section describes a categorical way of characterising the closed relations between stably compact spaces that actually correspond to continuous functions. This complements earlier results of Section 3.1.4 where we do this in logical, localic and topological terms.

Throughout the thesis there are *remarks* typeset in a slanted font. They offer a different, often categorical, angle on the topic under consideration. Their purpose is to provide a different point of view or to mention connections to other fields, usually without proofs. Other results in the thesis are always logically independent from these remarks.

Chapter 2 and a large part of Section 3.1 are based on the conference paper [JKM97]. The material has been improved in several points and some of these improvements will be included in the journal version [JKM99].

Acknowledgements

I would like to mention a number of people who contributed in different ways to this thesis. First of all I want express my gratitude towards my supervisor Achim Jung. Whenever I had an idea or a problem he went out of his way to make time for discussions. I also wish to thank the other members of my thesis group, Valeria dePaiva and Marta Kwiatkowska for their advice.

Andrew Moshier suggested the multi-lingual sequent calculus and the new cut rule. Moreover, the insight that cut elimination gives a proof-theoretic account of domain constructions in logical form is his.

I used Paul Taylor's macro packages `diagrams` and `prooftree` for typesetting this thesis. Paul also pointed out the connection to splitting idempotents mentioned on page 69. Peter Selinger introduced me to diagonals as a means of identifying functions in categories of relation as discussed in Section 3.3.1.

For proofreading I want to thank Thomas Erker and Paola Maneggia. In particular Thomas read large parts of this thesis, checked the proofs and made a number of valuable suggestions to improve it.

Finally, I thank Alessio Lomuscio and Axel Großmann for their friendship and support.

Chapter 1

The Basics

The material in this chapter provides the necessary prerequisites for most of this thesis. We begin with elementary topology and order theory, mainly to fix notation. Then we discuss Stone Duality in some detail as it is the foundation of domain theory in logical form. This allows us to define stably compact spaces, the class of topological spaces at the heart of this investigation. We study them from the point of view of topology, locale theory and domain theory. A quick review of Abramsky's original domain theory in logical form finishes the preliminaries.

1.1 Topology and Order

I assume that the reader is acquainted with basic topology, lattices, order theory and domain theory. The principal reference for the latter two is [AJ94]. I adhere to the terminology used there, and rely on familiarity with the material of its Chapters 2 and 3 which discuss basic facts and constructions for domains.

A good general introduction to order and lattice theory is [DP90]. As a reference for topology the handbook article [Smy92b] is particularly well suited for our purposes as it emphasises T_0 spaces.

1.1.1 Topology

Our motivating example of a topological space is a dcpo with the Scott topology. Unless such a dcpo is discrete the resulting topological space is not Hausdorff, or even T_1 , which means that the spaces at the centre of our attention will only satisfy the T_0 separation axiom.

Hence, we start by recalling some definitions and very basic facts of T_0 -topology. As these concepts do not appear in the Hausdorff case, they are less well known and usually not covered in textbooks on topology.

T_0 Spaces

For every topological space X we can define the *specialisation preorder* by

$$\begin{aligned} x \sqsubseteq y & : \iff \overline{\{x\}} \subseteq \overline{\{y\}} \\ & \iff x \in \overline{\{y\}} \\ & \iff \mathcal{N}(x) \subseteq \mathcal{N}(y) \\ & \iff (\forall U \in \mathcal{N}(x)) y \in U, \end{aligned}$$

where $\mathcal{N}(x)$ is the neighbourhood filter of the point x .

The connection between this preorder and convergence is as follows: If a net or a filter converges to a point x and $y \sqsubseteq x$ then y is also a limit.

The specialisation preorder is antisymmetric, and hence an order relation, if and only if the space is T_0 . As mentioned above we will focus on such spaces and, thus, will not state the following observations for the pre-ordered case. A topological space satisfies the T_1 axiom if and only if the specialisation order is trivial, i.e. the equality on the space.

Note that all open sets of a topological space are upper sets with respect to the specialisation order, and the closure of a point can be written as $\overline{\{x\}} = \downarrow x := \{y \mid y \sqsubseteq x\}$.

The order of specialisations also allows us to order the functions between two given topological spaces. We say $f \sqsubseteq g$ in the so called *extensional order* if this is the case point-wise, i.e. for all x we have $f(x) \sqsubseteq g(x)$.

Remark. *It is easily verified that every continuous function is automatically monotone for the respective preorders of specialisation. Hence, going from a topological space to a preordered set by using the specialisation preorder defines a functor. This functor restricts and co-restricts to the categories of T_0 spaces and of posets.*

In fact, using the extensional order on hom-sets it is easy to see that topological spaces form a (pre-)order enriched category.

In general, several topologies can induce the same order of specialisation. The coarsest such topology is the *lower topology* which has closed sets generated by principal ideals $\downarrow x$. The finest is the *Alexandrov topology* which has all lower sets as closed subsets.

Remark. *The Alexandrov topology gives rise to the left adjoint to the functor that takes the specialisation order. The lower topology, however, is not the right adjoint.*

We call a subset A of topological space *saturated* if it is an intersection of open sets. This is equivalent to A being equal to the intersection of all open sets that contain it. Using the specialisation preorder we can also characterise the saturated sets as being exactly the upper sets with respect to the specialisation order. Hence, in a T_1 space every subset is saturated. Note that arbitrary unions and intersections of saturated sets are again saturated.

Any subset A of a topological space has a saturation, i.e. a smallest saturated subset containing it, which is given by the intersection of all open supersets of A . In fact, the intersection of any basis—not necessarily made up from open sets—of the neighbourhood filter of A yields the saturation of A . Using the language of the specialisation order this saturation is also given by $\uparrow A := \{x \mid (\exists a \in A) a \sqsubseteq x\} = \bigcup_{a \in A} \uparrow a$.

Compactness

For us the term *compactness* comprises the Heine-Borel but not the Hausdorff property. As a consequence *local compactness* has to be defined to mean that for every point x and every neighbourhood $U \in \mathcal{N}(x)$ there is a compact neighbourhood $K \in \mathcal{N}(x)$ such that $K \subseteq U$. In the non-Hausdorff case this is strictly stronger than every point having a compact neighbourhood. Also, for a T_0 space compactness does not imply local compactness.

Compactness is well-behaved under saturation.

Lemma 1.1.1. *A subset of a topological space is compact if and only if its saturation is.*

Proof. This is an immediate consequence of the following observation: A union of open sets covers a given set K if and only if it covers the saturation $\uparrow K$. \square

As a corollary we get that a space is locally compact if and only if every point has a neighbourhood filter basis of compact saturated sets. Note that this implies that the neighbourhood filter of every compact set also has a basis of compact saturated sets.

Given a topological space X we denote its topology by $\Omega(X)$ and the collection of compact saturated sets by $\mathcal{K}(X)$. We can construct a new topology on X by taking $\mathcal{K}(X)$ as the subbasis for the closed sets. This topology is called the *co-compact topology* and we write X_κ for the resulting topological space. The coarsest refinement of the original topology and the co-compact topology is called the *patch topology* and we call the corresponding space X_π . It has $\Omega(X) \cup \{X \setminus K \mid K \in \mathcal{K}(X)\}$ as a subbasis for its topology.

The real importance of these concepts will become apparent later (Section 1.3) in the special case of stably compact spaces, and the same goes for pospaces that we consider briefly in the next section. The reason to introduce them here is that we want to have the terminology available when we discuss the Lawson topology in the following section.

The specialisation order and the co-compact topology are linked by the following observation.

Proposition 1.1.2. *For a space X the specialisation preorder \sqsubseteq_κ of the co-compact topology is the dual of that for the original topology.*

Proof. Suppose $x \sqsubseteq y$, with respect to the original order, and $K \in \mathcal{K}(X)$ is such that $y \notin K$. Then $x \notin K$ because K is a saturated, that is upper, set. The sets $X \setminus K$ are a subbasis of the co-compact topology, and so we have $y \sqsubseteq_{\kappa} x$ with respect to this topology generated by it.

Conversely, take $x \not\sqsubseteq y$. Then $\uparrow x$ is compact saturated and does not contain y . Since it contains x we see that $y \not\sqsubseteq_{\kappa} x$. \square

For stably compact spaces we will see that the co-compact topology is indeed the dual topology in a certain sense to be made precise in Section 1.3.2.

Pospaces

A (*partially*) *ordered space* or simply *pospace* is a topological space X with an order relation \sqsubseteq such that its graph is a closed subset of $X \times X$, endowed with the product topology. This can be recast equivalently in a way that makes it clearer that this is a condition on the way convergence and order interact: Given converging nets $(x_i)_{i \in I}$ and $(y_i)_{i \in I}$ such that $x_i \sqsubseteq y_i$ holds point-wise the limits also satisfy $\lim x_i \sqsubseteq \lim y_i$.

Proposition 1.1.3. *A pospace is Hausdorff.*

Proof. The definition of pospace is self-dual. So if (X, \sqsubseteq) is a pospace then so is (X, \supseteq) . This implies that the diagonal, as the intersection of \sqsubseteq and \supseteq , is closed in $X \times X$, and thus that X is Hausdorff. \square

As a consequence we see that T_0 spaces can only be pospaces with respect to their specialisation order if this order is trivial. Using the patch topology they are, however, a rich source of ordered spaces.

Proposition 1.1.4. *Let X be a locally compact T_0 space. Then X_{π} is a pospace with respect to the specialisation order of the original space X .*

Proof. Given $x \not\sqsubseteq y$ we find a compact saturated neighbourhood $K \in \mathcal{N}(x)$ such that $y \notin K$. Hence $\text{int}(K) \times (X \setminus K)$ is a patch open neighbourhood of $\langle x, y \rangle$. Suppose $\text{int}(K) \ni x' \sqsubseteq y'$, then y' is also in the saturated set $K = \uparrow K$. This shows that the graph of \sqsubseteq does not meet the neighbourhood $\text{int}(K) \times (X \setminus K)$. This holds for all $\langle x, y \rangle \in (X \times X) \setminus \sqsubseteq$, and thus the order is closed. \square

1.1.2 Domains

As the conventions in domain theory differ slightly from author to author we begin by repeating the basic definitions and results from [AJ94] to fix our notation. Then we cover some domain theoretic concepts that are not covered in the introductory chapters there.

Dcpo's

A subset D of a poset is called *directed* if it is non-empty and for any two elements in D there is an upper bound in D . As a consequence all finite subsets of D have upper bounds in D . A *directed complete partial order* or *dcpo* is a poset where every directed subset D has a supremum $\bigsqcup^\uparrow D$. Note that we do not require that a dcpo is non-empty or has a least element \perp . If it has a least element we call a dcpo *pointed*.

A function f between dcpo's is *Scott-continuous* if it preserves directed suprema, i.e. $f(\bigsqcup^\uparrow D) = \bigsqcup^\uparrow f(D)$ for directed sets D . In particular, f is monotone.

A subset C of a dcpo X is *Scott-closed* if it is a lower set, i.e. $C = \downarrow C$, and if $D \subseteq C$ is directed implies $\bigsqcup^\uparrow D \in C$. Equivalently, the *Scott topology* $\Sigma(X)$ contains the upper sets that are inaccessible by directed suprema. As the sets $\downarrow x$ are Scott-closed the specialisation order for the Scott topology is just the order of the dcpo. This topology justifies calling the above functions "continuous": A function is Scott-continuous if and only if it is continuous with respect to the Scott topologies on the dcpo's.

The specialisation order for the Scott topology is the original order as it is finer than the lower but coarser than the Alexandrov topology.

Approximation

The 'order' of approximation or *way-below relation* is derived from the order on a dcpo X as follows: $x \ll y$ if for all directed sets D , $y \sqsubseteq \bigsqcup^\uparrow D$ implies that there is a $d \in D$ such that $x \sqsubseteq d$. This order clearly satisfies the implications $x \ll y \implies x \sqsubseteq y$ and $x \sqsubseteq x' \ll y' \sqsubseteq y \implies x \ll y$. In particular, \ll is transitive.

A dcpo is *continuous* if every element is the directed supremum of elements that approximate it. This turns out to be equivalent to the set $\downarrow x := \{y \mid y \ll x\}$ being directed and $x = \bigsqcup^\uparrow \downarrow x$, for all x . Following [AJ94] we reserve the term *domain* for dcpo's that are at least continuous.

An element x is *compact* if $x \ll x$. We denote the set of compact elements by $K(X)$. If in a domain every element is the directed supremum of compact elements below it we call it *algebraic*. In an algebraic domain we have $x \ll y$ if and only if there is a compact element k such that $x \sqsubseteq k \sqsubseteq y$.

We now come to what [AJ94] calls the "single most important feature of the order of approximation" for continuous domains. As we have already seen \ll is always a transitive relation. For a continuous domain it is also *interpolative*, i.e. for $x \ll x'$ we can find an interpolating element y satisfying $x \ll y \ll x'$. This *interpolation property* can be stated more generally as follows (see [AJ94, Lemma 2.2.15]): Given $M \ll x$, where M is a finite set and the relation is required to hold for every element of M , we can find an interpolating y such that $M \ll y \ll x$.

As a consequence of interpolation the sets of the form $\uparrow x$ form a basis of the Scott topology. Moreover, for any element x the sets $\uparrow y$, where $y \ll x$, are compact neighbourhoods of x , and they are a basis of the neighbourhood filter $\mathcal{N}(x)$. This shows in particular that a continuous domain is locally compact. We will make use of the following consequence for the neighbourhood filters of compact saturated sets later:

Lemma 1.1.5. *Let O be an open and K a compact saturated subset of a domain X with $K \subseteq O$. Then there is a finite set M such that $K \subseteq \uparrow M \subseteq \uparrow O \subseteq O$.*

Proof. For every $x \in K$ we can find an $x' \in O$ such that $x' \ll x$ because of the continuity of X . The union of the sets $\uparrow x'$ covers K and because of the compactness of K finitely many of the x' suffice; we can take them as the finite set M . \square

A *basis* of a domain is a subset B such that all elements are the directed supremum of elements from B that approximate it. Again this is equivalent to $\downarrow x \cap B$ being directed and yielding x as its supremum. A dcpo is algebraic if and only if its compact elements form a basis.

Note that if we are given a basis B we get a more economical basis for the Scott topology; it is given by $\{\uparrow x \mid x \in B\}$. If we specialise this to the algebraic case we can use the sets $\uparrow k = \uparrow k$, for k compact, as basic open sets.

Lawson Topology

As we have already seen the Scott topology on a continuous domain is locally compact. By Proposition 1.1.4 this implies that the patch topology for the Scott topology is a pospace.

There is a different description of this patch topology. For any dcpo X the *Lawson topology* is given by the subbasis of Scott-open sets and sets of the form $X \setminus \uparrow x$. It is the refinement of the Scott topology and the *upper topology*; the latter has the sets $\uparrow x$ as a subbasic closed sets.

In general the upper topology is coarser than the co-compact topology for the Scott topology. For continuous domains, however, they agree.

Proposition 1.1.6. *For a continuous domain X the Lawson topology is the patch topology of the Scott topology.*

Proof. As the co-compact topology is finer than the upper topology the patch topology is at least as fine as the Lawson topology.

Conversely, let us suppose that U is Scott-open, K compact saturated and $x \in U \setminus K$. By the previous lemma there is a finite set M such that $K \subseteq \uparrow M \subseteq X \setminus \downarrow x$. This implies that the point x lies in the Lawson-open set $U \setminus \uparrow M$ which, in turn, is contained in $U \setminus K$. \square

Special Classes of Domains

Unless stated otherwise the proofs for the statements in this section can be found in [AJ94, Section 4].

We call a domain *Scott domain* if it has a least element \perp and any two bounded elements have a supremum. Note that some authors require Scott domains to be also ω -algebraic; we use the term in the more liberal fashion.

Remark. *Using induction and directed suprema we see that this implies that in such a domain every bounded subset has a supremum. Thus they are sometimes also known as bounded complete domains. An equivalent condition is that every non-empty subset has an infimum. Another characterisation is that they are precisely the Scott-closed subsets of continuous lattices, i.e. complete lattices that are also continuous domains.*

The ambient category of dcpo's is already cartesian closed. Products and exponentials are created by the forgetful functor to Set, the order in both cases is given point-wise. Unfortunately, neither the category of algebraic nor that of continuous domains is cartesian closed and so one has to restrict to the domains further (see [AJ94, Section 4]). Scott domains form a cartesian closed subcategory, and, moreover, they satisfy a very strong extension property [Esc98].

Sometimes, for example for several power domain constructions, one is forced to consider more general domains. We do the algebraic case first.

Proposition 1.1.7. *For an algebraic domain X with least element the following two conditions are equivalent:*

1. *$K(X)$ is MUB-complete, i.e. for every upper bound x of a finite subset M there is a minimal upper bound of M below x ; and every finite subset M has a finite MUB-closure which is the smallest set $N \supseteq M$ such that N contains all minimal upper bounds of all its subsets.*
2. *There is a directed family of continuous idempotents with finite image on X whose supremum is the identity.*

If an algebraic domain satisfies the equivalent conditions of the proposition we call it *bifinite*. They were introduced by Gordon Plotkin [Plo76].

Remark. *There is another characterisation of bifinite domains as bilimits (see [AJ94, Section 3]) of finite posets which explains the name "bifinite" and also the common acronym SFP domains which stands for "sequence of finite posets".*

As an immediate corollary to the proposition every algebraic Scott domain is bifinite. Moreover, bifinite domains form a cartesian closed category, in fact it is a maximal cartesian closed sub-category of pointed algebraic domains (see [Smy83a] and [Jun89]).

Let \mathcal{C} be a cartesian closed category. If \mathcal{C}' is a full subcategory of \mathcal{C} that is itself cartesian closed, then so is the full sub-category of retracts of objects from \mathcal{C}' (see [Sco80]). As the continuous domains are precisely the retracts of the algebraic domains, we get a fairly large category of continuous domains by applying this procedure to bifinite domains. The resulting domains are called *RSFP* for “retracts of SFP”.

Proposition 1.1.8. *A domain X is an RSFP domain if and only if there is a directed family of continuous endo-functions on X with finite image such that their supremum is the identity.*

Proof. [Jun89, Theorem 4.1] □

Every continuous Scott domain is an RSFP domain. The idea of the proof is as follows: Given any finite subset M of a Scott domain there are only finitely many suprema of subsets of M . Hence, we can assume, without loss of generality, that M is closed under bounded suprema. This allows us to define a projection as required in the proposition by mapping an element to the supremum, and hence the largest, of those members of M that approximate it. It is fairly straightforward to verify that this function is continuous. Furthermore, every infinite set is the directed union of its finite subsets, and hence the identity is the directed supremum of such functions.

It is not known whether the category of RSFP domains is a maximal cartesian closed sub-category of pointed continuous domains. It is contained in the category of FS domains—to be defined shortly—which is maximal, but it is an open problem whether this containment is proper.

We say a function $f: X \rightarrow X$ is *finitely separated* from the identity if there is a finite subset $M \subseteq X$ such that for all $x \in X$ there is an $m \in M$ satisfying $f(x) \sqsubseteq m \sqsubseteq x$. An *FS domain* is a pointed domain for which there is a directed family of endo-function that are finitely separated from the identity and whose supremum is the identity. This condition is clearly weaker than that of Proposition 1.1.8, hence RSFP is a subcategory of FS.

FS domains were first introduced in [Jun90]. For more information see there or [AJ94, Section 4]. They can also be described in purely topological terms [JS98]. This paper also contains more detailed proofs of some properties of FS domains than the two other sources.

A function that is finitely separated from the identity maps any element to one that approximates it:

Lemma 1.1.9. *Let X be a dcpo and $f: X \rightarrow X$ a continuous function that is finitely separated from the identity. Then for all $x \in X$ we have $f(x) \ll x$.*

Proof. Let $M \subseteq X$ be a finite set that shows that f is finitely separated from id and $x \in X$ an arbitrary element. Now, suppose D is directed and $x \sqsubseteq \bigsqcup^\uparrow D$ which clearly implies $f(x) \sqsubseteq f(\bigsqcup^\uparrow D) = \bigsqcup^\uparrow f[D]$. For every $y \in D$ there is an $m \in M$ such that $f(y) \sqsubseteq m \sqsubseteq y$. As M is finite and D directed

there must be at least one such m that satisfies $f[D] \sqsubseteq m$. This implies $f(x) \sqsubseteq \bigsqcup^\uparrow f[D] \sqsubseteq m$ and, moreover, there is an element $y \in D$ such that $f(y) \sqsubseteq m \sqsubseteq y$. This proves $f(x) \ll x$. \square

As an immediate corollary we get:

Corollary 1.1.10. *An FS domain is continuous.*

We will come back to these special classes of domains in the next section when we study their Stone duality. Later we will see that their Scott topologies give rise to stably compact spaces.

1.2 Stone Duality

We proceed with a brief introduction to Stone duality. This provides the link between topology and logic in the form of locales. They can be thought of as Lindenbaum algebras for a logic of observable properties that has (normal) conjunction and infinite disjunction. For details see [AJ94, Sections 7.1 and 7.2], [Joh82, Chapters 2 and 3] or [Vic89, Chapters 3–5].

1.2.1 The Adjunction $\Omega \dashv \text{pt}$

Frames and Locales

The open sets of a topological space are closed under arbitrary suprema and hence form a complete lattice. As only finite infima can be calculated by intersections it is natural to look at these finite intersections. It turns out that they distribute over arbitrary joins motivating the following definition.

Definition 1.2.1. A complete lattice is called a *frame* if it satisfies the following *frame distributivity*:

$$x \wedge \bigvee_i y_i = \bigvee_i (x \wedge y_i).$$

A *frame homomorphism* is a function between frames that preserves all suprema and finite infima. We denote the category of frames and frame homomorphisms by Frm . Its opposite is the category of *locales*, Loc .

Observe that the definition of a locale mirrors the logic of observable properties: The equations used in the definition are made up from arbitrary disjunctions but only finite conjunctions.

Remark. *The category of frames has a number of nice properties. It is, for example, algebraic over Set which means that free frames exist. More importantly, one can use generators and relations to construct frames with certain properties. For details see [Joh82, Chapter II].*

A continuous function f between topological spaces X and Y gives rise to a frame homomorphism $f^{-1}[\cdot]: \Omega(Y) \rightarrow \Omega(X)$ which takes an open set to its preimage. Hence we get a functor Ω from Top , the category of topological spaces, to Loc .

It is worth mentioning that open sets in a topological space X are in bijection to continuous functions from X to $\mathbf{2}$, the poset of two elements $0 \sqsubseteq 1$ equipped with the Scott topology. The bijection takes an open set U to its characteristic function χ_U and a map $\chi: X \rightarrow \mathbf{2}$ to the open set $\chi^{-1}[\{1\}]$. This correspondence is an order isomorphism since $U \subseteq V$ is equivalent to $\chi_U \sqsubseteq \chi_V$.

Taking the preimage under f can then be described as composing the characteristic function of an open set with f , i.e. $\chi_{f^{-1}[U]} = \chi_U \circ f$. This implies that Ω followed by the forgetful functor to \mathbf{Set} is isomorphic to the Hom-functor $\mathbf{Top}(-, \mathbf{2})$.

Reconstructing a Space

The natural question now is whether we can retrieve a space X from its locale of open subsets $\Omega(X)$. For a Hausdorff space all sets of the form $X \setminus \{x\}$ are open and co-atoms in $\Omega(X)$. These co-atoms turn out to correspond exactly to points of X .

For T_0 spaces, i.e. the ones we are particularly interested in, this is no longer true. In general we observe that $\overline{\{x\}}$ is an \cup -irreducible closed set which implies that $X \setminus \overline{\{x\}}$ is \wedge -irreducible in $\Omega(X)$. In a distributive lattice, in particular in a locale, the \wedge -irreducible and the \wedge -prime elements coincide. Hence, our candidates for points are the \wedge -prime elements.

Another approach is to observe that the open neighbourhood filter of a point x is a completely prime filter in $\Omega(X)$ which we denote by $\mathcal{N}^\circ(x)$. This suggests that we might want to take them as points. Fortunately, in a complete lattice the \wedge -prime elements and the completely prime filters are in one-to-one correspondence: If x is prime then the complement of $\downarrow x$ is a completely prime filter, and conversely if F is such a filter the supremum of its complement is still in the complement and prime.

There is a third way of describing points. The characteristic function of a completely prime filter is a frame homomorphism¹ to $\mathbf{2}$, now considered as a two element locale, and again the correspondence is one-to-one.

The last description allows for the simplest definition of the contravariant functor \mathbf{pt} from \mathbf{CLat} , the category of complete lattices with frame homomorphisms, to \mathbf{Top} : Its image under the forgetful functor can be taken to be $\mathbf{CLat}(-, \mathbf{2})$; an open subset of the space $\mathbf{pt}(L)$ is a set of ‘points’ which take a specific x to $1 \in \mathbf{2}$. Describing opens by their characteristic function they are of the form $\lambda f.f(x)$.

The two functors \mathbf{pt} and Ω form a dual adjunction. If we restrict \mathbf{pt} to locales we can consider it as an adjunction $\Omega \dashv \mathbf{pt}$ between \mathbf{Top} and \mathbf{Loc} . Since \mathbf{Loc} is simply the dual category of \mathbf{Frm} both functors are now covariant.

From now on we are going to use the term ‘point’ to refer to a completely prime filter in a locale or more generally in a complete lattice. We might also use the term to refer to the corresponding characteristic function as long as it is obvious from the context what is meant. Sometimes, if we want to make clear that we are talking about the elements of a locale and not its ‘points’ we refer to the elements as the ‘opens’ of the locale.

¹We use this as a shorthand for “preserves arbitrary suprema and finite infima” even if the lattices involved are just complete lattices and not frames.

Given a locale L one can directly read off the specialisation order of the topological space $\mathbf{pt}(L)$ it describes; it is just subset inclusion. The extensional order on functions between such spaces, introduced in Section 1.1.1, is also manifest on the localic side.

Proposition 1.2.2. *For two functions f and g between topological spaces we have $f \sqsubseteq g$ if and only if for all open sets U the preimages satisfy $f^{-1}[U] \subseteq g^{-1}[U]$, or in other words if $\Omega(f) \sqsubseteq \Omega(g)$.*

Conversely, if $f \sqsubseteq g$ holds point-wise for frame homomorphisms f and g then $\mathbf{pt}(f) \sqsubseteq \mathbf{pt}(g)$.

1.2.2 Sobriety and Spatiality

We still have not answered the question which spaces can be reconstructed from their lattice of opens. It is obvious that a space has at least to satisfy the T_0 separation axiom. For the dual problem of reconstructing a lattice from its ‘points’ it is clear that the lattice has to be at least a locale. So, the question is tantamount to restricting and co-restricting the above adjunction to an equivalence. That is to say we have to characterise those spaces and locales whose units and co-units are isomorphisms.

Before we can do this we need some information on how the units $\eta_X : X \rightarrow \mathbf{pt}(\Omega(X))$ and co-units $\epsilon_L : L \rightarrow \Omega(\mathbf{pt}(L))$ work. As we have already observed there is a perfect symmetry of the two functors \mathbf{pt} and Ω which allows us to write them both as Hom-functors $(-, \mathbf{2})$ in the respective categories. So, it may not come as a surprise that both η_X and ϵ_L can be written as the evaluation function $\lambda x. \lambda f. f(x)$, disregarding the types of the arguments.

For unit this is just a fancy way of saying that η_X takes a point to its neighbourhood filter. Thinking of ‘points’ as completely prime filters the co-unit takes the form $\epsilon_L(x) := \{P \in \mathbf{pt}(L) \mid x \in P\}$. When we want to stress that this is an open set in $\mathbf{pt}(L)$ we sometimes refer to it as O_x . As ϵ_L is a frame homomorphism we infer $O_{\bigvee x_i} = \bigcup O_{x_i}$ and $O_{x \wedge y} = O_x \cap O_y$.

Now, we can state the relevant facts concerning the reconstruction of spaces:

Proposition 1.2.3. *For a topological space X the following are equivalent:*

1. *The unit η_X is a homeomorphism.*
2. *The function η_X is bijective.*
3. *Every \cup -irreducible closed set is the closure of a unique point.*
4. *Every completely prime filter in $\Omega(X)$ is the open neighbourhood filter of a unique point.*

A space satisfying these equivalent conditions is called *sober*. Note that if we restrict ourselves to T_0 spaces we only have to ask for η_X to be surjective, or equivalently we can ignore uniqueness in conditions (3) and (4).

On the complete lattice side we get:

Proposition 1.2.4. *For a complete lattice L the following conditions are equivalent:*

1. *The co-unit ϵ_L is an order-isomorphism.*
2. *ϵ_L is injective.*
3. *ϵ_L is order-reflecting.*
4. *The elements of L are separated by completely prime filters.*
5. *For elements $x \not\leq y$ there exists a completely prime filter containing x but not y .*
6. *L is \wedge -generated by \wedge -prime elements, i.e. every element is the meet of the meet-prime elements above it.*

A complete lattice is called *spatial* if it satisfies the conditions of the above proposition.

The functors \mathbf{pt} and Ω give rise to closure operators in the following sense:

Proposition 1.2.5. *For a complete lattice L the space $\mathbf{pt}(L)$ is sober, and for a topological space X the locale $\Omega(X)$ is spatial.*

The functors \mathbf{pt} and Ω restrict and co-restrict to an equivalence of the full subcategories of sober spaces and spatial locales.

The composition $\mathbf{pt} \circ \Omega$ of these two functors is known as *sobrification*, $\Omega \circ \mathbf{pt}$ as *spatialisation*.

Proposition 1.2.6. *Arbitrary products and coproducts of sober spaces are sober.*

Proof. For products this is an immediate consequence of the fact that the functor \mathbf{pt} as a right adjoint preserves all limits.

For coproducts let $A \subseteq \coprod_i X_i$ be an irreducible closed subset of the disjoint union of the sober spaces X_i . We claim that A must be a subset of one of the X_i . This is true since $A \cap X_j \neq \emptyset$ and $A \cap X'_j \neq \emptyset$ implies that we can write it as a non-trivial union

$$A = \left(\left(\coprod_i X_i \right) \setminus X_j \cap A \right) \cup \left(\left(\coprod_i X_i \right) \setminus X'_j \cap A \right)$$

of closed sets, contradicting its irreducibility. The claim now follows immediately from the sobriety of the corresponding space $X_i \supseteq A$. \square

Sobriety and Domains

For the programme of domain theory in logical form we will concentrate on spaces which are domains with their Scott topology. With respect to our current viewpoint these spaces are ‘well-behaved’:

Proposition 1.2.7. *If D is a continuous domain then $\Sigma(D)$, the Scott topology on D , is sober.*

Continuity cannot be dropped from the previous proposition [Joh81]. Later on we will characterise the topologies that arise as Scott topologies for continuous domains.

There is also a connection between continuity and spatiality: For a distributive complete lattice continuity implies frame distributivity because intersection on a continuous domain is Scott-continuous. Hence, such lattices are locales. Moreover they turn out to be spatial.

Proposition 1.2.8. *A distributive, continuous lattice is a spatial locale.*

For a thorough discussion of continuous lattices see [GHK⁺80]. Most of the results about the hierarchy of Stone dualities, given later in this chapter, can also be found there.

As mentioned above the specialisation order on $\text{pt}(L)$ is just subset inclusion. As the directed union of ‘points’, i.e. completely prime filters, is again a ‘point’ this specialisation order is always a dcpo. We get even more:

Proposition 1.2.9. *The specialisation of a sober space is a dcpo, and the topology is coarser than the Scott topology.*

We can extend this to morphisms as well. The standard proof that for dcpo’s Scott-continuity and topological continuity with respect to the Scott topologies coincide actually shows slightly more. If $f: X \rightarrow Y$ is a continuous function from a dcpo X with the Scott topology to any space Y then f is Scott-continuous, i.e. f preserves directed suprema. In particular, suprema of images of directed sets in X exist in Y . As a consequence we see that continuous maps between sober spaces are Scott-continuous.

Remark. *Categorically, this makes the specialisation order a functor from sober spaces to dcpo’s.*

Considered as a functor to posets the left adjoint of the specialisation order is given by the sobrification of the Alexandrov topology. The resulting space is just the ideal completion of the poset with the Scott topology.

We can also take directed suprema of functions between sober spaces.

Lemma 1.2.10. *If X and Y are sober spaces and $f_i: X \rightarrow Y$ is a directed family of continuous maps then the point-wise supremum $\bigsqcup^\uparrow f_i$ is also continuous.*

This corresponds to taking the point-wise supremum of the locale morphisms:
 $\Omega(\bigsqcup^\uparrow f_i) = \bigsqcup^\uparrow \Omega(f_i)$.

Proof. We first claim that for a Scott-open set U we have $(\bigsqcup^\uparrow f_i)^{-1}[U] = \bigcup^\uparrow f_i^{-1}[U]$. The latter set is clearly contained in the former. For the other subset inclusion note that $x \in (\bigsqcup^\uparrow f_i)^{-1}[U]$ implies $\bigsqcup^\uparrow f_i(x) \in U$ and hence $f_i(x) \in U$ for an index i .

From this observation we instantly get that $\bigsqcup^\uparrow f_i$ is continuous as for a sober space every open set is Scott-open. The second claim is also an immediate consequence. \square

The proof of the first assertion of the lemma did not even use sobriety. It works whenever Y carries a topology that is coarser than the Scott topology.

Compact Saturated Sets

Given a compact subset K of a topological space X its open neighbourhood filter $\mathcal{N}^\circ(K) = \{O \in \Omega(X) \mid K \subseteq O\}$ is clearly a Scott-open filter in the locale $\Omega(X)$. (Without loss of generality we can assume K to be saturated as any set has the same open neighbourhood filter as its saturation.)

This suggests to think of Scott-open filters in any locale L as compact saturated ‘subsets’. To justify this we have to show that every such filter does indeed correspond to a compact saturated subset in $\mathbf{pt}(L)$. This is the statement of the (localic) Hofmann-Mislove Theorem [HM81]. We discuss the proof in detail so we can compare it to the logical version later. It is essentially a localic version of the proof given in [KP94]. The following is the key lemma.

Lemma 1.2.11. *If F is a Scott-open filter in a distributive complete lattice and $x \in L \setminus F$ then there is a completely prime filter $P \supseteq F$ such that $x \notin P$.*

Proof. The set $L \setminus F$ is Scott-closed and so in particular inductively ordered. By Zorn’s Lemma we find a maximal element $x' \geq x$ that does not lie in F . Any maximal element not in a filter is \wedge -irreducible, and in a distributive lattice \wedge -irreducible and \wedge -prime elements coincide. This in turn implies that $P := L \setminus \downarrow x'$ is again a filter. This filter is clearly completely prime as the complement of the principal ideal $\downarrow x'$. The conditions $x \notin P$ and $F \subseteq P$ follow immediately from $x \leq x' \notin F = \uparrow F$. \square

The lemma can be rephrased to say that completely prime filters, i.e. the ‘points’, separate elements from Scott-open filters, or that every Scott-open filter is the intersection of the completely prime filters containing it.

We can now prove the above claim that all Scott-open filters in a locale correspond to compact saturated sets.

Theorem 1.2.12. *Given a distributive complete lattice L the Scott-open filters on L are in one-to-one correspondence to the compact saturated subsets of $\mathbf{pt}(L)$.*

The isomorphism takes a filter $F \subseteq L$ to the set of all ‘points’ containing it and a compact saturated set $K \subseteq \mathbf{pt}(L)$ to $\bigcap K = \{x \in L \mid K \subseteq O_x\}$. These maps are order-isomorphisms with respect to inclusion on filters and reverse inclusion on compact saturated sets.

Proof. Suppose K is a compact saturated set. If we consider the equivalences

$$K \subseteq O_x \iff (\forall P \in K) x \in P \iff x \in \bigcap K$$

we see that the sets $\{x \mid K \subseteq O_x\}$ and $\bigcap K$ are indeed equal, and as the first is an intersection of filters it is again a filter. That it is Scott-open follows from the compactness of K and the fact that $O_{(\cdot)} = \epsilon_L$ is a frame homomorphism and thus commutes with arbitrary suprema.

Now take a Scott-open filter $F \subseteq L$. We only have to show that the set $K := \{P \in \mathbf{pt}(L) \mid F \subseteq P\}$ is compact, because it is obviously an upper set and thus saturated. Suppose it is covered by a directed union of open sets $\bigcup^\uparrow O_{x_i} = O_{\bigvee^\uparrow x_i}$ which is equivalent to $(\forall P \supseteq F) \bigvee^\uparrow x_i \in P$, where P is a ‘point’. By the previous lemma we get $\bigvee^\uparrow x_i \in F$ and as F is Scott-open there must be an index i for which $x_i \in F$. This implies $K \subseteq O_{x_i}$, and thus K is compact.

It remains to show that the above mappings are mutually inverse. From the lemma we already know that $\bigcap \{P \in \mathbf{pt}(L) \mid F \subseteq P\} = F$ holds. For the converse we get

$$\begin{aligned} \{P \in \mathbf{pt}(L) \mid \bigcap K \subseteq P\} &= \{P \mid (\forall x \in \bigcap K) x \in P\} \\ &= \bigcap \{O_x \mid x \in \bigcap K\} = \bigcap \{O_x \mid K \subseteq O_x\} \\ &= \uparrow K = K. \end{aligned}$$

□

Starting with a sober space, the more common, topological version of the theorem is an immediate consequence.

Corollary 1.2.13. *The compact saturated subsets of a sober space X are order isomorphic the Scott-open filters of $\Omega(X)$.*

The corresponding maps take a Scott-open filter to its intersection and a compact saturated set to its open neighbourhood filter.

Remark. *We have already observed that a subset of a space and its saturation have the same neighbourhood filter, and we can clearly reconstruct a saturated set from this filter by taking the intersection. This suggests a one-to-one correspondence between saturated sets and filters in the locale.*

Unfortunately, this does not quite work; in general, several filters represent the same saturated set. As an example consider the filter

$$\mathcal{F} := \{ \mathbb{R} \setminus \{1, 2, \dots, n\} \mid n \in \mathbb{N} \} \subseteq \Omega(\mathbb{R}).$$

Its intersection is $\bigcap \mathcal{F} = \mathbb{R} \setminus \mathbb{N}$, an open set that does not lie in \mathcal{F} .

The filters that do arise as neighbourhood filters of saturated sets are precisely those that are intersections of completely prime filters.

These observations can be used to give some meaning to infinite meets in locales, although they are not part of the language. The infimum of any subset in a complete lattice is the same as that of the filter generated by the set. Hence \bigwedge is essentially an operation on filters. In a locale the result yields the interior of the saturated set that such a filter represents.

We summarise the correspondence between topological and localic concepts that we have been discussing up to here in the following table. In the rest of this thesis we are going to make heavy use of these equivalences, sometimes without explicitly mentioning it.

Space	Locale
point	completely prime filter
open set	element
compact saturated set (saturated set)	Scott-open filter (filter)

In many cases the Hofmann-Mislove Theorem is used in the form of the following corollary. It is slightly stronger than just stating that for a sober space X the poset $(\mathcal{K}(X), \supseteq)$ is a dcpo.

Corollary 1.2.14. *In a sober space X a filtered intersection of compact saturated sets is compact saturated.*

Moreover, if such an intersection lies in an open set O than one of the compact saturated sets is already a subset of O . In particular, a filtered intersection of non-empty compact saturated sets is non-empty.

Proof. We can translate the directed intersection into a directed union of Scott-open filters. The result is then again such a filter. This implies that the infimum of the directed family of compact saturated sets exist in $(\mathcal{K}(X), \subseteq)$ and that it is given by the intersection of the union of the neighbourhood filters of the original compact saturated sets.

Now, it is generally true in a complete lattice that the meet of a union of subsets is the meet of the meets of the individual sets—this is ‘general associativity’. Applying this to the power set $\mathfrak{P}(X)$ we see that the above infimum is just the filtered intersection of the compact saturated sets.

Thinking on the side of Scott-open filters it is obvious that an open set is only in the directed union if it was already in one the original filters. \square

As another application of the Hofmann-Mislove Theorem we prove the localic equivalent of the well-known theorem that the continuous image of a compact set is compact. But first we need a lemma.

Lemma 1.2.15. *If P is a completely prime filter and $\bigcap Q_i \subseteq P$, where the Q_i are upper sets, then there is an index i such that $Q_i \subseteq P$.*

Proof. Suppose otherwise that for all i there is an element $x_i \in Q_i \setminus P$. Then $\bigvee x_i$ lies in $\bigcap Q_i$ as these are upper sets, but on the other hand the supremum cannot be a member of the completely prime filter P . Hence we have $\bigvee x_i \in \bigcap Q_i \setminus P$, a contradiction. \square

Topologically, the following proposition is almost trivial. The reason for a localic form of it is that we need it later for the proof of Lemma 1.3.16.

Proposition 1.2.16. *If $f: L \rightarrow M$ is a frame homomorphism and $F \subseteq M$ a Scott-open filter then so is its preimage $f^{-1}[F]$.*

Topologically this can be seen as map $f_{\mathcal{K}}: \mathcal{K}(\mathbf{pt}(L)) \rightarrow \mathcal{K}(\mathbf{pt}(M))$ that takes a compact saturated set K to $\uparrow f[K]$.

Proof. The first assertion is trivial because frame homomorphisms preserve finite infima and arbitrary, hence in particular directed, suprema.

Now, take a ‘point’ P in the compact saturated set K corresponding to a Scott-open filter F —in other words, let $F \subseteq P$. Then we get $f^{-1}[F] \subseteq f^{-1}[P] = (\mathbf{pt}f)(P)$ and we see that $f[K]$ is a subset of the compact saturated set corresponding to $f^{-1}[F]$. This implies that the saturation $\uparrow f[K]$ is also contained in this set.

Conversely, if we have $f^{-1}[F] \subseteq P$, then we can re-write the first filter using Lemma 1.2.11 as

$$f^{-1}[F] = f^{-1}\left[\bigcap\{Q \in \mathbf{pt}(M) \mid F \subseteq Q\}\right] = \bigcap\{f^{-1}[Q] \mid F \subseteq Q\}.$$

So, by the previous lemma we get a ‘point’ $Q \supseteq F$ such that $f^{-1}[Q] \subseteq P$ which shows the other subset inclusion. \square

1.2.3 Hierarchy of Stone Dualities I

We now investigate how the duality between sober spaces and spatial locales can be restricted and co-restricted. For a road map to the hierarchy of Stone dualities we are studying in this section see Figure 1.1 on page 35.

Local Compactness

From Proposition 1.2.8 we know that a continuous locale is spatial. We claim that these locales are exactly those which are isomorphic to the lattice of open subsets of a locally compact sober space. To see this we need more information about the order of approximation for the opens and compact saturated subsets of a space.

Lemma 1.2.17. *Let X be a locally compact space. Then $O \ll O'$ holds in the lattice $(\Omega(X), \subseteq)$ if and only if there is a compact (saturated) set K such that $O \subseteq K \subseteq O'$.*

In $(\mathcal{K}(X), \supseteq)$ the condition $K \ll K'$ implies that there is an open set O satisfying $K \supseteq O \supseteq K'$. The converse is true if, in addition, X is sober.

Proof. In the situation $O \subseteq K \subseteq O'$, where K is compact, we clearly have $O \ll O'$. For the reverse implication observe that local compactness implies that O' is the union of compact neighbourhoods of its points, and that finite unions of compact sets are compact.

Now, suppose $K \ll K'$. As a saturated set, K' is the intersection of any filter basis of its neighbourhood filter. By local compactness we can assume this to be made up from compact saturated sets, and so we find $L \in \mathcal{K}(X)$ such that $K' \subseteq \text{int}(L) \subseteq L \subseteq K$.

The implication $K \supseteq O \supseteq K' \implies K \ll K'$ is a direct consequence of Corollary 1.2.14. \square

From the first half of the lemma it follows that local compactness of a space entails the continuity of the corresponding locale.

Proposition 1.2.18. *The equivalence of Proposition 1.2.5 between sober spaces and spatial locales restricts and co-restricts to locally compact spaces and continuous locales.*

Proof. For the remaining direction consider $u \in P \in \text{pt}(L)$, where L is a continuous locale. We can think of $u \in P$ to say that the ‘point’ P lies in the ‘open set’ u . P is completely prime and thus Scott-open, and as L is continuous we find $u' \in P$ that approximates u . Using the interpolation property repeatedly we construct a sequence $u' \ll \dots \ll u_2 \ll u_1 \ll u$. It is now easy to check that $K := \bigcup_{n \in \mathbb{N}} \uparrow u_n$ is a Scott-open filter and that we have $u \in K \subseteq \uparrow u' \subseteq P$. Using the Hofmann-Mislove Theorem in the form of Lemma 1.2.11 we see that thinking in terms of the space $\text{pt}(L)$ this means that the compact saturated set corresponding to K is contained in O_u and that the interior of K is larger than O'_u which in turn contains the ‘point’ P . Consequently, $\text{pt}(L)$ is locally compact. \square

Remark. *This result can actually be slightly sharpened: Because of $\Omega(X) \cong \Omega(\text{pt}(\Omega(X)))$, the open set lattice $\Omega(X)$ is continuous if and only if $\text{pt}(\Omega(X))$, the sobrification of X , is locally compact.*

In the following compact saturated and open sets will play equally important roles; this link will feature more prominently in the next section. Hence, we state the analogue of the proposition for compact saturated sets.

Proposition 1.2.19. *For a locally compact sober space X , $(\mathcal{K}(X), \supseteq)$ is a continuous dcpo.*

The sets $\square O := \{K \in \mathcal{K}(X) \mid K \subseteq O\}$, for O open, form a basis of the Scott topology.

Proof. We have already seen in Corollary 1.2.14 that $\mathcal{K}(X)$ is a dcpo under reverse inclusion. Now from local compactness we can infer that every compact set has a neighbourhood filter basis of compact saturated sets. By Lemma 1.2.17 this shows that this dcpo is continuous.

A set $\square O$ is Scott-open by the second observation of Corollary 1.2.14, and interpolation together with Lemma 1.2.17 yields that such sets form a basis of the topology. \square

In general, we can say nothing about unions of compact sets. But if the set that we take the union over is itself compact the situation is much better. The following lemma is taken from [Sch93, Section 7.3.2] and will be important later.

Lemma 1.2.20. *Let X be a locally compact sober space. If \mathcal{C} is a compact subset of $\mathcal{K}(X)$ then $\bigcup \mathcal{C}$ is compact saturated.*

Proof. The set $\bigcup \mathcal{C}$ is an upper set as a union of upper sets. To see that it is compact suppose \mathcal{O} is a directed open cover, i.e. $\bigcup^\uparrow \mathcal{O} \supseteq \bigcup \mathcal{C}$. For every $K \in \mathcal{C}$ we can find an $O \in \mathcal{O}$, by compactness, such that $K \subseteq O$, or in other words $K \in \square O$. The set $\{\square O \mid O \in \mathcal{O}\}$ is clearly a directed family of open subsets of $\mathcal{K}(X)$, and by the previous observation it covers \mathcal{C} . Hence, we can use the compactness of \mathcal{C} to find an $O \in \mathcal{O}$ such that $\mathcal{C} \subseteq \square O$ and thus $\bigcup \mathcal{C} \subseteq O$. \square

Restricting locales further it is quite natural to ask what happens if we require them to be not just continuous but algebraic. The compact elements in $\Omega(X)$ are exactly the compact-open subsets of X , which we denote by $K\Omega(X)$. So, algebraicity of $\Omega(X)$ is tantamount to saying that the space X has a basis of compact-open subsets.

Proposition 1.2.21. *The functors Ω and pt restrict to an equivalence between spaces with a basis of compact-open sets and algebraic locales.*

Domains

Up to now we have discussed general topological spaces. In the following we want to concentrate on those that are continuous domains with their Scott topology. We let Dom denote the category of continuous domains with Scott-continuous functions and use Alg for the full subcategory of algebraic domains. Their Stone duals have been identified by Lawson [Law79] and Hoffmann [Hof81] to be the completely distributive lattices.

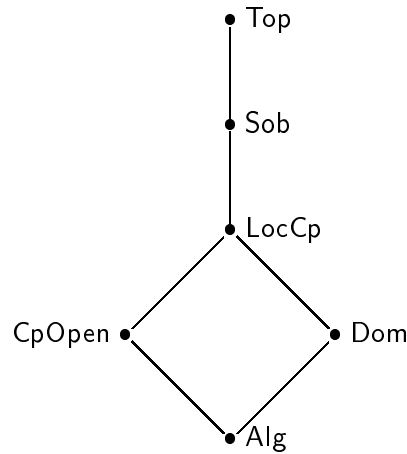


Figure 1.1: Hierarchy of categories of topological spaces I

Definition 1.2.22. A complete lattice is called *completely distributive* if the following infinite distributivity law holds:

$$\bigwedge_{i \in I} \bigvee A_i = \bigvee_{f \in \prod A_i} \bigwedge_{i \in I} f(i)$$

As with the finite version of this law, complete distributivity is a self-dual concept.

Proposition 1.2.23. *The Stone dual of Dom is the category of completely distributive lattices.*

To understand the intersection of Dom and CpOpen we have to study compact-open sets in continuous domains. We claim that they are exactly those of the form $\uparrow M$, for a finite set M of compact elements. Such sets are clearly compact-open. To see the converse take a compact-open set K , and note that being an open set it can be written as $K = \bigcup_{x \in K} \uparrow x$. By compactness we find a finite subset $M \subseteq K$ such that $K = \uparrow M$, and after eliminating superfluous elements from M we see that it consists of compact minimal elements of K .

Proposition 1.2.24. *The category Alg is the intersection of Dom and CpOpen. Its Stone dual is the category of algebraic completely distributive lattices.*

Proof. In an algebraic domain X sets of the form $\uparrow M$, for $M \subseteq_{\text{fin}} K(X)$, are a basis of the Scott topology.

Conversely, for any point x the sets $\uparrow y$, where $y \ll x$, form a basis of its neighbourhood filter. If the compact-open sets are a basis for the topology then the considerations prior to this proposition show that x is the directed supremum of compact elements. \square

1.3 Stably Compact Spaces

Another direction to take from locally compact spaces is to study stably compact spaces. They play an important role in domain theory as the commonly used classes of domains give rise to stably compact spaces when equipped with the Scott topology; for details see [AJ94, Section 4], [Jun89] and [Jun90]. Stably compact spaces are also closely related to compact pospaces, and they have many of the ‘nice’ properties of Hausdorff spaces.

In this section we will use some more advanced tools from general topology, in particular ultrafilters and nets, which are, for example, discussed in [Bou88] and [Kel75].

1.3.1 The Definition

There are several ways to define stably compact spaces. As these spaces are at the centre of this thesis we explicitly prove these definitions equivalent.

But first we need some auxiliary definitions. The order of approximation in a complete lattice is called *multiplicative* if $x \ll y, z$ implies $x \ll y \wedge z$.

A space is *supersober* if given an ultrafilter \mathcal{U} on it the set of its limits $L(\mathcal{U}) = \bigcap \{\overline{F} \mid F \in \mathcal{U}\}$ is either empty or has a largest element x . The latter is equivalent to the condition that the set of limits is closure of that point, i.e. $L(\mathcal{U}) = \downarrow x$.

This terminology is justified by the following observation.

Proposition 1.3.1. *A supersober space is sober.*

Proof. Let X be supersober and A an irreducible closed subset. We define

$$\mathcal{B} := \{U \cap A \mid U \in \Omega(X)\} \setminus \{\emptyset\}$$

which is a filter basis because of the irreducibility of A , and hence we can refine the filter it generates to an ultrafilter \mathcal{U} . By the construction of \mathcal{B} every element of A is a limit of this ultrafilter, and as $A \in \mathcal{B} \subseteq \mathcal{U}$ and A is closed no point outside of A can be in the set of its limits $L(\mathcal{U})$. We have thus shown $A = L(\mathcal{U})$ and by supersobriety this set is the closure of a unique point completing the proof that X is sober. \square

If X is a Hausdorff space then the specialisation order is simply equality. This implies that such spaces are trivially supersober and thus sober. In the T_0 setting these concepts become more interesting. A particularly important class of supersober spaces is given by the locally compact ones which can be characterised in a number of different ways:

Theorem 1.3.2. *For a locally compact sober space X the following conditions are equivalent:*

1. X is supersober;

2. the intersection of any two compact saturated sets is again compact;
3. the order of approximation on $\Omega(X)$ is multiplicative;
4. given two Scott-open filters $F, G \subseteq \Omega(X)$ the filter generated by $F \cup G$ is again Scott-open.

Proof. We begin by showing the last three conditions equivalent.

“(2) \Rightarrow (3)”: Let U, V, W be open sets such that $U \ll V, W$. Since X is locally compact we find compact saturated sets K and L satisfying $U \subseteq K \subseteq V$ and $U \subseteq L \subseteq W$. We get $U \subseteq K \cap L \subseteq V \cap W$, and as $K \cap L$ is compact by assumption we see $U \ll V \cap W$.

“(3) \Rightarrow (4)”: Suppose that F and G are Scott-open filters in the lattice $\Omega(X)$. First, we observe that the filter generated by $F \cup G$ is just $\{U \cap V \mid U \in F, V \in G\} =: H$. This set is clearly filtered and generates the same filter. Since $\Omega(X)$ is distributive we see that $W \supseteq U \cap V$ implies $W = W \cup (U \cap V) = (W \cup U) \cap (W \cup V)$ which shows that the set is already a filter.

We now verify that H is Scott-open. For a locally compact space $\Omega(X)$ is continuous, and so for $U \in F$ and $V \in G$ we can find $U' \in F$ and $V' \in G$ such that $U' \ll U$ and $V' \ll V$. We infer $U' \cap V' \ll U, V$ and thus $U' \cap V' \ll U \cap V$ by assumption. This shows that H is a union of Scott-open sets of the form $\uparrow(U' \cap V')$, hence Scott-open.

“(4) \Rightarrow (2)”: Let K and L be compact saturated sets. Then the neighbourhood filters $\mathcal{N}^\circ(K) := \{U \in \Omega(X) \mid K \subseteq U\}$ and $\mathcal{N}^\circ(L)$ are Scott-open. By Corollary 1.2.13 the set $K \cap L$ is simply the intersection over the filter generated by $\mathcal{N}^\circ(K) \cup \mathcal{N}^\circ(L)$. By assumption this filter is Scott-open and so because of the Hofmann-Mislove Theorem $K \cap L$ is compact.

“(1) \Rightarrow (2)”: Let K and L be compact saturated. We have to show that $K \cap L$ is compact. To this end let \mathcal{u} be an ultrafilter on $K \cap L$. Considered as a filter basis it induces ultrafilters $\mathcal{u}_K, \mathcal{u}_L$ and \mathcal{u}_X on K, L and X , respectively. As K and L are compact \mathcal{u}_K and \mathcal{u}_L have limits $x_K \in K$ and $x_L \in L$. They are also limits of \mathcal{u}_X in X . By supersobriety of X the set $\mathsf{L}(\mathcal{u}_X)$ has a largest element x , so $x_K, x_L \sqsubseteq x$ with respect to the specialisation order. Now, K and L are saturated which implies $x \in K \cap L$. Moreover, since \mathcal{u}_X is the filter in X generated by \mathcal{u} , the point x is also a limit of the original ultrafilter \mathcal{u} . This shows that every ultrafilter on $K \cap L$ converges.

“(2) \Rightarrow (1)”: Let \mathcal{u} be an ultrafilter with $\mathsf{L}(\mathcal{u}) \neq \emptyset$. Then the set $\mathsf{L}(\mathcal{u}) = \bigcap \{\overline{F} \mid F \in \mathcal{u}\}$ is closed, and we will show that it is irreducible. Assume $\mathsf{L}(\mathcal{u}) = A_1 \cup A_2$ where both are proper subsets. Then there are elements $a_i \in \mathsf{L}(\mathcal{u}) \setminus A_i$, for $i \in \{1, 2\}$, and as X is locally compact we find compact saturated neighbourhoods K_i satisfying $a_i \in \text{int}(K_i)$ and $K_i \cap A_i = \emptyset$. The set $K_1 \cap K_2$ is compact saturated by assumption and we get $(K_1 \cap K_2) \cap (A_1 \cup A_2) = (K_1 \cap K_2) \cap \mathsf{L}(\mathcal{u}) = \emptyset$. The sets K_1 and K_2

are neighbourhoods of the limits a_1 and a_2 and we infer $K_1, K_2 \in \mathcal{U}$ and thus $K_1 \cap K_2 \in \mathcal{U}$. But this means that the restriction

$$\mathcal{U}|_{K_1 \cap K_2} := \{F \cap K_1 \cap K_2 \mid F \in \mathcal{U}\} = \{F \in \mathcal{U} \mid F \subseteq K_1 \cap K_2\}$$

is an ultrafilter on the compact set $K_1 \cap K_2$ and so converges to a point x of that set. Thus \mathcal{U} also converges to $x \in K_1 \cap K_2$ which is disjoint from $L(\mathcal{U})$, a contradiction. This implies that $L(\mathcal{U})$ is an irreducible closed set, and by sobriety of X it must be the closure of a point. \square

Definition 1.3.3. A space is called *stably compact* if it is sober, compact, locally compact and satisfies the equivalent conditions of the proposition.

As supersobriety implies sobriety a space is stably compact if it is locally compact and every ultrafilter has a largest limit—the latter requirement comprises supersobriety and compactness. In domain theory compactness is quite a mild condition as, for example, every domain with \perp is trivially compact. On the localic side we can express compactness as $1 \ll 1$.

In a stably compact space we can take finite intersections of compact saturated sets. But in any sober space filtered intersections of compact saturated sets are again compact saturated by the Hofmann-Mislove Theorem, or more precisely Corollary 1.2.14. Hence, for a stably compact space X we can infer from Proposition 1.2.19 that $(\mathcal{K}(X), \supseteq)$ is a continuous lattice where suprema are given by intersection and finite infima by union.

As a consequence the closed sets of the co-compact topology for stably compact spaces are precisely the compact saturated sets for the original topology. We will improve on this result in the next section.

Like sober spaces, stably compact spaces are well-behaved under products and coproducts:

Proposition 1.3.4. *Stably compact spaces are closed under arbitrary products and finite coproducts.*

Proof. Sober spaces are closed under these constructions as we have seen in Proposition 1.2.6. Compactness and local compactness follow from Tychonov's Theorem in the case of products, and for finite coproducts they are trivial.

For the remaining condition take an ultrafilter \mathcal{U} on $\prod_i X_i$, a product of stably compact spaces. The projection of \mathcal{U} onto an individual compact supersober space X_i is also an ultrafilter and hence has a largest limit x_i . Since the product topology is the topology of point-wise convergence we see that \mathcal{U} converges to $\langle x_i \rangle$. Moreover, the projections are continuous and thus monotone which implies that $\langle x_i \rangle$ is the largest limit of \mathcal{U} . This shows that $\prod_i X_i$ is supersober.

Given an ultrafilter \mathcal{U} on $X + Y$, where both spaces are compact supersober, we must have either $X \in \mathcal{U}$ or $Y \in \mathcal{U}$. This implies that all limits of

\mathcal{U} must lie in one of the two spaces. As this spaces is supersober, \mathcal{U} must have a largest limit. \square

The only reason why we cannot take arbitrary coproducts of stably compact spaces is that infinite coproducts of non-empty spaces are never compact. The above proof is easily adapted to show all the other properties.

1.3.2 Compact Pospaces

A *compact pospace* is just an ordered space with a compact topology. They were introduced in [Nac48], and as we will see there is a very close connection to stably compact spaces. Unfortunately, the necessary details are hard to find in one place in the literature: The Compendium [GHK⁺80] discusses the link to Stone duality and also characterises them as compact supersober spaces (see also [Law88]). More recently they have been suggested as the adequate T_0 substitute for Hausdorff spaces [Law91, Kop94]. This section is an exposition of the relevant facts taken from these different sources.

As for pospaces the definition of compact ordered space is self-dual, i.e. if (X, \sqsubseteq) is such a space then so is (X, \supseteq) . Often this means that we only have to prove one half of the propositions. When it is as obviously the case as for example in the following lemma we do not usually mention it in the proof.

We begin with a series of observations that will allow us to make the link to stably compact spaces explicit. Some of them are also of independent interest as they illuminate the structure of compact pospaces.

Lemma 1.3.5. *If K is a compact subset of a compact pospace X then $\uparrow K$ and $\downarrow K$ are compact, as well.*

Proof. Suppose we are given a net $(x_i)_i$ in $\uparrow K$. We can assume that it converges to a point $x \in X$, otherwise we replace it by a converging subnet which exists by compactness of X . We find a net $(y_i)_i$ in K such that for all i we have $y_i \sqsubseteq x_i$. The latter net has a converging subnet $(y_{i_j})_j \rightarrow y \in K$ as K is compact. We still have $(x_{i_j})_j \rightarrow x$ and since X is an ordered space we infer $y \sqsubseteq x$. This shows that x lies in $\uparrow K$ and thus that $\uparrow K$ is compact. \square

Lemma 1.3.6. *Suppose $A = \downarrow A$ and $B = \uparrow B$ are closed subsets of a compact pospace X such that $A \cap B = \emptyset$. Then there are open sets $U = \downarrow U$ and $V = \uparrow V$ that do not intersect and satisfy $A \subseteq U$ and $B \subseteq V$.*

Proof. A compact Hausdorff space is normal. So there are disjoint open sets \tilde{U} and \tilde{V} such that $A \subseteq \tilde{U}$ and $B \subseteq \tilde{V}$. Hence we have compact sets $X \setminus \tilde{U}$ and $X \setminus \tilde{V}$, and because of the previous lemma the sets $\uparrow(X \setminus \tilde{U})$ and $\downarrow(X \setminus \tilde{V})$ are also compact. We set $U := X \setminus \uparrow(X \setminus \tilde{U})$ and $V := X \setminus \downarrow(X \setminus \tilde{V})$. It is now straight forward to verify that U and V have the desired property. \square

The following lemma is a fact from general topology.

Lemma 1.3.7. *If $\sigma \subseteq \tau$ are comparable topologies, where σ is Hausdorff and τ is compact, then the topologies coincide.*

Proof. Note that a topology that is finer than a Hausdorff topology is also Hausdorff, and that one coarser than a compact topology is also compact as there are less open covers. Hence, both σ and τ are compact Hausdorff topologies. Now take $U \in \tau$. Its complement $X \setminus U$ is closed and hence compact for τ . Thus it is also compact for σ which implies that it is closed, and we see $U \in \sigma$. We conclude $\sigma = \tau$. \square

Remark. *As a consequence, a compact Hausdorff topology is a maximal compact and a minimal Hausdorff topology. The converse is false. Neither is every minimal Hausdorff topology compact nor is every maximal compact topology Hausdorff [SW63].*

We can now show that there are ‘enough’ upper open and lower open sets.

Proposition 1.3.8. *The lower open and upper open sets of a compact pospace form a subbasis of the topology.*

Proof. The topology generated by these sets is clearly coarser than the original topology. Given $x \not\leq y$ the sets $\uparrow x$ and $\downarrow y$ are disjoint and compact by Lemma 1.3.5. Because of Lemma 1.3.6 they can be separated by upper, respectively, lower open sets. This shows that the generated topology is Hausdorff, and by Lemma 1.3.7 it must be the original topology. \square

The next two results provide some extra information about the patch topology for supersober spaces. We need them as the final prerequisites for the main proposition.

Proposition 1.3.9. *Let X be a supersober space and suppose $K \subseteq X$ is a compact saturated subset. Then K is also compact with respect to the patch topology.*

Proof. Take an ultrafilter \mathcal{U} on K and extend it to an ultrafilter \mathcal{U}' on X . As it has a limit in K its largest limit x lies in the saturated set K . We claim that \mathcal{U}' also converges to x with respect to the patch topology. For this it suffices to check that all co-compact neighbourhoods $x \in X \setminus L$, where L is compact, are in \mathcal{U}' . If such a neighbourhood is not in \mathcal{U}' then the ultrafilter has a $\Omega(X)$ -limit in the corresponding compact saturated set L . This limit must be smaller than or equal to x and hence $x \in L$, a contradiction. \square

Lemma 1.3.10. *Let X be a supersober space. Then the patch-open upper sets are precisely the original open sets.*

Proof. Closed sets are exactly the sets satisfying the property that all converging ultrafilters given on them converge only to points in the set. For a supersober space these are the lower sets A such that the largest limit of any such converging ultrafilter lies in A .

Let us say that a set A has property (\dagger) if the largest limit of converging ultrafilters on A is an element of A . In addition to closed sets the compact saturated sets also satisfy (\dagger) . We claim that the sets with (\dagger) are closed under arbitrary intersections and finite unions. For intersections this is clear. Suppose we are given a converging ultrafilter \mathcal{U} on $A \cup B$, where both sets have (\dagger) . Either A or B must be an element of \mathcal{U} and, consequently, the largest limit of \mathcal{U} lies in that set.

So, the sets satisfying (\dagger) form the closed sets of a topology which is clearly finer than the patch topology, and the closed sets of the original topology are the lower sets with (\dagger) . Hence the patch-closed lower sets must be closed with respect to the original topology.

By taking complements this we get that the original open sets are precisely the upper patch-open sets. \square

We now come to the central result of this section. Essentially it can be understood as yet another equivalent description of stably compact spaces.

Theorem 1.3.11. *For a stably compact space X the specialisation order makes X_π into a compact ordered space. Conversely, for a compact ordered space (X, \sqsubseteq) the open upper sets $\uparrow U = U \in \Omega(X)$ form the topology for a stably compact space X^\uparrow , and the two operations are mutually inverse.*

Moreover, for a stably compact space X the upper closed sets of X_π are precisely the compact saturated sets of X .

Proof. The space X_π with the specialisation order is a pospace by Proposition 1.1.4. That it is also compact follows from Proposition 1.3.9.

Let us conversely assume that (X, \sqsubseteq) is a compact pospace. To see that X^\uparrow is locally compact take $x \in U$ where $U = \uparrow U$ is open. Then $\uparrow x$ and $X \setminus U$ are disjoint closed upper, respectively, lower sets which by Lemma 1.3.6 we can separate by disjoint open sets $V = \uparrow V$ and $W = \downarrow W$. We get $x \in V \subseteq X \setminus W \subseteq U$ and $X \setminus W$ is a compact neighbourhood of x as required.

We next show compactness and supersobriety in one go. Every ultrafilter \mathcal{U} on X converges to a unique point x . We claim that x is also the largest limit of \mathcal{U} with respect to the coarser topology of X^\uparrow . The topology being coarser it is clearly a limit. For $y \not\sqsubseteq x$ we find disjoint open sets $U = \uparrow U$ and $V = \downarrow V$ such that $y \in U$ and $x \in V$. But as \mathcal{U} converges to x we get $V \in \mathcal{U}$ and hence $U \notin \mathcal{U}$. So, \mathcal{U} cannot converge to y in X^\uparrow . This shows that every ultrafilter on X^\uparrow has a largest limit completing the proof that X^\uparrow is stably compact.

It remains to show that the two translations are mutually inverse. In Lemma 1.3.10 we have already seen that $(X_\pi)^\uparrow = X$.

For the other composition take a compact pospace (X, \sqsubseteq) . We first look at the order: As in X^\uparrow all open sets are upper sets the topology is coarser than the Alexandrov topology. The sets $\downarrow x$ are compact by Lemma 1.3.5 and hence closed. This shows the topology of X^\uparrow is finer than the lower topology and thus has \sqsubseteq as its specialisation order. In the light of Proposition 1.3.8 the topology of $(X^\uparrow)_\pi$ is at least as fine as the original topology. But as we already know that $(X^\uparrow)_\pi$ carries a compact Hausdorff topology it must be the original topology because of Lemma 1.3.7.

The final assertion of the proposition is an immediate consequence of Proposition 1.3.9 and the fact that a compact set is also compact for any coarser topology. \square

The last observation of the proposition can be rephrased as follows: Under the above isomorphism, reversing the order of a compact ordered space corresponds to taking the co-compact topology on a stably compact space. This makes the slogan precise that in stably compact spaces open and compact saturated sets play dual roles.

Corollary 1.3.12. *If X is a stably compact space then the co-compact topology also defines a stably compact space X_κ . Taking the co-compact topology of a space is an involution and, moreover, we have $\Omega(X_\kappa) \cong \mathcal{K}(X)$ and $\mathcal{K}(X_\kappa) \cong \Omega(X)$. Both of these isomorphisms are given by complementation.*

Remark. *This can also be seen in the light of Lawson duality for complete semilattices [Law79]. One defines the dual of a semilattice to be the set of semi-lattice morphisms, i.e. Scott-continuous functions that preserve \wedge , to $\mathbf{2}$ ordered point-wise. These are essentially the Scott-open filters and hence for a locale the compact saturated sets of the corresponding space. Not all semilattices have duality, i.e. are isomorphic to their double dual. But as we have seen the topologies of stably compact spaces do. Going to the Lawson dual on the localic side corresponds to taking the co-compact topology for the space.*

Next, we are going to generalise the correspondence between stably compact spaces and compact ordered spaces to the non-compact case. As we will see most of the work has already been done in the above theorem. Some of the ideas needed for the locally compact case are taken from an unpublished note by Klaus Keimel and Regina Tix [KT96] where the authors tackle the general situation directly.

It is very easy to compactify a T_0 space X ; one just adds a bottom element \perp and one open set $X \cup \{\perp\}$. This construction is known as *lifting* and we call the resulting space X_\perp . It is clear that if X satisfies the equivalent conditions of Theorem 1.3.2 then so does X_\perp which then implies that X_\perp

is stably compact. Conversely, if a space X has a least element \perp then we can strip it off by removing the element from the space and the open set X from the topology as it is the only open set that contains \perp . We call the resulting space $X \setminus \{\perp\}$. It is again obvious that for a stably compact X with least element the resulting space $X \setminus \{\perp\}$ still satisfies the conditions of Theorem 1.3.2. Moreover, these operations are mutually inverse.

In the Hausdorff case we have to be slightly more subtle. We disregard the order at first since it is not an intrinsic part of the topology. Given a locally compact Hausdorff space X we get the *one point compactification* X_* by adding a new point $*$ and define the topology to be

$$\Omega(X) \cup \{X_* \setminus K \mid K \subseteq X \text{ compact}\}.$$

It is straightforward to verify that this does indeed define a topology and that it is compact. Given a compact Hausdorff space X and an element $x \in X$ we can delete this point by endowing the set $X \setminus \{x\}$ with the topology $\{U \in \Omega(X) \mid x \notin U\} = \{U \setminus \{x\} \mid U \in \Omega(X)\}$. This process yields a locally compact space and if we keep track of the distinguished point in the space the two operations are mutually inverse.

If X is a compact ordered space with a least element \perp then $X \setminus \{\perp\}$ is a locally compact ordered space. Unfortunately, we cannot simply add a bottom element to a locally compact ordered space to get a compact ordered space because in general the resulting space need not be a pospace.

As an example consider the one point compactification $\mathbb{R}_{-\infty}$ of the real numbers. We have $1 \not\leq -\infty$ but every neighbourhood of $-\infty$ contains arbitrarily large real numbers and hence this space cannot be an ordered space.

The following lemma characterises the pospaces for which we can add a least element in this fashion.

Lemma 1.3.13. *For a locally compact ordered space X the one point compactification X_\perp , where \perp is a new smallest element, is an ordered space if and only if for every compact set $K \subseteq X$ the upper set $\uparrow K$ it generates is again compact.*

Proof. The latter condition is necessary because it holds in every compact ordered space (Lemma 1.3.5).

Suppose, for the converse implication that for every compact K the set $\uparrow K$ is compact. We want to show that \sqsubseteq_X is closed, and to this end we take two elements $x \not\sqsubseteq y$. We have to find neighbourhoods U and V for these points such that for all $x' \in U$ and all $y' \in V$ we have $x' \not\sqsubseteq y'$. If both x and y come from X then this is no problem as X is an ordered space. The only other case is $x \in X$ and $y = \perp$. We take a compact neighbourhood K of x in X and by assumption $\uparrow K$ is also a compact neighbourhood. Then $\uparrow K$ and $X_\perp \setminus \uparrow K$ are neighbourhoods of x and $y = \perp$, respectively, that have the required property. \square

We call the spaces that satisfy the equivalent conditions of the lemma *properly ordered*. There is obviously an order-dual version of this lemma characterising the spaces for which we can add the new point as a top element. Since we focus on the link to stably compact spaces this is, however, not as interesting for us.

With the machinery we have just set up we get the general correspondence between locally compact pospaces and stably compact spaces *sans compactness* as a corollary of Theorem 1.3.11.

Corollary 1.3.14. *If X is a locally compact supersober space, then the specialisation order makes X_π into a locally compact properly ordered space. Conversely, for a locally compact properly ordered space (X, \sqsubseteq) the open upper sets form the topology for a locally compact supersober space X^\uparrow . Furthermore, these two operations are mutually inverse.*

Proof. Given such a T_0 space X , we can form $(X_\perp)_\pi \setminus \{\perp\}$ which is a locally compact pospace by Theorem 1.3.11 and the considerations of the previous paragraphs. As $(X_\perp)_\pi$ is a compact pospace it is also properly ordered, and it is easy to check that we have $(X_\perp)_\pi \setminus \{\perp\} = X_\pi$.

Conversely, suppose X is a properly ordered locally compact space. Then we get the space $(X_\perp)^\uparrow \setminus \{\perp\} = X^\uparrow$. The operations are mutually inverse because their integral parts are. \square

1.3.3 Stably Compact Domains

A *stably compact* domain is just a continuous domain whose Scott topology makes it a stably compact space. These domains can be characterised via their Lawson topologies.

Proposition 1.3.15. *A Scott-compact domain is stably compact if and only if it is Lawson-compact.*

Proof. If X is a stably compact domain then its patch topology is a compact pospace by Theorem 1.3.11. As the patch and the Lawson topology agree by Proposition 1.1.6, X is Lawson-compact.

For the converse, let us assume that X is Lawson-compact. By the Propositions 1.1.4 and 1.1.6 the Lawson topology with the original order forms a pospace, and because of Theorem 1.3.11 the upper Lawson-open sets form a stably compact topology. We claim that this is precisely the Scott topology. The proof of this claim is similar to that of Lemma 1.3.10. It suffices to show that Lawson-closed sets are closed under directed suprema. This is clear for Scott-closed sets and the other subbasic closed sets $\uparrow x$ that generate the closed sets of the Lawson topology. The collection of all \sqcap^\uparrow -closed sets is clearly closed under taking arbitrary intersections. It is also closed under finite unions:

Let A and B be two \sqcup^\uparrow -closed sets and $D \subseteq A \cup B$ a directed subset of their union. If for every $x \in D$ there is an $x' \in D \cap A$ such that $x \sqsubseteq x'$, then $A \cap D$ is directed and satisfies $\sqcup^\uparrow D = \sqcup^\uparrow(A \cap D) \in A$. Otherwise there is an $x \in D$ such that $\uparrow x \cap A \cap D = \emptyset$, and in this case we have $\uparrow x \cap D \subseteq B$ and thus $\sqcup^\uparrow D = \sqcup^\uparrow(B \cap D) \in B$.

This implies that all Lawson-closed are closed under directed suprema which concludes the proof. \square

The next proposition shows that all the classes of special domains discussed in Section 1.1.2 are stably compact. All the categories considered there are subcategories of FS. Hence, we begin by studying FS domains in some more detail. Most of this material is taken from [JS98].

Lemma 1.3.16. *Let $f: X \rightarrow Y$ be a continuous function between sober spaces. Then the preimage function $f^{-1}[\cdot]: \Omega(Y) \rightarrow \Omega(X)$ and the induced function*

$$f_{\mathcal{K}}: (\mathcal{K}(X), \supseteq) \rightarrow (\mathcal{K}(Y), \supseteq) \quad K \mapsto \uparrow f[K]$$

are Scott-continuous.

Proof. The preimage function is a frame morphism regardless of X and Y , so it is in particular Scott-continuous. For $f_{\mathcal{K}}$ this is an immediate consequence of Proposition 1.2.16 since on the localic side this function is given by taking preimages of Scott-open filters. \square

We know from Lemma 1.1.9 that for a finitely separated endo-function f we have $f(x) \ll x$. The next lemma shows that for such an f the induced functions of Lemma 1.3.16 provide approximants for open and compact saturated sets as well.

Lemma 1.3.17. *Let $f: X \rightarrow X$ be a continuous function on a dcpo X that is finitely separated from the identity. For each open set O there is a compact saturated set K such that $f^{-1}[O] \subseteq K \subseteq O$, and for each compact saturated set K there is an open set O such that $K \subseteq O \subseteq f_{\mathcal{K}}(K)$.*

Proof. If M is a finite separating set for f and $O \subseteq Y$ open, then we can infer $f^{-1}[O] \subseteq \uparrow(M \cap O) \subseteq O$. The set $M \cap O$ is finite, which implies that $\uparrow(M \cap O)$ is compact, thus showing the first assertion.

Now suppose that K is compact. We already know that $f(x) \ll x$ holds for all $x \in X$ and so we infer $K \subseteq \uparrow f[K] \subseteq \uparrow f[K] = f_{\mathcal{K}}(K)$. \square

The next lemma essentially says that for an FS domain there are enough such approximants for open and compact saturated sets.

Lemma 1.3.18. *Let X be any domain and $\{f_i \mid i \in I\}$ a directed set of continuous endo-functions on X such that $\sqcup^\uparrow f_i = \text{id}_X$. Then we have $O = \bigcup^\uparrow f_i^{-1}[O]$ and $K = \bigcap_\downarrow ((f_i)_{\mathcal{K}})(K)$ for every open set O and every compact saturated set K .*

Proof. The claim for open sets follows immediately from Lemma 1.2.10.

For compact saturated sets we use the localic form of the functions $(f_i)_{\mathcal{X}}$ which simply take preimages of Scott-open filters $F \subseteq \Omega(X)$ under the corresponding frame homomorphisms (see Proposition 1.2.16). Hence, we have to calculate $\bigcup^{\uparrow} (\Omega f_i)^{-1}[F]$. We apply Lemma 1.2.10 to the continuous lattice $\Omega(X)$, also endowed with the Scott topology, and the original space X and get

$$\bigcup^{\uparrow} (\Omega f_i)^{-1}[F] = \left(\bigsqcup^{\uparrow} (\Omega f_i)\right)^{-1}[F] = \Omega\left(\bigsqcup^{\uparrow} f_i\right)^{-1}[F] = \Omega(\text{id}_X)^{-1}[F] = F.$$

□

Now we can prove the main result of this section.

Proposition 1.3.19. *Every FS domain is stably compact.*

Proof. Let $(f_i)_i$ be a family of finitely separated functions on the domain X whose supremum is the identity, and let $(M_i)_i$ be corresponding finite separating sets. For any i we have $\uparrow M_i = X$ which shows that X is compact.

Suppose K and K' are compact saturated sets, i.e. the intersection of the open sets that contain them. Thus, we can write $K \cap K'$ as a filtered intersection of open sets $O \cap O'$ where $K \subseteq O$ and $K' \subseteq O'$. By the previous lemma we have $K = \bigcap_{\downarrow} (f_i)_{\mathcal{X}}(K) \subseteq O$, and analogously for K' , and from the Hofmann-Mislove Theorem in the form of Corollary 1.2.14 we can infer that there is an i such that $(f_i)_{\mathcal{X}}(K) \subseteq O$ and $(f_i)_{\mathcal{X}}(K') \subseteq O'$. This implies $K \subseteq f_i^{-1}[O]$, $K' \subseteq f_i^{-1}[O']$ and thus

$$K \cap K' \subseteq f_i^{-1}[O] \cap f_i^{-1}[O'] = f_i^{-1}[O \cap O'] \ll O \cap O',$$

where the final ‘way below’ follows from Lemma 1.3.17 and Lemma 1.2.17. This shows that the filter that these sets $O \cap O'$ generate in $\Omega(X)$ is Scott-open. As $K \cap K'$ is the intersection of this filter the set $K \cap K'$ is also compact saturated. This completes the proof that X is stably compact. □

Complete Distributivity

As we know the completely distributive lattices are precisely the topologies of continuous domains (Proposition 1.2.23). Throughout our investigation of stably compact spaces we have stressed the point that for them compact saturated and open sets play equally important roles. So one might hope that for a stably compact domain X the lattice $(\mathcal{K}(X), \supseteq)$ is also completely distributive. Unfortunately, this is generally not the case.

Theorem 1.3.20. *If X is a stably compact domain then $\mathcal{K}(X)$ is completely distributive if and only if X is bi-continuous and the co-compact topology agrees with the Scott topology on the order dual of X .*

Proof. The lattice $\mathcal{K}(X)$ is isomorphic to $\Omega(X_\kappa)$ which by Corollary 1.3.12 is the topology of a stably compact space. This sober space has the same points as X and the order of specialisation is the dual of the original order because of Proposition 1.1.2.

The lattice $\mathcal{K}(X) \cong \Omega(X_\kappa)$ is completely distributive if and only if X_κ is a domain with the Scott topology (Proposition 1.2.23). \square

To give an explicit example that shows that continuous domains, in general, are far from being bi-continuous we make the link to complete distributivity explicit: The completely distributive lattices are exactly the distributive bi-continuous lattices. This result appears in several places in the Compendium [GHK⁺80]. Since this is not too difficult to show we prove it here.

In a complete lattice L we define $x \lll y$ if whenever $y \leq \bigvee M$ holds there is an element $m \in M$ such that $x \leq m$. With this auxiliary relation we get Raney's characterisation of complete distributivity [Ran53].

Proposition 1.3.21. *A complete lattice is completely distributive if and only if for every x we have $x = \bigvee \{y \mid y \lll x\}$.*

Proof. See [AJ94, Theorem 7.1.3 and 7.1.1]. \square

The relation \lll is closely related to the ordinary order of approximation. The following connection is an immediate consequence of the definitions.

Lemma 1.3.22. *If p is a \vee -prime element of a complete lattice and $p \ll x$ then we have $p \lll x$.*

To get enough primes in our situation we use the following proposition that goes essentially back to Birkhoff [Bir67]. It appears as Theorem 7.1.7 in [AJ94].

Proposition 1.3.23. *In a continuous lattice every element is an infimum of \wedge -irreducible elements.*

Proposition 1.3.24. *A complete lattice is completely distributive if and only if it is distributive and bi-continuous.*

Proof. One implication follows directly from Proposition 1.3.21 and the fact that complete distributivity is a self-dual concept.

For the other direction let L be distributive and bi-continuous. By the previous theorem this implies that there are enough \vee -irreducible elements and they are \vee -primes as L is distributive. Every element $x \in L$ is the directed supremum of elements $y \ll x$ and these y 's are suprema of the \vee -primes below them. This implies that x is the supremum of such primes p satisfying $p \leq y \ll x$, and by Lemma 1.3.22 this means $p \lll x$. It follows from Proposition 1.3.21 that L is completely distributive. \square

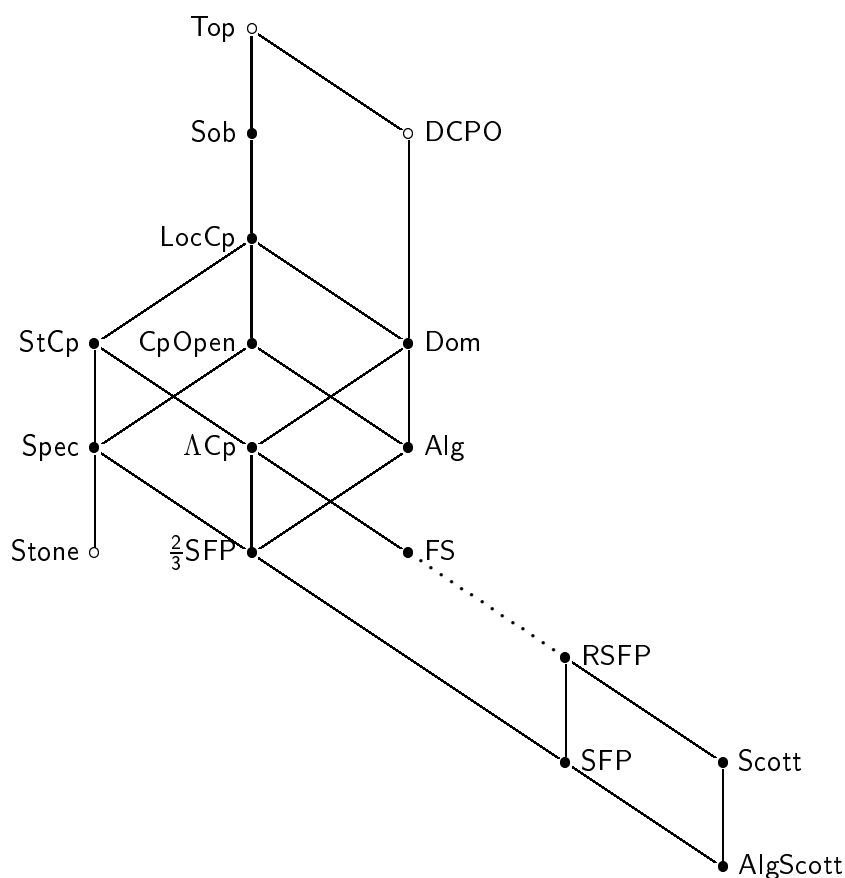


Figure 1.2: Hierarchy of categories of topological spaces II

Now let X be the topology of a locally compact sober space that is not endowed with the Scott topology, for example any non-discrete locally compact Hausdorff space. Then X is not completely distributive and hence not bi-continuous by the previous proposition. If we consider the Scott topology on X we get a stably compact space such that $\mathcal{K}(X)$ is not completely distributive because of Theorem 1.3.20.

1.3.4 Hierarchy of Stone Dualities II

We now complete the hierarchy of Stone dualities that we started in Section 1.2.3. For a survey of categories involved see Figure 1.2, which might also be helpful as a map to the rest of this section.

We call a locale *arithmetic* if it is continuous, compact, i.e. $1 \ll 1$, and the order of approximation is multiplicative. As we have seen in Theorem 1.3.2 these locales correspond to stably compact spaces.

Proposition 1.3.25. *The Stone duality for locally compact spaces as given in Proposition 1.2.18 restricts and co-restricts to stably compact spaces and arithmetic locales.*

Combining this with algebraicity of the locale yields stably compact spaces which have a basis of compact-opens because of Proposition 1.2.21. The importance of this concept is that we now have a basis which is closed under finite intersections and finite unions. Hence, the elements of $K\Omega(X)$ can be used as tokens for a finitary description of these spaces; the space itself, or its corresponding locale to be more precise, arises as the ideal completion of this lattice. We have thus in a way changed our localic description from an infinitary to a finitary one. The analogue for the non-algebraic case is the main theme of this thesis.

Definition 1.3.26. *A spectral space is a stably compact space which has a basis of compact-open sets.*

From our discussion it is clear that these spaces correspond to algebraic arithmetic locales. But we can go one step further by just looking at the distributive lattice $K\Omega(X)$. In the corresponding category the morphisms are not functions but certain *approximable relations* (see [AJ94, Def. 7.2.24]) which code frame morphisms, much in the spirit of approximable relations between bases of domains (see [AJ94, Def. 2.2.27]). We do not go into the details here, but we will come back to these relations in Section 3.1.2 where we present them in a slightly different form.

Proposition 1.3.27. *The equivalence of Proposition 1.2.18 (co-)restricts to spectral spaces and arithmetic algebraic locales.*

These categories are also equivalent to the category of distributive lattices with least element and approximable relations.

The space corresponding to such a distributive lattice is known as its *spectrum*. The functor from the category of distributive lattices to the category of spectral spaces giving rise to the equivalence of the above theorem is called *spec*.

Remark. *At this point we are very close to Stones original representation theorem for Boolean algebras [Sto36]. The distributive lattices of the previous theorem turn out to be boolean algebras if and only if the corresponding spaces are compact and Hausdorff. In this case the morphisms corresponding to continuous maps are Boolean algebra homomorphisms and not just approximable relations. Thus we get the classical duality between the category of Boolean algebras and the category of compact, totally disconnected Hausdorff spaces, the so-called Stone spaces.*

We now consider spaces that are algebraic domains with their Scott topology. In Section 1.2.3 we have seen that algebraic domains correspond

to completely distributive algebraic lattices. As observed before, in an algebraic domain every compact-open set is a finite union of principal filters $\uparrow x$ for compact x . These principal filters in turn are exactly the \vee -prime elements in the lattice of compact open subsets. The condition of being generated by prime elements is stronger than complete distributivity, and thus characterises the locales that come from algebraic domains.

Proposition 1.3.28. *The equivalence of Proposition 1.2.18 (co-)restricts to an stably compact algebraic domains with their Scott topology and arithmetic algebraic completely distributive lattices.*

The corresponding category of approximable relations is the one of distributive lattices in which every element is a finite supremum of \vee -prime elements.

We now want to combine this with stable compactness. The category of stably compact algebraic domains is not usually studied in domain theory because it fails to be cartesian closed. Such domains might be called 2/3-bifinite domains in the light of Plotkin's "2/3-SFP Theorem" [Plo81]. It says that a stably compact algebraic domain is MUB-complete, and that for every finite set of compact elements the set of minimal upper bounds is again finite (compare Proposition 1.1.7). The only extra condition required for bifinite domains is that finite sets of compact elements must have finite MUB-closures. In general, this is not the case for stably compact domains. To complete our localic description of bifinite domains we have to translate precisely this condition. This is the contents of the following proposition which is also the most restricted Stone duality we investigate in this section:

Proposition 1.3.29. *A lattice L is isomorphic to the compact-opens of a bifinite domain if and only if every element is a finite supremum of \vee -primes and if for every finite subset $M \subseteq L$ there is a finite set $N \supseteq M$ such that*

$$(\forall A \subseteq N)(\exists B \subseteq N) \bigwedge A = \bigvee B.$$

This concludes our exposition of Stone duality. Abramsky's domain theory in logical form uses a localic description of domains which is very close to the one indicated in the previous proposition. Later on we will discuss a finitary description of stably compact spaces for our extension of the programme to continuous domains.

1.4 Abramsky's Domain Theory in Logical Form

This section gives a quick overview of Samson Abramsky's work [Abr91b], and we will not rely on any of the material discussed here later in the text. Its purpose is to give an exposition of the established domain theory in logical form, so we can compare it to the results of this thesis. This also allows us to discuss the differences in approach. The introduction given here only sketches the main ideas; for details see [Abr91b] or [AJ94, Section 7.3].

Domain theory in the context of the title means performing the usual constructions of domain theory like products, sums, function spaces, solving recursive domain equations, and power domains. In the light of the previous section we now want to describe these constructions solely in terms of the lattices of compact-open sets of the domains involved.

1.4.1 Prelocales

As we think of the open sets of a domain as certain observable properties the constructions are presented in the form of a logical system. Hence, we do not handle the elements of the lattices corresponding to the domains directly, but these lattices can be understood as the Lindenbaum algebras for the logic. This means that we have to formulate the preordered equivalent of the lattices described in Proposition 1.3.29.

Let us begin with some notes on preordered sets. If (P, \lesssim) is a preorder, then \lesssim induces an equivalence relation \approx which is the intersection of \lesssim and \gtrsim . The quotient P/\approx is a poset, and from a categorical viewpoint this quotient can also be seen as a skeleton of the preorder seen as a category. In a preorder we understand the usual order theoretic concepts as meaning that the corresponding property holds in the quotient. For example, an element z is *a* (rather than *the*) supremum of x and y if $x, y \lesssim z$ and for all $z' \gtrsim x, y$ we have $z \lesssim z'$. Clearly, all suprema of x and y are equivalent.

Definition 1.4.1. A *domain prelocale* $(A; \lesssim; \vee, \wedge, 0, 1; C)$ is an algebra of type $(2, 2, 0, 0)$ with a unary predicate C such that $a \vee b$ is a supremum for a and b , $a \wedge b$ is an infimum, 0 is a least, 1 is a largest element, and C characterises the \vee -prime elements.

Furthermore, every element must be equivalent to a finite supremum of \vee -primes, $C(1)$ must hold for all finites subsets M of $C(A)$ —the set of all \vee -primes in A —there has to be a finite superset $N \subseteq C(A)$ of M such that $(\forall O \subseteq N)(\exists P \subseteq N) \bigwedge O \approx \bigvee P$.

The predicate C characterises the elements of the prelocale that can be thought of as points: As we have observed before, the \vee -prime compact open sets in an algebraic domain are precisely the upper sets $\uparrow k$ for a compact k . This predicate is only needed for the function space construction as it is notoriously hard to describe function spaces in purely logical terms.

In [Abr91b] there is an additional unary predicate to perform the coalesced sum construction. Currently, we cannot tell whether a ‘point’ of the domain, i.e. a \vee -prime element, represents \perp , but we will not go into the details of the different constructions anyway.

The above definition caters for two extremes: As we have seen in Proposition 1.3.29 the compact-opens of a bifinite domain gives rise to a domain prelocale where \lesssim is actually an order. The logical approach yields prelocales whose elements are terms made up from certain generators and the operations \vee , \wedge , 0 and 1 . These prelocales can then be seen as purely syntactic objects.

For a prelocale P , the relation \approx is not just an equivalence but a congruence with respect to the operations of P . The quotient P/\approx is a prime generated and hence distributive lattice.

1.4.2 Prelocalic Description of Domains

Prelocales are used to represent concrete domains and hence we are interested to describe when two prelocales represent the same domain.

Definition 1.4.2. A *pre-isomorphism* between domain prelocales A and B is a monotone and order-reflecting function $\phi: A \rightarrow B$ such that every element in B is equivalent to an element in the image of ϕ .

If B is of the form $\mathbf{K}\Omega(X)$ for a bifinite domain X then we call A a *prelocalic description* of X via ϕ . In this case the pre-isomorphism is usually denoted by semantic brackets $\llbracket \cdot \rrbracket$.

Remark. *The definition of pre-isomorphism is tantamount to the two prelocales A and B being equivalent as categories (see [Mac71, IV.4, Theorem 1]). In the case of a prelocalic description the conditions boil down to surjectivity, monotonicity and order-reflection.*

From the remark it is obvious that a pre-isomorphism $\phi: A \rightarrow B$ preserves and reflects suprema, infima, least and largest elements, as well as primeness. Hence, it restricts to a map ϕ^0 between $C(A)$ and $C(B)$.

Given a prelocalic description $\llbracket \cdot \rrbracket: A \rightarrow \mathbf{K}\Omega(X)$ of a domain X we get an isomorphism of domains $\tau: \text{spec}(A) \rightarrow X$ as indicated in the following diagram.

$$\begin{array}{ccc}
 & \text{spec}(A) & \\
 & \downarrow & \searrow \tau \\
 (\text{spec}(\llbracket \cdot \rrbracket))^{-1} & & \\
 & \downarrow & \\
 \text{pt}(\Omega(X)) \cong \text{spec}(\mathbf{K}\Omega(X)) & \xrightarrow{\eta^{-1}} & X
 \end{array}$$

Formally, the functor spec can only be applied to distributive lattices and not to prelocales, but this can easily be remedied by either pre-composing with the factorisation by \approx or by a straightforward extension of the definition of spec to prelocales.

The converse of this observation holds as well, which confirms our intuition about the meaning of prelocalic descriptions:

Proposition 1.4.3. *For a domain prelocale A and a bifinite domain X there exists a pre-isomorphism $\llbracket \cdot \rrbracket: A \rightarrow \mathbf{K}\Omega(X)$ if and only if $\text{spec}(A)$ and X are isomorphic.*

We now turn our attention to embeddings and sub-prelocales. The importance of embeddings is that they are needed for the bilimit construction of domains.

Definition 1.4.4. We say that a prelocale A is a *sub-prelocale* of another one B if it is a subalgebra with respect to the four prelocale-operations and if the relations C and \lesssim on A are the restrictions of their counterparts on B .

This definitions captures our intended meaning:

Proposition 1.4.5. *If A is a sub-prelocale of B then there is an embedding projection pair between $\text{spec}(A)$ and $\text{spec}(B)$.*

1.4.3 Domain Constructions

We now have the machinery to outline what, concretely, doing domain theory by prelocalic descriptions is about. We illustrate it by considering any binary domain constructor $F: \text{SFP} \times \text{SFP} \rightarrow \text{SFP}$.

Let us assume we have two domains X and X' and prelocalic descriptions $\llbracket \cdot \rrbracket_A: A \rightarrow \mathbf{K}\Omega(X)$ and $\llbracket \cdot \rrbracket_{A'}: A' \rightarrow \mathbf{K}\Omega(X')$ for them. We want to construct a prelocale $T(A, A')$ from A and A' —not by looking at X , X' or $F(X, X')$, of course—and a prelocalic description

$$\llbracket \cdot \rrbracket: T(A, A') \rightarrow \mathbf{K}\Omega(F(X, X')).$$

Furthermore, this construction has to be natural with respect to sub-prelocales so that we can solve recursive domain equations which make use of this construction. That is to say for sub-prelocales B and B' of A and A' describing sub-domains of X and X' we demand that $T(B, B')$ is a sub-prelocale of $T(A, A')$ and that

$$\begin{array}{ccc} \text{spec}(T(B, B')) & \xrightarrow{I} & \text{spec}(T(A, A')) \\ \tau_B \downarrow & & \downarrow \tau_A \\ F(Y, Y') & \xrightarrow{F(e, e')} & F(X, X') \end{array}$$

commutes, where e , e' and I are the embeddings corresponding to sub-prelocale inclusions via Proposition 1.4.5.

There is a general technique that one can follow to define a type constructor on the prelocalic side. One advantage of having such a programme is that it identifies the points that follow from general principles and those steps that have to be devised and checked for each individual type constructor. To give an impression of how this is done we go through one rather straightforward example. For an explicit list of the individual steps of the programme see [AJ94, Section 7.3.2] and [Abr91b, Section 3.4].

Example 1.4.6 (Product Prelocale). We assume we are given two domains X and X' and their prelocalic descriptions via A and A' as in the generic scenario above. The compact-opens of $X \times X'$ are finite unions of products of compact-opens of X and X' . Hence, we use the elements $|A| \times |A'|$ as generators of the term algebra which is going to be the underlying set of our domain prelocale $A \times A'$.

Apart from the general rules for \lesssim , \vee , \wedge , 0 and 1 which guarantee that our structure becomes the preordered equivalent of a distributive lattice we state the following rules:

$$\frac{a \lesssim b \quad a' \lesssim b'}{(a, a') \lesssim (b, b')} \qquad \frac{}{0 \approx (0, a') \approx (a, 0)}$$

$$\frac{C(a) \quad C(a')}{C(a, a')}$$

We define the interpretation $\llbracket \cdot \rrbracket$ for generators as $\llbracket a, a' \rrbracket := \llbracket a \rrbracket \times \llbracket a' \rrbracket$ and extend it to the term algebra as a (pre-)lattice homomorphism.

The rules are obviously sound which implies that $\llbracket \cdot \rrbracket$ is monotone and restricts and co-restricts to a map $\llbracket \cdot \rrbracket^0 : C(A \times A') \longrightarrow \mathbf{K}\Omega(X \times X')$.

Because of distributivity we can prove $(a \wedge b, a' \wedge b') \approx (a, a') \wedge (b, b')$ and $(a \vee b, a' \vee b') \approx (a, a') \vee (b, a') \vee (a, b') \vee (b, b')$, and thus by induction we get that every element in $A \times A'$ is equivalent to a finite supremum of elements which satisfy the predicate C .

The map $\llbracket \cdot \rrbracket^0$ is clearly surjective and order-reflecting. It now follows from general considerations that this entails that $\llbracket \cdot \rrbracket$ is surjective and order-reflecting, as well. Hence, it is an pre-isomorphism, but as we already know that $\mathbf{K}\Omega(X \times X')$ is a domain pre-locale this provides a semantic proof that $A \times A'$ is also a domain prelocale.

Now, suppose we are given sub-prelocales B and B' . It is clear that $B \times B'$ is a subset of $A \times A'$ and that everything that can be proved in the former can also be proved in the latter. Again, it follows in general that the interplay of sub-prelocales and the prelocalic descriptions $\llbracket \cdot \rrbracket$ have only

to be verified for primes, i.e. we have to check the commutativity of the diagram

$$\begin{array}{ccc}
 C(B \times B') & \hookrightarrow & C(A \times A') \\
 \downarrow \llbracket \cdot \rrbracket^0 & & \downarrow \llbracket \cdot \rrbracket^0 \\
 C(\mathbf{K}\Omega(Y \times Y')) & \xrightarrow{i} & C(\mathbf{K}\Omega(X \times X'))
 \end{array}$$

where Y and Y' are the sub-domains induced by B and B' and i is the map induced by the two sub-prelocale containments: If b and b' represent $\uparrow x$ and $\uparrow x'$, respectively, and we denote the embeddings by $e: Y \hookrightarrow X$ and $e': Y' \hookrightarrow X'$, then both compositions map the pair (b, b') to $\uparrow e(x) \times \uparrow e(x')$.

This shows that the above rules do indeed define the product on the prelocalic side: The syntactic construction yields a prelocale $A \times A'$ and it corresponds to the product of the spaces corresponding to A and A' , respectively.

1.4.4 Logic

Once we have all the constructions for the prelocalic side we can interpret each type built up from these constructors as a domain prelocale. This is analogous to the usual procedure of interpreting a type by a domain.

We can then compare these two interpretations. Assuming that for each free variable the prelocales we choose are prelocalic descriptions for the corresponding domain it is clear from our setup that the resulting prelocale is a prelocalic description of the domain interpreting that type.

The constructions also yield a logical system for 'properties', i.e. the points of the domain prelocales. We have seen some of the rules in our example; for the complete system see [Abr91b, p. 49ff]. Our previous work immediately yields a semantics for this logic by mapping an element of a prelocale to the compact-open subset it represents in the corresponding domain. It turns out that the syntactic relation \lesssim and subset containment in the semantics agree; in other words, the semantics is sound and complete. Another interesting result is that the language of properties is decidable.

On top of that we can add terms of an untyped lambda calculus, extended by term constructors corresponding to the type constructors. We get judgements of the form $\Gamma \vdash M : \phi$, where M is a lambda term, ϕ a property (an element in a prelocale of appropriate type) and Γ a list of assumptions $x \mapsto \psi_x$ on the free variables of M . The intended meaning is that assuming the variables meet the requirements stated in Γ the term M satisfies ϕ .

There is a straightforward logical system for these judgements, and as usual most of the terms can be seen as encodings of the derivation of the

property ϕ in the previous system. From classical domain theory we know how to interpret such lambda terms in the domain interpreting the type of that term. Again this semantics is sound and complete with respect to the logical system for these judgements.

In [Abr91b, Section 4.4] there is also a corresponding exogenous logic in which the terms describe morphisms rather than elements of domains. This is an extension of the meta-language for cartesian closed categories (see [LS86]), but we will not go into this here as it is not relevant for our further considerations.

Chapter 2

Syntax

We now change our point of view to a purely syntactic one and study a Gentzen style sequent calculus. Its special feature is that the formulae on the left and right of the turnstile need not come from the same logical system. Such a sequent can be seen as a consequence between different domains of reasoning. The usual identity and cut rules do not make sense for sequents which connect different logical systems because they mix formulae from antecedent and succedent. This necessitates certain syntactic adjustments.

After an independent motivation for such a logical system we investigate the ingredients needed to set up this logic. In particular, we discuss an appropriate cut rule. It can be used as a basis for composition in a suitable category MLS of logical systems, and as we will see in the next chapter on semantics this category is equivalent to stably compact spaces with certain closed relations between them.

We also study cut elimination in this set-up. The upshot is—as might be expected—that we can push up cuts in a proof to cuts between atomic formulae. This can be used to do define new objects and morphisms from known ones which can be understood as performing domain constructions in logical form. As an example we construct products and coproducts in MLS. The procedure is a bit tedious, but shows exactly what is involved if one wants to do such constructions in a purely syntactic fashion.

2.1 Multi-Lingual Sequent Calculus

Our main objects of study are *sequents* in tradition of Proof Theory [Gen34]

$$\phi_1, \dots, \phi_n \vdash \psi_1, \dots, \psi_m$$

and we read this as “if all ϕ_i hold then at least one ψ_j holds”, as usual.

We want to allow the formulae ϕ_i to come from a different language than the ψ_j , i.e. we want to be able to consider a situation where inferences

are to be drawn between different logical systems. There are many situations where such a separation is desirable or even necessary. We discuss two of them.

Consider ordinary propositional logic. Someone could say

“It is very cold in here. I need to put on a jumper.”

thus drawing an inference from an observation about the temperature to a certain action. Note that there is nothing *logical* about this inference and, indeed, someone else might say

“It is very cold in here. I will turn on the heating.”

The inference relation in this example is a subjective one and there can be many different such relations. Although it is common to combine arbitrary propositions in logic we may wish to distinguish in a situation like this between propositions about the state of the environment and propositions about actions of a certain individual.

A second example is given by Hoare Logic. When we write a triple like

$$\{x > 0\} \quad x := -x \quad \{x < 0\}$$

we certainly do not mean that $x > 0$ logically implies $x < 0$, rather, we read this as

“If $x > 0$ holds before the execution of $x := -x$ then $x < 0$ holds afterwards.”

In this example every program fragment gives rise to a characteristic relationship between preconditions and postconditions. The logical formulae are (typically) all about the contents of program variables and there is no syntactic reason to keep pre- and postconditions separate, as in the previous example, but the separation becomes necessary because the formulae refer to the state at different times.

Having different languages for formulae on the left and the right of the turnstile \vdash forces us to restrict ourselves to *positive logic*, i.e. conjunction and disjunction. Moreover, we leave out the identity axiom scheme, $\phi \vdash \phi$. If the two domains of reasoning related by \vdash are different it is impossible to formulate it.

But even if they are the same this is justified by the fact that *observing* a certain state of the world does not always imply that the corresponding proposition is actually true, the reason being that our instruments for observing the world are not precise enough.

Many features of our logical system are a direct consequence of these special properties. Leaving out the identity axiom, for example, necessitates to check carefully how to retain some of its essential consequences.

The cut rule also has to be adjusted to this new situation. There are several related formulations of the appropriate cut rule, and we are going to study how they are related. Considering several versions of essentially the same rule may seem redundant at first, but the development of the theory shows that they all have their particular strengths. So we can always use the rule that is most suitable in any given situation.

Note that we allow classical sequents. At first glance there seems to be no point in this because there is no difference between intuitionistic positive logic and classical positive logic. However, this formulation emphasises the rather pretty self-symmetry of the whole set-up and mirrors the duality between open and compact subsets of the next chapter.

2.1.1 The Logic

We take a very liberal approach as to what structures the actual formulae of our sequents are drawn from. All we require is the presence of conjunction and disjunction, and the units \perp (falsity) and \top (truth). Each system embodies a certain ‘logic’ in the sense that certain formulae imply others. We capture the internal logic by referring to arbitrary $(2, 2, 0, 0)$ -algebras instead of syntactically defined sets of formulae.

Such an algebra could, for example, have been obtained as the Lindenbaum algebra by factorising the set of formulae of the system by logical equivalence, and in the next chapter we will study how to construct such algebras from any given stably compact space. At the other extreme, the syntactically defined set of formulae for a logical system can be regarded as a such an algebra, providing the logic contains the connectives of positive logic. Proposition 2.2.5 and Theorem 3.1.16 show that we could have settled for either extreme, but as we will see they have different applications. Hence, we choose to work in this more flexible framework. Regardless of the particular algebra at hand we will refer to its elements as *formulae*.

The logical part of our system is given by the rules

$\frac{}{\perp \vdash} \text{ (L}\perp\text{)}$	$\frac{\Gamma \vdash \Delta}{\Gamma \vdash \Delta, \perp} \text{ (R}\perp\text{)}$
$\frac{\Gamma \vdash \Delta}{\top, \Gamma \vdash \Delta} \text{ (L}\top\text{)}$	$\frac{}{\vdash \top} \text{ (R}\top\text{)}$
$\frac{\phi, \psi, \Gamma \vdash \Delta}{\phi \wedge \psi, \Gamma \vdash \Delta} \text{ (L}\wedge\text{)}$	$\frac{\Gamma \vdash \Delta, \phi \quad \Gamma \vdash \Delta, \psi}{\Gamma \vdash \Delta, \phi \wedge \psi} \text{ (R}\wedge\text{)}$
$\frac{\phi, \Gamma \vdash \Delta \quad \psi, \Gamma \vdash \Delta}{\phi \vee \psi, \Gamma \vdash \Delta} \text{ (L}\vee\text{)}$	$\frac{\Gamma \vdash \Delta, \phi, \psi}{\Gamma \vdash \Delta, \phi \vee \psi} \text{ (R}\vee\text{)}$

where, as usual, small Greek letters refer to individual and capital Greek letters to finite sets of formulae. Furthermore, a double line indicates that the rule can be used in both directions. The *backwards* rules are not present in the usual sequent calculus since there they are consequences of the identity and the cut rule. The essential difference in character between the *forwards* and the *backwards* rules will become apparent in Section 2.2.

The rules (L \wedge) and (R \vee) forces the comma which separates formulae on the left to refer to conjunction and the comma on the right to disjunction.

Note that we cannot refer to implication or negation in the logical systems as the corresponding rules would make it necessary to transfer formulae from one side to the other. However, the logical systems themselves may still support these connectives.

On the side of structural rules we will only refer to weakening

$$\boxed{\frac{\Gamma \vdash \Delta}{\Gamma', \Gamma \vdash \Delta, \Delta'} (W)}$$

and keep exchange and contraction implicit. Thus we are working with *sets* of formulae rather than sequences. The *forwards* rules (R \perp) and (L \top) are special cases of weakening.

As the examples above suggest, this calculus is not about finding tautologies but rather, each relation \vdash between formulae of two logical systems embodies a particular, possibly subjective, inference. Whatever the reasons are for holding such an inference as true, there are other inferences which should in such a situation also be held as true. The rules above formalise precisely this reasoning: If ϕ, ψ, Γ entails Δ then $\phi \wedge \psi, \Gamma$ should also entail Δ , and so on. Our objects of study are therefore relations between sets of formulae which are closed under the rules from above. We fix this in a definition:

Definition 2.1.1. For two algebras $\langle L; \wedge, \vee, \top, \perp \rangle$ and $\langle M; \wedge', \vee', \top', \perp' \rangle$ of type $(2, 2, 0, 0)$, a *consequence relation* \vdash from L to M is a binary relation between finite sets from L and M closed under (L \perp), (R \perp), (L \top), (R \top), (L \wedge), (R \wedge), (L \vee), (R \vee), and (W).

If, according to a consequence relation \vdash , the formula ϕ implies ψ , and if, according to a second relation \vdash' , ψ implies σ , then it makes sense to combine these two inferences and to say that ϕ implies σ according to the composition $\vdash \circ \vdash'$ of the two given consequence relations.

The obvious way to formulate the cut rule is the following:

$$\boxed{\frac{\Gamma \vdash \phi \quad \phi \vdash' \Delta}{\Gamma \vdash (\vdash \circ \vdash') \Delta} (\text{Cut})}$$

As consequence relations relate *sets* of formulae rather than single formulae one might look for a version that allows sets of formulae on both sides of the turnstile. A logically correct alternative is given by the following rule whose premise is meant to be read as two families of sequents, not as proof trees:

$$\boxed{\begin{array}{ccc} \Gamma \vdash \Delta_1 & \Theta_1 \vdash' \Lambda & \\ \vdots & \vdots & \\ \Gamma \vdash \Delta_n & \Theta_m \vdash' \Lambda & \\ \hline \Gamma (\vdash \circ \vdash') \Lambda & & \text{(Cut*)} \end{array}}$$

subject to the condition that for every choice function $f \in \prod_i \Delta_i$ there exists an index j so that $\Theta_j \subseteq \{f_1, \dots, f_n\}$

The intuition behind this cut rule and its side condition is that if Γ entails all the Δ_i 's then at least one formula in each Δ_i is true. If for every possibility, coded by a choice function, these formulae cover one of the Θ_j 's then Γ also entails Λ .

Remark. *The side condition includes the degenerate case $n = 0$. Then there is only one choice function, namely the empty one. Its image is \emptyset and thus the side condition says that one of the premises Θ_j must also be empty. This agrees with the intuition that $\vdash \Lambda$ means that the disjunction of the formulae in Λ is true and hence is a consequence of any set of premises Γ .*

Similarly, $m = 0$ is admissible if and only if one of the Δ_i is empty. Note, however, that $m = n = 0$ does not satisfy the side condition.

This cut rule looks rather asymmetric whereas the other rules are perfectly symmetric. That is to say if we take a rule and interchange left and right as well as the connectives \vee and \perp with their duals \wedge and \top then we again get a rule. We can, however, reformulate the side condition in a way that shows that it is also inherently symmetric:

Lemma 2.1.2. *The side condition of (Cut*) is equivalent to*

$$(\forall f \in \prod_i \Delta_i) (\forall g \in \prod_j \Theta_j) \{f_1, \dots, f_n\} \cap \{g_1, \dots, g_m\} \neq \emptyset.$$

Proof. Assume f and g are two such choice functions. If the side condition of (Cut*) holds then there is an index j such that $\Theta_j \subseteq \{f_1, \dots, f_n\}$ which implies that $g_j \in \Theta_j$ is in the intersection.

Conversely, if the side condition fails then there is a choice function $f \in \prod_i \Delta_i$ that does not cover any of the Θ_j . This implies that all sets $\Theta_j \setminus \{f_1, \dots, f_n\}$ are non-empty which yields a choice function $g \in \prod_j \Theta_j$ satisfying $\{f_1, \dots, f_n\} \cap \{g_1, \dots, g_m\} = \emptyset$. \square

As the next proposition shows both cut rules are in fact equivalent. The first rule (Cut) is usually easier to handle in proofs. An advantage of (Cut*) is that many properties of the system can be proved without using any logical rules, thus exposing their combinatorial character. Moreover, as we will see in Section 2.2, the seemingly more complicated rule (Cut*) is easier to handle from the point of view of cut elimination. We are going to use both cuts—and even a third formulation later—depending on which one is more convenient.

Proposition 2.1.3. *The cut rules (Cut) and (Cut*) are inter-definable.*

Proof. The rule (Cut) is clearly an instance of (Cut*) as the side condition is trivially satisfied.

For the other direction consider a cut

$$\frac{\begin{array}{c} \Gamma \vdash \Delta_1 \quad \Theta_1 \vdash' \Lambda \\ \vdots \quad \quad \quad \vdots \\ \Gamma \vdash \Delta_n \quad \Theta_m \vdash' \Lambda \end{array}}{\Gamma (\vdash \circ \vdash') \Lambda.} \text{ (Cut*)}$$

We prove that we can rewrite it in terms of (Cut) by induction on the number of choice functions in the product $\prod_{i=1}^n \Delta_i$.

The case where there is no such function has to be dealt with separately: We infer that one of the Δ_i is empty and we get

$$\frac{\frac{\Gamma \vdash}{\Gamma \vdash \perp} \text{ (W)} \quad \frac{\frac{}{\perp \vdash'} \text{ (L}\perp)}{\perp \vdash' \Lambda} \text{ (W)}}{\Gamma (\vdash \circ \vdash') \Lambda.} \text{ (Cut)}$$

If there is precisely one such function then all Δ_i are singletons $\{\delta_i\}$, and we can assume without loss of generality that $m = 1$ and $\Theta_1 \subseteq \{\delta_1, \dots, \delta_n\}$. We construct the proof

$$\frac{\frac{\begin{array}{c} \Gamma \vdash \delta_1 \\ \vdots \\ \Gamma \vdash \delta_n \end{array}}{\Gamma \vdash \delta_1 \wedge \dots \wedge \delta_n} \text{ (R}\wedge\text{)*} \quad \frac{\frac{\Theta_1 \vdash' \Lambda}{\delta_1, \dots, \delta_n \vdash' \Lambda} \text{ (W)}}{\delta_1 \wedge \dots \wedge \delta_n \vdash' \Lambda} \text{ (L}\wedge\text{)*}}{\Gamma (\vdash \circ \vdash') \Lambda} \text{ (Cut)}$$

where the bold lines indicate multiple applications of the respective rules.

Now for the induction step: If there is more than one choice function then one Δ_i contains at least two elements, say Δ_n . We are going to reduce the number of choice functions by replacing the sequent $\Gamma \vdash \Delta_n$ by

$\Gamma \vdash \bigvee \Delta_n$ as follows: For each (restricted) choice function $f \in \prod_{i=1}^{n-1} \Delta_i$ and $\phi \in \Delta_n$ there is an index j such that $\Theta_j \subseteq \{f_1, \dots, f_{n-1}, \phi\}$, and because of weakening we can assume that the two sets are equal.

We apply (LV) to these sequents and get $\bigvee \Delta_n, f_1, \dots, f_{n-1} \vdash' \Lambda$. Then we repeat this procedure for all elements of $\prod_{i=1}^{n-1} \Delta_i$ and consider the cut

$$\frac{\begin{array}{c} \Gamma \vdash \Delta_1 \\ \vdots \\ \Gamma \vdash \Delta_{n-1} \\ \Gamma \vdash \bigvee \Delta_n \end{array} \quad \begin{array}{c} \bigvee \Delta_n, f_1, \dots, f_{n-1} \vdash' \Lambda \\ \bigvee \Delta_n, g_1, \dots, g_{n-1} \vdash' \Lambda \\ \vdots \\ \bigvee \Delta_n, h_1, \dots, h_{n-1} \vdash' \Lambda \end{array}}{\Gamma (\vdash \circ \vdash') \Lambda} \text{ (Cut*)}$$

where f, g, \dots, h range over all such choice functions. The side condition is still satisfied as all possible choices are explicitly listed on the right, but the number of choice functions on the left is now the previous number divided by $|\Delta_n| > 1$. Hence, we can apply the induction hypothesis which concludes the proof. \square

We make one more observation about the logic to make some proofs later more convenient. We have weakening in the system and hence the following alternative multiplicative formulations of several logical rules are derivable. By slight abuse of notation we refer to them by the same names.

$$\boxed{\begin{array}{cc} \frac{}{\perp, \Gamma \vdash \Delta} \text{ (L}\perp\text{)} & \frac{}{\Gamma \vdash \Delta, \top} \text{ (R}\top\text{)} \\ \frac{\phi, \Gamma \vdash \Delta \quad \psi, \Gamma' \vdash \Delta'}{\phi \vee \psi, \Gamma, \Gamma' \vdash \Delta, \Delta'} \text{ (LV)} & \frac{\Gamma \vdash \Delta, \phi \quad \Gamma' \vdash \Delta', \psi}{\Gamma, \Gamma' \vdash \Delta, \Delta', \phi \wedge \psi} \text{ (R}\wedge\text{)} \end{array}}$$

The advantage of these rules is that we do not have to apply weakening explicitly when it is only peripheral to the proof. Note, however, that there is no multiplicative formulation of the backwards rules.

The second version of the cut rule can also be similarly reformulated as

$$\boxed{\frac{\begin{array}{c} \Gamma_1 \vdash \Delta_1 \\ \vdots \\ \Gamma_n \vdash \Delta_n \end{array} \quad \begin{array}{c} \Theta_1 \vdash' \Lambda_1 \\ \vdots \\ \Theta_m \vdash' \Lambda_m \end{array}}{\Gamma_1, \dots, \Gamma_n (\vdash \circ \vdash') \Lambda_1, \dots, \Lambda_m} \text{ (Cut*)}}$$

which, of course, is again subject to the side condition.

Having different formulations of the rules available allows us to use the following general strategy in the next section: When we have to transform one proof into another we will always assume that the given proof uses

the most restrictive formulation of the rules. In the construction of the new proof we can, however, feel free to use the most convenient version of the rules. This enables us to get rid of a lot of the bookkeeping that goes on in these proofs and to focus on the essential features.

2.1.2 A Category of Consequence Relations

We are interested in constructing a category of consequence relations. As a first step we observe that (Cut) preserves consequence relations and that it is associative.

Lemma 2.1.4. *Given consequence relations \vdash from L to M and \vdash' from M to N the sequents $\Gamma (\vdash \circ \vdash') \Delta$ that arise from cuts form a consequence relation.*

Proof. Weakening is obvious because no sequence is altered that is relevant for the cut, and for the same reason (R \perp), (R \vee) and backwards (R \wedge) are trivial.

For forwards (R \wedge) assume we have two cuts:

$$\frac{\Gamma \vdash \alpha \quad \alpha \vdash' \Delta, \phi}{\Gamma (\vdash \circ \vdash') \Delta, \phi} \text{ (Cut)} \qquad \frac{\Gamma \vdash \beta \quad \beta \vdash' \Delta, \psi}{\Gamma (\vdash \circ \vdash') \Delta, \psi} \text{ (Cut)}$$

Then we can directly construct the new proof:

$$\frac{\Gamma \vdash \alpha \quad \frac{\alpha \vdash' \Delta, \phi \quad \beta \vdash' \Delta, \psi}{\alpha, \beta \vdash' \Delta, \phi \wedge \psi} \text{ (R}\wedge\text{)}}{\Gamma (\vdash \circ \vdash') \Delta, \phi \wedge \psi} \text{ (Cut*)}$$

For (R \top) we have the proof

$$\frac{\frac{}{\vdash \top} \text{ (R}\top\text{)} \quad \frac{}{\top \vdash' \top} \text{ (R}\top\text{)}}{(\vdash \circ \vdash') \top} \text{ (Cut)}$$

Since all rules are self-dual as discussed in the paragraph preceding Lemma 2.1.2, it is now clear that $\vdash \circ \vdash'$ is closed under the remaining rules. \square

Lemma 2.1.5. *The composition of consequence relations induced by (Cut) is associative.*

Proof. We can easily rewrite a proof of the form

$$\frac{\frac{\Gamma \vdash \phi \quad \phi \vdash' \psi}{\Gamma (\vdash \circ \vdash') \psi} \text{ (Cut)} \quad \psi \vdash'' \Delta}{\Gamma ((\vdash \circ \vdash') \circ \vdash'') \Delta} \text{ (Cut)}$$

as

$$\frac{\Gamma \vdash \phi \quad \frac{\phi \vdash' \psi \quad \psi \vdash'' \Delta}{\phi (\vdash' \circ \vdash'') \Delta} (\text{Cut})}{\Gamma (\vdash \circ (\vdash' \circ \vdash'')) \Delta.} (\text{Cut})$$

And vice versa. □

A Proof-Theoretic Analysis of Identities

What remains to be done, to make consequence relations into a category, is finding identities. One might be tempted to employ ordinary logical implication between formulae of one world for this. However, this is somewhat against the spirit of our whole setup where we want to suppress purely logical equivalences in order to exhibit the properties of inferences which are, in some sense, subjective or observational. As we have argued, for such inferences it is not necessarily the case that a formula ϕ implies itself. That is, we refuse the identity axioms

$$\overline{\phi \Vdash \phi.}$$

We reserve the symbol \Vdash to represent a consequence relation that has identical source and target algebra L .

On the other hand, Gentzen's original cut rule

$$\boxed{\frac{\Gamma \Vdash \Delta, \phi \quad \phi, \Theta \Vdash \Lambda}{\Gamma, \Theta \Vdash \Delta, \Lambda} (\text{Cut}')} \quad (\text{Cut}')$$

makes sense even in an observational interpretation. The point of introducing yet another cut rule is the following: The rule (Cut) is clearly structurally simpler than (Cut*), and hence it is also easier to check whether a given relation \Vdash is closed under the former. Unfortunately, (Cut) is not very well suited for cut elimination as we will see in Section 2.2. The new rule (Cut') is a compromise that to some extent combines the advantages of the two other cut rules; it is easy as it has only two premises and it is still well-behaved under cut elimination.

We might hope that (Cut') holds for an identity consequence relation. More precisely, if $\Vdash \circ \Vdash = \Vdash$ for a consequence relation on some algebra L then one might expect that \Vdash is closed under Gentzen's cut. As it turns out, this can be shown if consequence relations are assumed to be *interpolative* in the following sense:

Definition 2.1.6. We say that \Vdash has *interpolants* if the following are satisfied:

(L-Int) If $\phi, \Gamma \Vdash \Delta$ then there exists $\phi' \in L$ so that $\phi \Vdash \phi'$ and $\phi', \Gamma \Vdash \Delta$.

(R-Int) If $\Gamma \Vdash \Delta, \phi$ then there exists $\phi' \in L$ so that $\Gamma \Vdash \Delta, \phi'$ and $\phi' \Vdash \phi$.

Of course, if the identity axioms of sequent calculus are adopted then interpolation is trivial. Looking at this from the other end, we can say that interpolation will provide us with some of the consequences of the identity axiom scheme.

Lemma 2.1.7. If \Vdash has interpolants, then $\Vdash \subseteq \Vdash \circ \Vdash$.

Proof. A sequent of the form $\Vdash \Delta$ can be cut with itself using (Cut*) to give $(\Vdash \circ \Vdash) \Delta$ since the side condition is trivial (see also the remark on page 61). If it is of the form $\gamma_1, \dots, \gamma_n \Vdash \Delta$ we use interpolation n times to get $\gamma'_1, \dots, \gamma'_n \Vdash \Delta$ and $\gamma_i \Vdash \gamma'_i$, for $i = 1, \dots, n$. This allows us to form the proof

$$\frac{\begin{array}{c} \gamma_1 \Vdash \gamma'_1 \\ \vdots \\ \gamma_n \Vdash \gamma'_n \end{array} \quad \gamma'_1, \dots, \gamma'_n \Vdash \Delta}{\gamma_1, \dots, \gamma_n (\Vdash \circ \Vdash) \Delta.} \text{ (Cut*)}$$

□

Proposition 2.1.8. A consequence relation \Vdash with interpolants is closed under (Cut'), if and only if $\Vdash \circ \Vdash \subseteq \Vdash$.

Proof. Suppose that $\Vdash \circ \Vdash \subseteq \Vdash$ holds. We want to show that it is closed under (Cut'). To this end, let $\Gamma \Vdash \Delta, \phi$ and $\phi, \Theta \Vdash \Lambda$ be two sequents with $\Delta = \{\delta_1, \dots, \delta_n\}$ and $\Theta = \{\theta_1, \dots, \theta_m\}$. As \Vdash has interpolants there are elements δ'_i and θ'_j such that $\delta'_i \Vdash \delta_i$; $\Gamma \Vdash \delta'_1, \dots, \delta'_n, \phi$; $\theta_j \Vdash \theta'_j$; and $\phi, \theta'_1, \dots, \theta'_m \Vdash \Lambda$. Hence, we can form the cut

$$\frac{\begin{array}{c} \Gamma \Vdash \delta'_1, \dots, \delta'_n, \phi \\ \theta_1 \Vdash \theta'_1 \\ \vdots \\ \theta_m \Vdash \theta'_m \end{array} \quad \begin{array}{c} \delta'_1 \Vdash \delta_1 \\ \vdots \\ \delta'_n \Vdash \delta_n \\ \phi, \theta'_1, \dots, \theta'_m \Vdash \Lambda \end{array}}{\Gamma, \Theta (\Vdash \circ \Vdash) \Delta, \Lambda} \text{ (Cut*)}$$

as the side condition is straightforward to verify. By assumption this implies $\Gamma, \Theta \Vdash \Delta, \Lambda$.

The converse follows immediately from the observation that (Cut) is an instance of (Cut'). □

Corollary 2.1.9. *If \Vdash has interpolants, then it is an idempotent with respect to (Cut) if and only if it is closed under Gentzen's cut rule.*

From this, we take our cue to define the objects of a category.

Definition 2.1.10. $\langle L; \wedge, \vee, \top, \perp; \Vdash \rangle$ is a *continuous sequent calculus* if \Vdash is a consequence relation from L to L such that \Vdash has interpolants and is closed under (Cut), or equivalently (Cut').

The relations \Vdash are in fact idempotents but not identities for all consequence relations. This is not surprising because, as yet, we do not have any axioms that make sure that identities and other consequence relations interact in a sensible way.

Definition 2.1.11. Say that a consequence relation \vdash from L to M is *compatible* with \Vdash_L and \Vdash_M if

$$\Vdash_L \circ \vdash = \vdash = \vdash \circ \Vdash_M$$

We let MLS (for multi-lingual sequents) be the category that has continuous sequent calculi as objects and compatible consequence relations between them as arrows.

The facts that \Vdash_L is self-compatible on both sides, and that composition of compatible consequent relations preserves compatibility are both evident from the definition. From the discussion preceding Lemma 2.1.2 it is clear that MLS is self-dual.

It is also obvious that for any consequence relation $\vdash: L \rightarrow M$ the derived consequence relation $\Vdash_L \circ \vdash \circ \Vdash_M$ is compatible.

The properties (L-Int), (R-Int), and being closed under (Cut') of continuous sequent calculi are inherited by compatible consequence relations:

Proposition 2.1.12. *For a consequence relation \vdash between continuous sequent calculi L and M the following are equivalent:*

1. *The relation \vdash is compatible with \Vdash_L and \Vdash_M .*

2. *It is closed under*

(L-Int') $\phi, \Gamma \vdash \Delta$ implies that there exists $\phi' \in L$ so that $\phi \Vdash_L \phi'$ and $\phi', \Gamma \vdash \Delta$;

(R-Int') $\Gamma \vdash \Delta, \phi$ implies that there exists $\phi' \in L$ so that $\phi' \Vdash_L \phi$ and $\Gamma \vdash \Delta, \phi'$;

(L-Cut) if $\Gamma \Vdash_L \phi$ and $\phi, \Theta \vdash \Lambda$, then $\Gamma, \Theta \vdash \Lambda$; and

(R-Cut) if $\Gamma \vdash \Delta, \phi$ and $\phi \Vdash_M \Lambda$, then $\Gamma \vdash \Delta, \Lambda$.

3. *It satisfies $\Vdash_L \circ \vdash \subseteq \vdash$ and $\vdash \circ \Vdash_M \subseteq \vdash$, and is closed under (L-Int') and (R-Int').*

Proof. If \vdash is compatible with \Vdash_L , then any sequent $\phi, \Gamma \vdash \Delta$ can be written as a cut

$$\frac{\phi, \Gamma \Vdash_L \alpha \quad \alpha \vdash \Delta}{\phi, \Gamma (\Vdash_L \circ \vdash) \Delta} \text{ (Cut)}$$

Because of interpolation in L there exists a $\phi' \in L$ such that $\phi \Vdash_L \phi'$ and $\phi', \Gamma \Vdash_L \alpha$. We cut and get

$$\frac{\phi', \Gamma \Vdash_L \alpha \quad \alpha \vdash \Delta}{\phi', \Gamma \underbrace{(\Vdash_L \circ \vdash)}_{\vdash} \Delta} \text{ (Cut)}$$

which proves (L-Int').

Now, suppose $\Gamma \Vdash_L \phi$ and $\phi, \Theta \vdash \Lambda$, where $\Theta = \{\theta_1, \dots, \theta_m\}$. By the existence of interpolants that we have just proved, there are elements θ'_1 through θ'_m so that $\phi, \theta'_1, \dots, \theta'_m \vdash \Lambda$ and $\theta_i \Vdash_L \theta'_i$. These sequents allow the proof

$$\frac{\begin{array}{c} \Gamma \Vdash_L \phi \\ \theta_1 \Vdash_L \theta'_1 \\ \vdots \\ \theta_m \Vdash_L \theta'_m \end{array} \quad \phi, \theta'_1, \dots, \theta'_m \vdash \Lambda}{\Gamma, \Theta \underbrace{(\Vdash_L \circ \vdash)}_{\vdash} \Lambda} \text{ (Cut*)}$$

The two dual rules are proved analogously. This concludes the proof of (1) \implies (2).

The implication (2) \implies (3) is clear since from (L-Cut) and (R-Cut) we immediately get $\Vdash_L \circ \vdash \subseteq \vdash$ and $\vdash \circ \Vdash_M \subseteq \vdash$.

For (3) \implies (1) we only have to observe that (L-Int') and (R-Int') yield $\vdash \subseteq \Vdash_L \circ \vdash$ and $\vdash \subseteq \vdash \circ \Vdash_M$ —the trivial case, where the relevant side contains the empty sequent, is taken care of by (W), (L \perp) and (R \top). \square

As with (Cut'), the advantage of introducing the new rules (L-Cut) and (R-Cut) will only become clear in Section 2.2.

A closer inspection of the above proof shows that it did not really depend on \vdash being a consequence relation. We only needed (W), (L \perp) and (R \top) to prove that a relation \vdash satisfying the conditions of the proposition is compatible. Such a compatible relation automatically satisfies most rules for consequence relations, since \Vdash_L and \Vdash_M do:

Corollary 2.1.13. *A relation between continuous sequent calculi is a compatible consequence relation if and only if it satisfies (W), (L \perp), (R \top), forwards (L \vee) and (R \wedge), and the rules of Proposition 2.1.12, namely (L-Int'), (R-Int'), (L-Cut) and (R-Cut).*

Proof. The conditions of the corollary are clearly necessary. If we assume conversely that the conditions hold for such a relation \vdash then we know already that it is compatible. That it also satisfies all the remaining logical rules is shown using the pattern of first interpolating, then applying the appropriate rule for \Vdash and then cutting. We consider the only interesting case:

For a sequent $\Gamma \vdash \Delta, \phi, \psi$ interpolation yields the sequents $\Gamma \vdash \Delta, \phi', \psi'$; $\phi' \Vdash \phi$; and $\psi' \Vdash \psi$ enabling us construct the proof:

$$\begin{array}{c}
 \frac{\phi' \Vdash \phi}{\phi' \Vdash \phi, \psi} \text{ (W)} \\
 \frac{\Gamma \vdash \Delta, \phi', \psi'}{\phi' \Vdash \phi \vee \psi} \text{ (RV)} \\
 \frac{\Gamma \vdash \Delta, \phi', \psi' \quad \phi' \Vdash \phi \vee \psi}{\Gamma \vdash \Delta, \phi \vee \psi, \psi'} \text{ (R-Cut)} \\
 \frac{\psi' \Vdash \psi}{\psi' \Vdash \phi, \psi} \text{ (W)} \\
 \frac{\psi' \Vdash \phi, \psi}{\psi' \Vdash \phi \vee \psi} \text{ (RV)} \\
 \frac{\Gamma \vdash \Delta, \phi \vee \psi, \psi' \quad \psi' \Vdash \phi \vee \psi}{\Gamma \vdash \Delta, \phi \vee \psi, \phi \vee \psi} \text{ (R-Cut)}
 \end{array}$$

□

It is worthwhile to note that all definitions and theorems up to this point still make sense if one does not allow the application of the logical rules from the lower sequent to the upper sequent. We will come back to the significance and problems of backwards rules in the next section where we discuss cut elimination and later in Section 3.1.2 when we study semantics.

Remark. *A different perspective on the definition of MLS is given by the following. We consider consequence relations and the composition given by (Cut). As we observed before this is not quite a category, but only because it fails to have identities. We now split the class of idempotents that have interpolants, a technique which is well-known in category theory, [FS90, 1.28]. It works even if the original ‘category’ fails to have identities. The resulting category in our case is precisely MLS.*

Alternatively, we can build a proper category before we split the idempotents: We can restrict the ‘internal’ logic to situations where a proposition ϕ does imply itself. As the identity morphism on a $(2, 2, 0, 0)$ -algebra we can then take the smallest consequence relation generated by the identity rules, which will yield precisely the classically valid sequents of the system. Compatibility is then not an issue and we immediately obtain a (self-dual) category RMLS. Now we can again obtain MLS by splitting idempotents in RMLS that have interpolants.

2.2 Cut Elimination

The famous Cut Elimination Theorem of Gentzen [Gen34] states that every valid sequent in the sequent calculus can be derived without employing the cut rule. Sequents in our setting, however, are not about absolute validity but about derivability of sequents from assumed sequents. The analogous theorem for this situation says that in every such derivation cuts between arbitrary sequents can be eliminated in favour of cuts between assumed sequents. We will exhibit a similar result which applies to the rule (Cut*).

2.2.1 Simple Elimination

If R is any relation between finite sets of elements of $(2, 2, 0, 0)$ -algebras let R^+ denote the smallest such relation which contains R and is closed under application of *forwards* rules, including weakening. In this definition we do not close R under backwards rules like

$$\frac{\phi \wedge \psi, \Gamma \vdash \Delta}{\phi, \psi, \Gamma \vdash \Delta}$$

as they eliminate connectives and hence decrease the complexity of a sequent.

For algebras themselves we say that $B \subseteq L$ is a *generating set* if the smallest subalgebra B^+ of L containing B is L itself. This generation process can also be described by finitary rules:

$$\frac{\phi \in B}{\phi \in B^+} \quad \frac{\phi, \psi \in B^+}{\phi \wedge \psi \in B^+} \quad \frac{\phi, \psi \in B^+}{\phi \vee \psi \in B^+} \quad \frac{}{\top, \perp \in B^+}$$

That is to say, the elements of B^+ are precisely those whose membership in B^+ can be derived using these rules.

If $B \subseteq L$ and $C \subseteq M$ and if \vdash is a consequence relation from L to M , write $\vdash \Big|_B^C$ to denote the restriction of \vdash to sequents made up entirely from the respective subsets, $\vdash \cap (\mathfrak{P}_{\text{fin}}(B) \times \mathfrak{P}_{\text{fin}}(C))$.

For a fixed set of generators B we define the *rank* $r(\phi)$ of a formula ϕ to be the minimum height of a derivation for $\phi \in B^+$ using the rules given above, setting $r(\phi) = 0$ for $\phi \in B$. For a finite set $\Gamma \subseteq L$, let $r(\Gamma)$ be the maximum rank of any member of Γ . Also let $rc(\Gamma)$ denote the number of elements of Γ of maximum rank. For finite Γ , define the *grade* $g(\Gamma)$ as the pair $g(\Gamma) := \langle r(\Gamma), rc(\Gamma) \rangle$, with the lexicographical order on the set of all grades to make it into a well-order.

We come to the first important lemma relating sets of generators for algebras and freely generated consequence relations. Note that although the definition of $(\cdot)^+$ excludes backwards rules we nonetheless require consequence relations to be closed under these rules. In fact, many of the results

on cut elimination depend on them. As we will see, for example, in the proof of the following lemma they are essential for many induction arguments to go through.

Lemma 2.2.1. *Let \vdash be a consequence relation from L to M . If B and C are generating sets of L and M , respectively, then $\vdash = (\vdash \mid_B^C)^+$.*

Proof. We obviously have $(\vdash \mid_B^C)^+ \subseteq \vdash$. For the interesting containment we have to verify that $\Gamma \vdash \Delta$ implies $\tilde{\Gamma} (\vdash \mid_B^C)^+ \Delta$. We do so by induction on the grades $g(\Gamma)$ and $g(\Delta)$: The basis is immediate, and if $r(\Gamma) > 0$, then we can write Γ as $\phi, \tilde{\Gamma}$ for some token ϕ such that $r(\phi) = r(\Gamma)$. Now, ϕ is of the form $\top, \perp, \psi \wedge \chi$ or $\psi \vee \chi$ where ψ and χ are of lower rank.

Let us look at the case $\phi = \psi \vee \chi$ as an example, where this is the decomposition giving rise to the rank of ϕ . We get

$$\frac{\psi \vee \chi, \tilde{\Gamma} \vdash \Delta}{\psi, \tilde{\Gamma} \vdash \Delta \quad \chi, \tilde{\Gamma} \vdash \Delta} (\text{LV})$$

by using backwards (LV). As for the grades involved we have

$$g(\psi, \tilde{\Gamma}), g(\chi, \tilde{\Gamma}) < g(\psi \vee \chi, \tilde{\Gamma}) = g(\Gamma)$$

and so the induction hypothesis yields

$$\frac{\psi, \tilde{\Gamma} (\vdash \mid_B^C)^+ \Delta \quad \chi, \tilde{\Gamma} (\vdash \mid_B^C)^+ \Delta}{\psi \vee \chi, \tilde{\Gamma} (\vdash \mid_B^C)^+ \Delta} (\text{LV})$$

The other cases (and their duals in Δ) are proved similarly. □

The lemma shows that we can restrict our attention to the behaviour of consequence relations on generators for the algebras involved. Note that this is, in particular, true for relations \Vdash that define the objects of our category MLS. In the following we examine how far this idea can be pushed.

We start with the composition of consequence relations. Suppose R and S are binary relations between finite sets of formulae. We write $R \circ S$ for the set of sequents that can be derived by using (Cut*) which, of course, includes all sequents derivable by (Cut).

In the context of cut elimination it actually matters which formulation of cut we use: If we go back to the proof of Proposition 2.1.3 we see that cuts using (Cut*) can be reformulated as (Cut)-cuts by using logical forward rules. Unfortunately, this operation increases the grades of the sequents involved and hence cannot be used in an induction proof based on these grades. If we had formulated our logical system only in terms of the simple cut rule (Cut) most of the following results would not hold.

Lemma 2.2.2. *If C is a generating set for M and \vdash, \vdash' are consequence relations from L to M , and M to N , respectively, then*

$$\vdash \circ \vdash' = \vdash|_C^C \circ \vdash|_C,$$

where the composition on the right hand side stands for the sequents that can be derived by either cut rule.

Proof. One containment is obvious. For the other one the proof is by induction on the maximal grade of Δ 's and Θ 's involved in a cut

$$\frac{\begin{array}{c} \Gamma \vdash \Delta_1 \\ \vdots \\ \Gamma \vdash \Delta_n \end{array} \quad \begin{array}{c} \Theta_1 \vdash' \Lambda \\ \vdots \\ \Theta_m \vdash' \Lambda \end{array}}{\Gamma (\vdash \circ \vdash') \Lambda} \text{ (Cut*)}$$

and the number of sequents with the maximal grade. If all Δ 's and Θ 's have rank 0, i.e. they contain only formulae from C , then we clearly have $\Gamma (\vdash|_C^C \circ \vdash|_C) \Lambda$. For the induction step we look at the sequent Δ_i or Θ_j with the maximal grade. Without loss of generality we can assume that $\Delta_n = \tilde{\Delta}_n, \phi$ is that sequent, where ϕ is a formula of maximal rank. We distinguish cases on the basis of the decomposition determining the rank of ϕ .

If $\phi = \psi \vee \chi$ we replace the sequent $\Gamma \vdash \tilde{\Delta}_n, \psi \vee \chi$ by $\Gamma \vdash \tilde{\Delta}_n, \psi, \chi$ and consider the choice functions that now arise. The only ones that might be a problem are those that include ψ or χ . Suppose $f \in \prod_{i=1}^{n-1} \Delta_i$ is a choice function such that $\{f_1, \dots, f_{n-1}, \psi \vee \chi\}$ covers a Θ_j but $\{f_1, \dots, f_{n-1}\}$ does not. Then Θ_j must be of the form $\psi \vee \chi, \tilde{\Theta}_j$. We deduce

$$\frac{\psi \vee \chi, \tilde{\Theta}_j \vdash' \Lambda}{\psi, \tilde{\Theta}_j \vdash' \Lambda \quad \chi, \tilde{\Theta}_j \vdash' \Lambda} \text{ (LV)}$$

and adding these two sequents ensures that the side condition is satisfied whether we extend f by ψ or χ . Note that we have to add these sequents rather than just replacing the old one since we cannot be sure that none of the f_i is $\psi \vee \chi$. However, the new sequents that arise in this way by 'splitting' $\psi \vee \chi$ where it appears in a Θ do not have the maximal grade. Hence, having already decreased the maximal grade or the number of its occurrences by substituting $\Gamma \vdash \tilde{\Delta}_n, \psi, \chi$ for $\Gamma \vdash \tilde{\Delta}_n, \psi \vee \chi$, we can simply add all these sequents and apply the induction hypothesis to get $\Gamma (\vdash|_C^C \circ \vdash|_C) \Lambda$.

For $\phi = \psi \wedge \chi$, analogously, we replace $\Gamma \vdash \tilde{\Delta}_n, \psi \wedge \chi$ by $\Gamma \vdash \tilde{\Delta}_n, \psi, \chi$ and add all results of applying (L \wedge) to sequents of the form $\psi \wedge \chi, \tilde{\Theta}_j \vdash' \Lambda$.

If $\phi = \perp$ we apply (R \perp) to get the new sequent $\Gamma \vdash \tilde{\Delta}_n$. This reduces the number of choice functions, and the side condition is valid without making any modifications to sequents on the right.

Finally we have to consider the case $\phi = \top$. Here we delete all sequents $\Gamma \vdash \tilde{\Delta}_i, \top$ and replace all sequents $\top, \tilde{\Theta} \vdash' \Lambda$ by $\tilde{\Theta} \vdash' \Lambda$. Any choice function \tilde{f} for the remaining Δ 's can be extended to one $f \in \prod_{i=1}^n \Delta_i$ by picking \top for the other i 's. Since $\{f_1, \dots, f_n\}$ covers one Θ_j the image of f contains $\Theta_j \setminus \{\top\}$. Note that in the case that there is no longer a sequent on the left there was a choice function picking only \top . This implies that a sequent on the right is of the form $\top \vdash' \Lambda$ or $\vdash' \Lambda$, the later being derivable from the former using (L \top). This is precisely the degenerate case of the side condition discussed in the remark on page 61.

The dual conditions, where one of the Θ 's is the sequent with the maximal rank, are argued the same. \square

Theorem 2.2.3 (Cut Elimination). *Let $B \subseteq L$, $C \subseteq M$ and $D \subseteq N$ be sets of generators. Then for any consequence relations \vdash and \vdash' between L , M and N it is the case that*

$$\vdash \circ \vdash' = (\vdash|_B^C \circ \vdash'|_C^D)^+.$$

Proof. Because of the previous two lemmata we immediately get

$$\vdash \circ \vdash' = \vdash|_C^C \circ \vdash'|_C^C = ((\vdash|_C^C \circ \vdash'|_C^C)|_B^D)^+.$$

The set $(\vdash|_C^C \circ \vdash'|_C^C)|_B^D$ contains precisely the sequents that arise as cuts:

$$\frac{\begin{array}{cc} \Gamma \vdash \Delta_1 & \Theta_1 \vdash' \Lambda \\ \vdots & \vdots \\ \Gamma \vdash \Delta_n & \Theta_m \vdash' \Lambda \end{array}}{\Gamma (\vdash \circ \vdash') \Lambda} \text{ (Cut*)}$$

where Γ is a subset of B , the Δ_i and Θ_j are subsets of C , and Λ is a subset of D . In other words $(\vdash|_C^C \circ \vdash'|_C^C)|_B^D = \vdash|_B^C \circ \vdash'|_C^D$ which together with the first equation proves the theorem. \square

2.2.2 Construction of Consequence Relations

So far we have looked at given consequence relations. Now, we want to use similar techniques to construct consequence relations by specifying them on generators. The question arises, what is needed so we can guarantee that for a given such relation R the resulting R^+ is actually a morphism in our category. As it turns out, it is very difficult to derive general rules in this direction. The problem is that originally we allowed formulae to be drawn from general $(2, 2, 0, 0)$ -algebras. The situation becomes much more manageable if we restrict ourselves to term algebras over a given set of generators. In logic, formulae are usually freely defined, so this is a quite normal restriction. As we will see in the following discussion, it is not even a serious one.

Term Algebras

For any set B let $T(B)$ be the term algebra for the signature $(\wedge, \vee, \top, \perp)$. We are now going to show that in MLS every object is isomorphic to one where the underlying algebra is a term algebra. It is well known from universal algebra that every algebra L is a quotient of the term algebra $T(L)$ [MT92, Corollary 5.1.7]. The canonical quotient map takes a term from $T(L)$ and evaluates it in L . Writing $[\cdot]: T(L) \rightarrow L$ for this map we can define

$$\gamma_1, \dots, \gamma_n \Vdash_{T(L)} \delta_1, \dots, \delta_m \quad :\iff \quad [\gamma_1], \dots, [\gamma_n] \Vdash_L [\delta_1], \dots, [\delta_m].$$

We usually abbreviate expressions of the form $[\gamma_1], \dots, [\gamma_n]$ by $[\Gamma]$.

Lemma 2.2.4. *If (L, \Vdash_L) is a continuous sequent calculus, then $(T(L), \Vdash_{T(L)})$ is as well.*

Proof. To prove that $\Vdash_{T(L)}$ is a consequence relation we have to check each of the rules of the calculus. This can be done by using the fact that \Vdash_L satisfies these rules. We do one case to give the flavour of the argument: Let us suppose

$$\Gamma \Vdash_{T(L)} \Delta, \phi \quad \Gamma \Vdash_{T(L)} \Delta, \psi.$$

which, by definition, is equivalent to

$$[\Gamma] \Vdash_L [\Delta], [\phi] \quad [\Gamma] \Vdash_L [\Delta], [\psi].$$

We now apply the relevant rule in L

$$\frac{[\Gamma] \Vdash_L [\Delta], [\phi] \quad [\Gamma] \Vdash_L [\Delta], [\psi]}{[\Gamma] \Vdash_L [\Delta], \underbrace{[\phi] \wedge [\psi]}_{[\phi \wedge \psi]}}$$

and the result is defined to be

$$\Gamma \Vdash_{T(L)} \Delta, \phi \wedge \psi.$$

Hence $\Vdash_{T(L)}$ is a consequence relation. The proof that it is closed under (Cut) follows exactly the same pattern.

The argument for the interpolation axioms (L-Int) and (L-Int) is also very similar: Suppose $\phi, \Gamma \Vdash_{T(L)} \Delta$, or equivalently $[\phi], [\Gamma] \Vdash_L [\Delta]$. Interpolation in L yields an element $\psi \in L$ such that $[\phi] \Vdash_L \psi$ and $\psi, [\Gamma] \Vdash_L [\Delta]$. Now, ψ is also a term in $T(L)$ and it satisfies $[\psi] = \psi$. Thus we get $\phi \Vdash_{T(L)} \psi$ and $\psi, \Gamma \Vdash_{T(L)} \Delta$. \square

Proposition 2.2.5. *The continuous sequent calculi (L, \Vdash_L) and $(T(L), \Vdash_{T(L)})$ are isomorphic in MLS.*

Proof. Given the definition of $\Vdash_{T(L)}$ it is easy to come up with the isomorphisms between L and $T(L)$. We define $\vdash: L \rightarrow T(L)$ and $\vdash': T(L) \rightarrow L$ by

$$\Gamma \vdash \Delta \quad :\Leftrightarrow \quad \Gamma \Vdash_L [\Delta] \quad \text{and} \quad \Gamma \vdash' \Delta \quad :\Leftrightarrow \quad [\Gamma] \Vdash_L \Delta.$$

To prove that \vdash and \vdash' are compatible consequence relation we have to check the rules listed in Corollary 2.1.13, namely (W), (L \perp), (R \top), forwards (LV) and (R \wedge), (L-Int'), (R-Int'), (L-Cut) and (R-Cut).

Apart from the last two rules we can use essentially the same proof as in the previous lemma. And the cut rules require only a minor new ingredient: Consider two sequents $\Gamma \vdash \Delta, \phi$ and $\phi \Vdash_{T(L)} \Lambda$. We infer $\Gamma \Vdash_L [\Delta], [\phi]$ and $[\phi] \Vdash_L [\Lambda]$, and because of Proposition 2.1.8 we can form

$$\frac{\Gamma \Vdash_L [\Delta], [\phi] \quad [\phi] \Vdash_L [\Lambda]}{\Gamma \Vdash_L [\Delta], [\Lambda]} \text{ (Cut')}$$

which shows $\Gamma \vdash \Delta, \Lambda$. The other cases involving cut rules are analogous.

It remains to show that \vdash and \vdash' are mutually inverse. The containment $(\vdash' \circ \vdash) \subseteq \Vdash_{T(L)}$ follows from the fact that cuts of the form

$$\frac{\Gamma \vdash' \phi \quad \phi \vdash \Delta}{\Gamma (\vdash' \circ \vdash) \Delta} \text{ (Cut)}$$

correspond directly to cuts

$$\frac{[\Gamma] \Vdash_L \phi \quad \phi \Vdash_L [\Delta]}{[\Gamma] \underbrace{(\Vdash_L \circ \Vdash_L)}_{\Vdash_L} [\Delta]} \text{ (Cut)}$$

and the last sequent, by definition, means $\Gamma \Vdash_{T(L)} \Delta$. The other containment follows by the same argument since $\Vdash_L = \Vdash_L \circ \Vdash_L$.

The proof of $(\vdash \circ \vdash') = \Vdash_L$ is practically the same. \square

The proposition shows that we can restrict our study of continuous sequent calculi to the purely syntactic ones, i.e. the term algebras. We can express this categorically by saying that the full subcategory of MLS whose objects are algebras of the form $T(B)$ is equivalent to MLS.

We now return to the main thrust of this section. Our first lemma shows that $(\cdot)^+$ yields consequence relations:

Lemma 2.2.6. *Let R be a relation between finite subsets of B and C . Then R^+ is a consequence relation from $T(B)$ to $T(C)$.*

The restriction $R^+ \upharpoonright_B^C$ is just the closure of R under weakening with formulae from B and C , respectively.

Proof. The algebras $T(B)$ and $T(C)$ are freely generated. So every sequent in R^+ is derived from the sequent or the sequents required by the corresponding backwards rule. The only exception are formulae that have been introduced by weakening. In this case an inspection of the rules shows that the sequents resulting from a backwards application of a rule can then likewise be derived by weakening.

All rules other than (W) introduce composite formulae which shows the second claim. \square

A condition that is needed for compatibility and identities in MLS is interpolation. We prove a slightly technical lemma that is strong enough so we can use it to show both, compatibility of generated consequence relations and interpolation for candidates of continuous sequent calculi.

Lemma 2.2.7. *Let $R \subseteq \mathfrak{F}_{\text{fin}}(B) \times \mathfrak{F}_{\text{fin}}(B)$ and $S \subseteq \mathfrak{F}_{\text{fin}}(B) \times \mathfrak{F}_{\text{fin}}(C)$ be relations such that $\phi, \Gamma (S) \Delta$ implies the existence of an interpolant $\phi' \in T(B)$ satisfying $\phi (R^+) \phi'$ and $\phi', \Gamma (S^+) \Delta$. Then for all S^+ -sequents there are R^+ -interpolants: I.e. for all $\phi, \Gamma (S^+) \Delta$ there is a $\phi' \in T(B)$ such that $\phi (R^+) \phi'$ and $\phi', \Gamma (S^+) \Delta$.*

Proof. The proof proceeds by induction the rank of ϕ and the height of the derivation of $\phi, \Gamma (S^+) \Delta$, the former taking precedence. The base case concerning sequents that contain only elements of B and C is trivial, and so are the cases where ϕ is not the principal formula and the rule depends on only one sequent.

For rules taking two sequents, consider the derivation

$$\frac{\phi, \Gamma (S^+) \Delta, \psi \quad \phi, \Gamma (S^+) \Delta, \chi}{\phi, \Gamma (S^+) \Delta, \psi \wedge \chi} \text{ (R}\wedge\text{)}$$

as an example. The induction hypothesis yields two formulae ϕ' and ϕ'' such that $\phi (R^+) \phi'$; $\phi (R^+) \phi''$; $\phi', \Gamma (S^+) \Delta, \psi$ and $\phi'', \Gamma (S^+) \Delta, \chi$. From this we can deduce

$$\frac{\phi (R^+) \phi' \quad \phi (R^+) \phi''}{\phi (R^+) \phi' \wedge \phi''} \text{ (R}\wedge\text{)}$$

and

$$\frac{\frac{\frac{\phi', \Gamma (S^+) \Delta, \psi}{\phi', \phi'', \Gamma (S^+) \Delta, \psi} \text{ (W)}}{\phi' \wedge \phi'', \Gamma (S^+) \Delta, \psi} \text{ (L}\wedge\text{)}}{\phi' \wedge \phi'', \Gamma (S^+) \Delta, \psi \wedge \chi} \text{ (R}\wedge\text{)} \quad \frac{\frac{\frac{\phi'', \Gamma (S^+) \Delta, \chi}{\phi', \phi'', \Gamma (S^+) \Delta, \chi} \text{ (W)}}{\phi' \wedge \phi'', \Gamma (S^+) \Delta, \chi} \text{ (L}\wedge\text{)}}{\phi' \wedge \phi'', \Gamma (S^+) \Delta, \psi \wedge \chi} \text{ (R}\wedge\text{)}$$

For the rules actually introducing ϕ , the cases for the constants \top , \perp and weakening are trivial because we can interpolate using either \top or \perp .

Next we consider the case that the principal formula is $\phi = \psi \wedge \chi$. Given the proof

$$\frac{\psi, \chi, \Gamma (S^+) \Delta}{\psi \wedge \chi, \Gamma (S^+) \Delta} (\text{L}\wedge)$$

we interpolate twice using the induction hypothesis and get ψ', χ' such that $\psi (R^+) \psi'; \chi (R^+) \chi'$ and $\psi', \chi', \Gamma (S^+) \Delta$. It is here that we actually need that the induction is not only on the derivation of the sequent but also on the rank of ϕ . The first interpolation may well increase the height of the derivation, but as we have already reduced the rank of the formula this is not a problem. From these sequents we construct the two derivations

$$\frac{\frac{\psi (R^+) \psi' \quad \chi (R^+) \chi'}{\psi, \chi (R^+) \psi' \wedge \chi'} (\text{R}\wedge)}{\psi \wedge \chi (R^+) \psi' \wedge \chi'} (\text{L}\wedge)$$

and

$$\frac{\psi', \chi', \Gamma (S^+) \Delta}{\psi' \wedge \chi', \Gamma (S^+) \Delta} (\text{L}\wedge)$$

Mutatis mutandis, the same argument also proves the last case (LV). \square

To construct continuous sequent calculi and to show compatibility of constructed consequence relations we have to reconsider cut elimination. Looking at the generators of consequence relations, the original rule (Cut*) and the multiplicative one with implicit weakening are not interchangeable unless the relations in question is already closed under (W). In the following if we refer to the rule (Cut*) or the composition defined in terms of it we understand this in the stronger sense of the multiplicative formulation of (Cut*) as given on page 63.

The only reason that we cannot directly use Theorem 2.2.3 is that there we started from a consequence relation, whereas now we want to begin with a relation R that generates one. Lemma 2.2.6 tells us that R contains almost all sequents of $R^+ \upharpoonright_B^C$; the ones that are missing can be derived by weakening. Hence, we have to study how weakening and (Cut*) interact. Consider a cut

$$\frac{\begin{array}{c} \Gamma', \Gamma \vdash \Delta_1, \Delta'_1 \\ \vdots \\ \Gamma', \Gamma \vdash \Delta_n, \Delta'_n \end{array} \quad \begin{array}{c} \Theta'_1, \Theta_1 \vdash' \Lambda, \Lambda' \\ \vdots \\ \Theta'_m, \Theta_m \vdash' \Lambda, \Lambda' \end{array}}{\Gamma', \Gamma (\vdash \circ \vdash') \Lambda, \Lambda'} (\text{Cut}^*)$$

where all the primed formulae were added by weakening. Then the proof

$$\frac{\frac{\Gamma \vdash \Delta_1 \quad \Theta_1 \vdash' \Lambda}{\vdots} \quad \frac{\Gamma \vdash \Delta_n \quad \Theta_m \vdash' \Lambda}{\Gamma (\vdash \circ \vdash') \Lambda} (\text{Cut}^*)}{\Gamma', \Gamma (\vdash \circ \vdash') \Lambda, \Lambda'} (\text{W})$$

is also valid: It is easy to verify the side condition in the formulation of Lemma 2.1.2 as before weakening there are less choice functions on both sides.

Proposition 2.2.8. *For a relation R between finite subsets of B and C and a relation S between finite subsets of C and D we have $(R^+) \circ (S^+) = (R \circ S)^+$.*

Proof. Theorem 2.2.3 yields $(R^+) \circ (S^+) = ((R^+) |^C_B \circ (S^+) |^D_C)^+$. The upshot of the discussion preceding this proposition is that the only difference between $(R^+) |^C_B \circ (S^+) |^D_C$ and $R \circ S$ are sequents that are derivable from the latter by weakening. Hence we conclude

$$(R^+) \circ (S^+) = ((R^+) |^C_B \circ (S^+) |^D_C)^+ = (R \circ S)^+.$$

□

Corollary 2.2.9. *If a binary relation R on finite subsets of B is closed under (Cut^*) then so is R^+ .*

If we want to construct continuous sequent calculi, then it is sometimes more convenient to use (Cut') . The following lemma shows that we can do so.

Lemma 2.2.10. *Let R be binary relation on finite subsets of a set B and R' its closure under weakening with formulae from B . If R is closed under (Cut') then R' is closed under (Cut^*) .*

Proof. We show that R' is closed under (Cut^*) by rewriting cuts using this rule in terms of (Cut') . But first we observe that R' is also closed under (Cut') : This follows from the fact that every such cut over a formula introduced by weakening can be entirely replaced by weakening.

Now, we prove by induction on n that if $\Delta_1, \dots, \Delta_n$ and $\Theta_1, \dots, \Theta_m$ satisfy the side condition, then the sequents $\Gamma (R') \Delta_1 \cdots \Gamma (R') \Delta_n$ and $\Theta_1, \Gamma (R') \Lambda \cdots \Theta_m, \Gamma (R') \Lambda$ imply $\Gamma (R') \Lambda$. Note that this immediately shows that R' is closed under cuts of the form

$$\frac{\frac{\Gamma (R') \Delta_1 \quad \Theta_1 (R') \Lambda}{\vdots} \quad \frac{\Gamma (R') \Delta_n \quad \Theta_m (R') \Lambda}{\Gamma (R') \Lambda}}{\Gamma (R') \Lambda}$$

as we can weaken the sequents on the right with Γ .

For $n = 0$ the side condition boils down to one Θ_j being empty and we get $\Gamma (R') \Delta$ by weakening.

In the induction step we ‘eliminate’ the sequent $\Gamma (R') \Delta_{n+1}$. Given any choice function $f \in \prod_{i=1}^n \Delta_i$ and an element $\phi \in \Delta_{n+1}$ there is a j such that $\{f_1, \dots, f_n, \phi\} \supseteq \Theta_j$. Cutting $\Gamma (R') \Delta_{n+1}$ with the sequent $\Theta_j, \Gamma (R') \Delta$ —weakening it to $\Theta_j, \phi, \Gamma (R') \Delta$ if ϕ does not occur in Θ_j —produces a sequent

$$\Gamma, \Theta_j \setminus \{\phi\} (R') \Delta_{n+1} \setminus \{\phi\}, \Delta.$$

The important observation is now that $\Theta_j \setminus \{\phi\} \subseteq \{f_1, \dots, f_n\}$. We pick the next formula $\psi \in \Delta_{n+1}$ and cut the sequent we just generated with the corresponding $\Theta_{j'}, \Gamma (R') \Delta$ to get

$$\Gamma, \Theta_j \setminus \{\phi\}, \Theta_{j'} \setminus \{\psi\} (R') \Delta_{n+1} \setminus \{\phi, \psi\}, \Delta.$$

We iterate this procedure of cutting the resulting sequents with the appropriate $\Theta_k, \Gamma (R') \Delta$ and finally get a sequent of the form

$$\Gamma, \Xi_f (R') \Delta$$

where $\Xi \subseteq \{f_1, \dots, f_n\}$.

Repeating this process for all such choice functions f yields sequents such that $\Delta_1, \dots, \Delta_n$ and $(\Xi_f)_{f \in \prod \Delta_i}$ satisfy the side condition and hence the induction hypothesis. \square

The proof of the lemma essentially shows that (Cut^*) is derivable from (Cut') and (W) without any other rules. The statement of the lemma is reminiscent of one direction of Proposition 2.1.8. There the proof was much simpler since we could use (Cut) . As the rewriting of (Cut^*) in terms of (Cut) uses logical rules we cannot use it in the current context.

Corollary 2.2.11. *If a binary relation R on finite sets is closed under (Cut') then R^+ is closed under (Cut^*) .*

We can put all these results together to get the central theorem about the construction of continuous sequent calculi:

Theorem 2.2.12. *Let R be a binary relation R on finite subsets of a set B that is either closed under (Cut^*) or (Cut') . If it also satisfies the condition that for all $\phi, \Gamma (R) \Delta$ there is a $\phi' \in T(B)$ such that $\phi', \Gamma (R^+) \Delta$ and $\phi (R^+) \phi'$, and dually for interpolation on the right, then R^+ makes $T(B)$ into a continuous sequent calculus.*

Proof. The relation R^+ is a consequence relation by Lemma 2.2.6. From Lemma 2.2.7 and its dual we infer the existence of interpolants. That it is an idempotent then follows from Corollary 2.2.9 or Corollary 2.2.11 together with Lemma 2.1.7. \square

An analogous result for the generation of arbitrary morphisms, rather than just identities, is now also easy to prove.

Theorem 2.2.13. *Let $T(B)$ and $T(C)$ be continuous sequent calculi and R a binary relation between finite subsets of B and C that is closed under (Cut^*) with respect to generators of $\Vdash_{T(B)}$ and $\Vdash_{T(C)}$. If, moreover, for all $\phi, \Gamma (R) \Delta$ there is a $\phi' \in T(B)$ such that $\phi', \Gamma (R^+) \Delta$ and $\phi \Vdash_{T(B)} \phi'$, and the dual interpolation on the right, then R^+ is a compatible consequence relation from $T(B)$ to $T(C)$.*

Proof. Because of Lemma 2.2.6, R^+ is a consequence relation, and Proposition 2.2.8 implies that R^+ is closed under cuts with \Vdash . The remaining conditions of Proposition 2.1.12 are consequences of Lemma 2.2.7. \square

In applications it is quite common that we want to describe a consequence relation between $T(B)$ and $T(C)$ by referring to arbitrary formulae and not just the generators, i.e. elements from B and C . The following lemma explores when we can do this without any additional overhead:

Lemma 2.2.14. *Let R be a relation between finite subsets of $T(B)$ and $T(C)$ that is closed under the backwards rules. Then R^+ is a consequence relation, and moreover it is equal to $(R \upharpoonright_B^C)^+$.*

Proof. Since R is closed under backward rules and $T(B)$ as well as $T(C)$ are freely generated we have $R \subseteq (R \upharpoonright_B^C)^+$. The claims of the lemma are an immediate consequence of this observation. \square

Suppose R is an ‘over-specified’ relation as in this lemma. If we look at the formulation of Theorem 2.2.12 and 2.2.13, we see that if R satisfies the premises of these theorems—apart from being over-specified—then so does its restriction to the generators.

There is another condition that we can relax, namely being closed under the appropriate cut rules. In the applications it often happens that the result of a cut does not quite lie in R but can be derived from such a sequent by weakening. Adding such sequents does not create new problems for interpolation. Moreover, the discussion before Proposition 2.2.8 shows that this does not introduce essentially new cuts, either. Hence, we get slightly more liberal versions of the above theorems:

Corollary 2.2.15. *Let R be a binary relation on finite subsets of a freely generated $(2, 2, 0, 0)$ -algebra that is closed under backwards rules. If R satisfies the interpolation conditions of Theorem 2.2.12 and is closed under (Cut^*) or $(\text{Cut}!)$ up to weakening, then R^+ is a continuous sequent calculus.*

Corollary 2.2.16. *Let R be a binary relation between finite subsets of continuous sequent calculi $T(B)$ and $T(C)$ which is closed under backwards rules. If it satisfies the interpolation conditions of Theorem 2.2.13 and is closed under (Cut^*) with respect to generators of $\Vdash_{T(B)}$ and $\Vdash_{T(C)}$ up to weakening, then R^+ is a compatible consequence relation from $T(B)$ to $T(C)$.*

2.2.3 Coproducts and Products

As an example application of the results of the previous section we construct coproducts in MLS. We do this in considerable detail to show what is involved in performing a construction like this in a purely syntactic fashion.

Let L and M be continuous sequent calculi whose underlying algebras are freely generated. We take the disjoint union $\{0\} \times L \cup \{1\} \times M$ and let $L + M$ be the algebra that is freely generated by it. We already know that we only have to come up with a suitable relation on the basis to define the continuous sequent calculus on $L + M$. For this, we specify the rules:

$$\boxed{\begin{array}{c} \frac{\Gamma \Vdash_L \Delta}{\{0\} \times \Gamma \Vdash_{L+M} \{0\} \times \Delta} \quad (\Gamma \neq \emptyset) \qquad \frac{\Gamma \Vdash_M \Delta}{\{1\} \times \Gamma \Vdash_{L+M} \{1\} \times \Delta} \quad (\Gamma \neq \emptyset) \\ \\ \frac{\Vdash_L \Gamma \quad \Vdash_M \Delta}{\Vdash_{L+M} \{0\} \times \Gamma, \{1\} \times \Delta} \qquad \frac{}{\langle 0, \phi \rangle, \langle 1, \psi \rangle \Vdash_{L+M}} \end{array}}$$

To understand these rules it might be helpful to jump ahead a bit and to consider them semantically: In the next chapter we will see that coproducts in MLS correspond to taking the disjoint union of spaces. Let us suppose L corresponds to a space X and M to Y . In these terms the restriction $\Gamma \neq \emptyset$ in the first rule says, in effect, that the embedding of X in $X + Y$ is only a part of the latter; and analogously for the second rule. The two other rules can be read as saying that the union of the embedding of X and Y is all of $X + Y$ and that the intersection of any part of X with any part Y is the empty set.

Lemma 2.2.17. *The relation $(\Vdash_{L+M})^+$ is a continuous sequent calculus.*

Proof. Clearly, the relation \Vdash_{L+M} has interpolants precisely because \Vdash_L and \Vdash_M do. It is also closed under Cut' as is readily checked by considering the different cuts that arise. Let us verify one case explicitly to give the flavour of the argument: Consider the proof

$$\frac{\frac{\Vdash_L \phi, \Gamma \quad \Vdash_M \Delta}{\Vdash_{L+M} \langle 0, \phi \rangle, \{0\} \times \Gamma, \{1\} \times \Delta} \quad \frac{\phi, \Theta \Vdash_L \Lambda}{\langle 0, \phi \rangle, \{0\} \times \Theta \Vdash_{L+M} \{0\} \times \Lambda}}{\{0\} \times \Theta \Vdash_{L+M} \{0\} \times \Gamma, \{0\} \times \Lambda, \{1\} \times \Delta.} \quad (\text{Cut}')$$

If Θ is not empty we can construct

$$\frac{\frac{\Vdash_L \phi, \Gamma \quad \phi, \Theta \Vdash_L \Lambda}{\Theta \Vdash_L \Gamma, \Lambda} \quad (\text{Cut}')}{\frac{\{0\} \times \Theta \Vdash_{L+M} \{0\} \times \Gamma, \{0\} \times \Lambda}{\{0\} \times \Theta \Vdash_{L+M} \{0\} \times \Gamma, \{0\} \times \Lambda, \{1\} \times \Delta} \quad (\text{W})}$$

otherwise we simply use a different rule and get

$$\frac{\frac{\Vdash_L \phi, \Gamma \quad \phi \Vdash_L \Lambda}{\Vdash_L \Gamma, \Lambda} \text{ (Cut')} \quad \Vdash_M \Delta}{\Vdash_{L+M} \{0\} \times \Gamma, \{0\} \times \Lambda, \{1\} \times \Delta.}$$

All other cases are equally straightforward. So $(\Vdash_{L+M})^+$ is a continuous sequent calculus by Theorem 2.2.12. \square

We claim that this is the identity for a coproduct of \Vdash_L and \Vdash_M . The embeddings from L and M to $L + M$ are defined by the behaviour on the generators $\{0\} \times L \cup \{1\} \times M$. We define them by:

$$\boxed{\frac{\Gamma \Vdash_L \Delta}{\Gamma \vdash_{\iota_0} \{0\} \times \Delta} \quad \frac{\Gamma \Vdash_M \Delta}{\Gamma \vdash_{\iota_1} \{1\} \times \Delta}}$$

For these rules we do not need the restriction $\Gamma \neq \emptyset$ since the preimage of $X \subseteq X + Y$ under the embedding is all of X .

Note that on the left hand side of these sequents arbitrary formulae from L and M may appear. Of course, we could restrict them to elements from the generators of L and M , but this would actually make it harder for us to show the necessary interpolation properties.

Lemma 2.2.18. *The relations $(\vdash_{\iota_0})^+$ and $(\vdash_{\iota_1})^+$ are compatible consequence relations.*

Proof. They are closed under backwards rules since \Vdash_L and \Vdash_M are. For the same reason they satisfy the interpolation properties needed in Corollary 2.2.16.

If we consider (Cut^*) -cuts between \vdash_{ι_0} and \Vdash_{L+M} -sequents we see that those of the form $\langle 0, \phi \rangle, \langle 1, \psi \rangle \Vdash_{L+M}$ are always redundant since no choice function can cover the formula $\langle 1, \psi \rangle$. If we have to use a sequent that is derived as

$$\frac{\Vdash_L \Gamma \quad \Vdash_M \Delta}{\Vdash_{L+M} \{0\} \times \Gamma, \{1\} \times \Delta}$$

we get

$$\frac{\frac{\Vdash_L \Gamma}{\vdash_{\iota_0} \{0\} \times \Gamma}}{\Theta \vdash_{\iota_0} \{0\} \times \Gamma, \{1\} \times \Delta} \text{ (W)}$$

for any $\Theta \subseteq_{\text{fin}} L$. The remaining case follows immediately from the fact that \Vdash_L is closed under (Cut^*) , and so does closure under cuts with \Vdash_L . Thus, all conditions of Corollary 2.2.16 are satisfied. \square

To show that $L + M$ is indeed the coproduct we have to come up with unique mediating morphisms for arbitrary co-cones. Suppose that $\vdash_f: L \rightarrow N$ and $\vdash_g: M \rightarrow N$ are compatible consequence relations, where N is also freely generated. We define \vdash_{fg} by:

$$\boxed{\begin{array}{cc} \frac{\Gamma \vdash_f \Delta}{\{0\} \times \Gamma \vdash_{fg} \Delta} (\Gamma \neq \emptyset) & \frac{\Gamma \vdash_g \Delta}{\{1\} \times \Gamma \vdash_{fg} \Delta} (\Gamma \neq \emptyset) \\ \\ \frac{\vdash_f \Gamma \quad \vdash_g \Delta}{\vdash_{fg} \Gamma, \Delta} & \frac{}{\langle 0, \phi \rangle, \langle 1, \psi \rangle \vdash_{fg}} \end{array}}$$

Lemma 2.2.19. *The rules generate a compatible consequence relation $(\vdash_{fg})^+$, and it satisfies $(\vdash_{\iota_0})^+ \circ (\vdash_{fg})^+ = \vdash_f$ and $(\vdash_{\iota_1})^+ \circ (\vdash_{fg})^+ = \vdash_g$.*

Proof. As before the existence of interpolating formulae for \vdash_{fg} follows directly from the fact that \vdash_f and \vdash_g have interpolants, and for the same reason it is clear that \vdash_{fg} is closed under backwards rules.

The main difficulty of the proof that \vdash_{fg} is closed under cuts with \Vdash_{L+M} and \Vdash_N is to make sure that we consider all the different ways in which such cuts can arise. Let us begin with cuts of the form $\vdash_{fg} \circ \Vdash_N$. If such a cut uses a sequent $\langle 0, \phi \rangle, \langle 1, \psi \rangle \vdash_{fg}$, then we can get the result simply by weakening this formula, whatever the other formulae that are involved in the cut. The same is the case if we have at least one sequent of the form $\{0\} \times \Gamma \vdash_{fg} \Delta$ and one of the form $\{1\} \times \Gamma' \vdash_{fg} \Delta'$, where $\Gamma, \Gamma' \neq \emptyset$. Now, suppose that on the left hand side of (Cut*) there is at least one sequent of the form $\{i\} \times \Gamma \vdash_{fg} \Delta$ with $\Gamma \neq \emptyset$. Then any sequent derived as

$$\frac{\vdash_f \Theta_0 \quad \vdash_g \Theta_1}{\vdash_{fg} \Theta_0, \Theta_1}$$

can be replaced by $\{i\} \times \Gamma \vdash_{fg} \Theta_i$ by applying weakening to the corresponding premise. This shows that we can assume, without loss of generality, that the sequents on the left are all derived by the same rule.

If it is the first or the second rule then closure under cut with \Vdash_N follows immediately from the compatibility of \vdash_f and \vdash_g . Otherwise, the cut must be of the form

$$\frac{\begin{array}{cc} \vdash_{fg} \Gamma_1, \Delta_1 & \Theta_1 \Vdash_N \Lambda \\ \vdots & \vdots \\ \vdash_{fg} \Gamma_n, \Delta_n & \Theta_m \Vdash_N \Lambda \end{array}}{\vdash_{fg} \circ \Vdash_N \Lambda.} \text{ (Cut*)}$$

The sequents $\vdash_{fg} \Gamma_i, \Delta_i$ are derived from $\vdash_f \Gamma_i$ and $\vdash_g \Delta_i$, and we construct

the proof

$$\frac{\frac{\frac{\vdash_f \Gamma_1 \quad \Theta_1 \Vdash_N \Lambda}{\vdots} \quad \frac{\vdash_f \Gamma_n \quad \Theta_m \Vdash_N \Lambda}{\vdots}}{\vdash_f \Lambda} \text{ (Cut*)} \quad \frac{\frac{\vdash_g \Delta_1 \quad \Theta_1 \Vdash_N \Lambda}{\vdots} \quad \frac{\vdash_g \Delta_n \quad \Theta_m \Vdash_N \Lambda}{\vdots}}{\vdash_g \Lambda} \text{ (Cut*)}}{\vdash_{fg} \Lambda}.$$

The side conditions are satisfied since there are less choice functions to consider in each of the two cuts. We have now shown that, up to weakening, we have $\vdash_{fg} \circ \Vdash_N \subseteq \vdash_{fg}$.

For cuts between \Vdash_{L+M} and \vdash_{fg} essentially the same argument shows that we can assume the \Vdash_{L+M} -sequents involved in such a cut to be of one type. This restricts the \vdash_{fg} -sequents that can appear on the right of such a cut. The only new case is

$$\frac{\frac{\frac{\Vdash_{L+M} \{0\} \times \Gamma_1, \{1\} \times \Delta_1}{\vdots} \quad \frac{\Vdash_{L+M} \{0\} \times \Gamma_k, \{1\} \times \Delta_k}{\vdots}}{\Vdash_{L+M} \{0\} \times \Gamma_1, \{1\} \times \Delta_1, \dots, \{0\} \times \Gamma_k, \{1\} \times \Delta_k} \quad \frac{\frac{\{0\} \times \Theta_1 \vdash_{fg} \Lambda}{\vdots} \quad \frac{\{0\} \times \Theta_l \vdash_{fg} \Lambda}{\vdots} \quad \frac{\{1\} \times \Theta'_1 \vdash_{fg} \Lambda'}{\vdots}}{\{0, \phi_1\}, \langle 1, \psi_1 \rangle \vdash_{fg}} \quad \frac{\frac{\{1\} \times \Theta'_m \vdash_{fg} \Lambda'}{\vdots} \quad \langle 0, \phi_n \rangle, \langle 1, \psi_n \rangle \vdash_{fg}}{\vdots}}{\langle 0, \phi_n \rangle, \langle 1, \psi_n \rangle \vdash_{fg}} \text{ (Cut*)}}{\Vdash_{L+M} \circ \vdash_{fg} \Lambda, \Lambda'}.$$

We observe that the side condition must in particular be satisfied for choice functions that pick elements exclusively from the $\{0\} \times \Gamma_i$ or exclusively from the $\{1\} \times \Delta_j$. For such choice functions only certain sequents on the right can be relevant. Since we know how the individual sequents in the cut were derived from the respective continuous sequent calculi $\Vdash_L, \Vdash_M, \vdash_f$ and \vdash_g this implies that we can construct the proof

$$\frac{\frac{\frac{\Vdash_L \Gamma_1 \quad \Theta_1 \vdash_f \Lambda}{\vdots} \quad \frac{\Vdash_L \Gamma_k \quad \Theta_l \vdash_f \Lambda}{\vdots}}{\vdash_f \Lambda} \text{ (Cut*)} \quad \frac{\frac{\Vdash_L \Delta_1 \quad \Theta'_1 \vdash_g \Lambda'}{\vdots} \quad \frac{\Vdash_L \Delta_k \quad \Theta'_m \vdash_g \Lambda'}{\vdots}}{\vdash_g \Lambda'} \text{ (Cut*)}}{\vdash_{fg} \Lambda, \Lambda'}.$$

We have thus proved $\Vdash_{L+M} \circ \vdash_{fg} \subseteq \vdash_{fg}$, up to weakening. So, \vdash_{fg} satisfies all conditions of Corollary 2.2.16 showing that $(\vdash_{fg})^+$ is a compatible consequence relation.

Now we check $(\vdash_{\iota_0})^+ \circ (\vdash_{fg})^+ = (\vdash_{\iota_0} \circ \vdash_{fg})^+ \stackrel{!}{=} \vdash_f$, where the first equality follows from Proposition 2.2.8. For $\Gamma \vdash_f \Delta$ we can find a $\phi \in L$ such that $\Gamma \Vdash_L \phi$ and $\phi \vdash_f \Delta$ since \vdash_f is compatible. We construct the proof

$$\frac{\frac{\Gamma \Vdash_L \phi}{\Gamma \vdash_{\iota_0} \langle 0, \phi \rangle} \quad \frac{\phi \vdash_f \Delta}{\langle 0, \phi \rangle \vdash_{fg} \Delta}}{\Gamma (\vdash_{\iota_0} \circ \vdash_{fg}) \Delta.}$$

which shows $\vdash_f \subseteq \vdash_{\iota_0} \circ \vdash_{fg} \subseteq (\vdash_{\iota_0} \circ \vdash_{fg})^+$.

Because of the structural similarities between the rules for \Vdash_{L+M} and \vdash_{fg} the proof of the other containment is almost identical to the argument for $\vdash_{\iota_0} \circ \Vdash_{L+M} \subseteq \vdash_{\iota_0}$ in the previous lemma. \square

Unfortunately, the uniqueness of the mediating morphism $(\vdash_{fg})^+$ is harder to show. For this we need more information about the exact relationship between \Vdash_{L+M} and \vdash_{ι_i} . We establish the necessary prerequisites in a series of lemmata.

Lemma 2.2.20. *If $(\vdash_{\iota_0})^+ \Delta$ then we can prove $\{0\} \times \Gamma (\Vdash_{L+M})^+ \Delta$ for all non-empty $\Gamma \subseteq_{\text{fin}} L$.*

Proof. The proof is by induction on the derivation of $(\vdash_{\iota_0})^+ \Delta$. For the base case consider

$$\frac{\Vdash_L \tilde{\Delta}}{\vdash_{\iota_0} \{0\} \times \tilde{\Delta}.}$$

We can immediately construct the proof

$$\frac{\frac{\Vdash_L \tilde{\Delta}}{\Gamma \Vdash_L \tilde{\Delta}} (\mathbf{W})}{\{0\} \times \Gamma \Vdash_{L+M} \{0\} \times \tilde{\Delta}.}$$

The rest of the induction poses no problems since the term $\{0\} \times \Gamma$ on the left does not interfere with the logical rule that is used on the right. \square

Lemma 2.2.21. *For any derivable sequent $\langle 0, \phi \rangle, \Gamma (\Vdash_{L+M})^+ \Delta$ and any $\psi \in L$ we can derive $\langle 0, \phi \wedge \psi \rangle, \Gamma (\Vdash_{L+M})^+ \Delta$.*

Proof. The proof by induction is straightforward: For the base case

$$\frac{\phi, \tilde{\Gamma} \Vdash_L \tilde{\Delta}}{\langle 0, \phi \rangle, \{0\} \times \tilde{\Gamma} \Vdash_{L+M} \{0\} \times \tilde{\Delta}}$$

we construct

$$\frac{\frac{\phi, \tilde{\Gamma} \Vdash_L \tilde{\Delta}}{\phi, \psi, \tilde{\Gamma} \Vdash_L \tilde{\Delta}} \text{ (W)}}{\phi \wedge \psi, \tilde{\Gamma} \Vdash_L \tilde{\Delta}} \text{ (L}\wedge\text{)}}{\langle 0, \phi \wedge \psi \rangle, \{0\} \times \tilde{\Gamma} \Vdash_{L+M} \{0\} \times \tilde{\Delta}}.$$

No logical rule can introduce the term $\langle 0, \phi \rangle$, so it is always a side formula. And for weakening it is clear that if we can weaken with $\langle 0, \phi \rangle$ then we may as well weaken with $\langle 0, \phi \wedge \psi \rangle$ instead. \square

Lemma 2.2.22. *If the sequent $\langle 0, \phi \rangle, \langle 0, \psi \rangle, \Gamma \Vdash_{L+M}^+ \Delta$ is derivable then so is $\langle 0, \phi \wedge \psi \rangle, \Gamma \Vdash_{L+M}^+ \Delta$.*

Proof. As before the only interesting bits of the induction are the cases where either $\langle 0, \phi \rangle$ or $\langle 0, \psi \rangle$ is introduced. For the base

$$\frac{\phi, \psi, \tilde{\Gamma} \Vdash_L \tilde{\Delta}}{\langle 0, \phi \rangle, \langle 0, \psi \rangle, \{0\} \times \tilde{\Gamma} \Vdash_{L+M} \{0\} \times \tilde{\Delta}}$$

we immediately get

$$\frac{\frac{\phi, \psi, \tilde{\Gamma} \Vdash_L \tilde{\Delta}}{\phi \wedge \psi, \tilde{\Gamma} \Vdash_L \tilde{\Delta}} \text{ (L}\wedge\text{)}}{\langle 0, \phi \wedge \psi \rangle, \{0\} \times \tilde{\Gamma} \Vdash_{L+M} \{0\} \times \tilde{\Delta}}$$

and for weakening we can use Lemma 2.2.21 \square

Lemma 2.2.23. *If $\langle 0, \phi \rangle, \Gamma \Vdash_{L+M}^+ \Delta$ and $\langle 0, \psi \rangle, \Gamma' \Vdash_{L+M}^+ \Delta'$ is derivable, then so is $\langle 0, \phi \vee \psi \rangle, \Gamma, \Gamma' \Vdash_{L+M}^+ \Delta, \Delta'$.*

Proof. The proof is essentially the same as before; we perform an induction on the derivation of both sequents. The base case is done by using multiplicative (L \vee), and all the other cases, even weakening, are trivial. \square

We can now make the connection between \Vdash_{L+M} and \vdash_{ι_0} :

Proposition 2.2.24. 1. *If $\{0\} \times \Gamma \Vdash_{L+M}^+ \Delta$ then we also have $\Gamma \vdash_{\iota_0}^+ \Delta$.*

2. *If $\Gamma \vdash_{\iota_0}^+ \Delta$ and $\Gamma \neq \emptyset$ then we can derive $\{0\} \times \Gamma \Vdash_{L+M}^+ \Delta$.*

Proof. Both claims are proved by induction on the derivation of the sequent in question. The first claim is straightforward and does not need any of the previous lemmata.

For the second claim the base case is not a problem because of the side condition $\Gamma \neq \emptyset$. The only difficulty with weakening is the case

$$\frac{\vdash_{\iota_0} \Delta}{\Gamma \vdash_{\iota_0} \Delta.} \text{ (W)}$$

This is, however, taken care of by Lemma 2.2.20. The only non-trivial logical rules are $(L\wedge)$ and $(L\vee)$; this is where we need Lemma 2.2.22 and 2.2.23, \square

This allows us to show finally that we have actually constructed the coproduct of L and M :

Theorem 2.2.25. *The continuous sequent calculus $L+M$, together with the compatible consequence relations $(\vdash_{\iota_0})^+$ and $(\vdash_{\iota_1})^+$, is a coproduct of L and M in MLS.*

Proof. Suppose that $\vdash_f: N \rightarrow L$ and $\vdash_g: N \rightarrow M$ are compatible consequence relations, where N is also freely generated. Lemma 2.2.19 says that $(\vdash_{fg})^+$ is a compatible consequence relation and, moreover, that it satisfies $(\vdash_{\iota_0})^+ \circ (\vdash_{fg})^+ = \vdash_f$ and $(\vdash_{\iota_1})^+ \circ (\vdash_{fg})^+ = \vdash_g$. Hence, we only have to show that it is unique.

To this end let us suppose that $\vdash_d: L+M \rightarrow N$ is any morphism with this property. We first show that the generators \vdash_{fg} are contained in \vdash_d . If we have $\Gamma \vdash_f \Delta$, where $\Gamma \neq \emptyset$, then compatibility of \vdash_f allows us to find a $\phi \in L$ such that $\Gamma \Vdash_L \phi$ and $\phi \vdash_f \Delta$. Since we assume $\vdash_f = (\vdash_{\iota_0})^+ \circ \vdash_d$ we find a $\psi \in L+M$ such that $\phi (\vdash_{\iota_0})^+ \psi$ and $\psi \vdash_d \Delta$. From $\phi (\vdash_{\iota_0})^+ \psi$ we get $\langle 0, \phi \rangle (\Vdash_{L+M})^+ \psi$ because of Proposition 2.2.24.(2) and we construct the proof

$$\frac{\frac{\Gamma \Vdash_L \phi}{\{0\} \times \Gamma \Vdash_{L+M} \langle 0, \phi \rangle} \quad \frac{\langle 0, \phi \rangle (\Vdash_{L+M})^+ \psi \quad \psi \vdash_d \Delta}{\langle 0, \phi \rangle \vdash_d \Delta} \text{ (Cut)}}{\{0\} \times \Gamma \vdash_d \Delta.} \text{ (Cut)}$$

We still have to check two rules that generate \vdash_{fg} -sequents. Because of

$$\frac{\frac{\overline{\langle 0, \phi \rangle, \langle 1, \psi \rangle \Vdash_{L+M}}}{\langle 0, \phi \rangle, \langle 1, \psi \rangle (\Vdash_{L+M})^+ \perp} \text{ (W)} \quad \frac{}{\perp \vdash_d} \text{ (L}\perp\text{)}}{\langle 0, \phi \rangle, \langle 1, \psi \rangle \vdash_d} \text{ (Cut)}$$

those of the form $\langle 0, \phi \rangle, \langle 1, \psi \rangle \vdash_{fg}$ are not a problem. We consider the last rule

$$\frac{\vdash_f \Gamma \quad \vdash_g \Delta}{\vdash_{fg} \Gamma, \Delta.}$$

Because of weakening and what we have already shown we get the sequents $\langle 0, \top \rangle \vdash_d \Gamma$ and $\langle 1, \top \rangle \vdash_d \Delta$ and we can use them in the following derivation:

$$\frac{\frac{\frac{}{\vdash_L \top} (\text{RT}) \quad \frac{}{\vdash_M \top} (\text{RT})}{\vdash_{L+M} \langle 0, \top \rangle, \langle 1, \top \rangle} \quad \frac{\langle 0, \top \rangle \vdash_d \Gamma \quad \langle 1, \top \rangle \vdash_d \Delta}{\vdash_d \Gamma, \Delta} (\text{Cut}^*)}{\vdash_d \Gamma, \Delta} (\text{Cut}^*)$$

This completes the proof of $\vdash_{fg} \subseteq \vdash_d$.

For the reverse containment we know from Lemma 2.2.1 that we only have to show that \vdash_d -sequents that relate generators lie in $(\vdash_{fg})^+$. As before the only interesting such sequents are those of the form $\{i\} \times \Gamma \vdash_d \Delta$; those containing $\langle 0, \phi \rangle$ and $\langle 1, \psi \rangle$ formulae on the left are trivial, and we can take care of the case where the left hand side is empty by (LT) and (RT).

So, we take a sequent $\{0\} \times \Gamma \vdash_d \Delta$, where $\Gamma \neq \emptyset$. We exploit compatibility of \vdash_d to get $\{0\} \times \Gamma (\vdash_{L+M})^+ \phi$ and $\phi \vdash_d \Delta$ for a $\phi \in L + M$. By Proposition 2.2.24.(1) this implies $\Gamma (\vdash_{l_0})^+ \phi$ and we get

$$\frac{\frac{\Gamma (\vdash_{l_0})^+ \phi \quad \phi \vdash_d \Delta}{\Gamma \vdash_f \Delta} (\text{Cut})}{\{0\} \times \Gamma \vdash_{fg} \Delta} (\text{Cut})$$

Hence, we have shown that \vdash_d and \vdash_{fg} are equal.

If N is not freely generated, we construct $T(N)$ which is isomorphic to N by Proposition 2.2.5; let us call the isomorphism $\vdash_\epsilon: T(L) \rightarrow L$. We get a mediating morphism \vdash_d for \vdash_f and \vdash_g composed with $(\vdash_\epsilon)^{-1}$. It is straightforward to check that $\vdash_d \circ \vdash_\epsilon$ is the unique mediating morphism for \vdash_f and \vdash_g . So, $L \times M$ is a coproduct in the category MLS. \square

It is clear from Lemma 2.2.1 that the free algebra on no generators supports exactly two continuous sequent calculi; one is created by the empty relation, the other by the empty sequent $\emptyset \vdash \emptyset$. Because of weakening, the continuous sequent calculus $\mathbf{0}$ generated by the latter one relates all finite subsets of formulae. We infer that it is a *zero object* in MLS, i.e. it is initial and terminal.

Corollary 2.2.26. *MLS has all finite coproducts.*

Proof. The only thing we have to show is that we can also form the coproduct of continuous sequent calculi that are not necessarily freely generated. This is a consequence of Proposition 2.2.5 and the following observation: If objects X_i and X'_i are isomorphic in any category and the coproduct of the X_i exists, then it is also a coproduct of the X'_i . \square

Remark. *There is also a more conceptual way of explaining why the objects of MLS that are not freely generated do not cause any problems: Our constructions show that the full subcategory of freely generated continuous sequent calculi has coproducts. As we already observed after Proposition 2.2.5 this category is equivalent to MLS. But as equivalent categories have the same categorical properties (see [FS90, 1.39]), MLS has all coproducts.*

The construction for coproducts also yields products:

Corollary 2.2.27. *The category MLS has all finite products.*

Proof. As pointed out directly after the definition of MLS on page 67 the category is self-dual. Unfortunately, the corresponding involution $(\cdot)^{\text{op}}$ does not fix objects; it reverses \Vdash and interchanges \wedge , \vee , \top and \perp . Hence, we only know $L \times M \cong (L^{\text{op}} + M^{\text{op}})^{\text{op}}$. \square

There is no reason why $L + M$ and the product $(L^{\text{op}} + M^{\text{op}})^{\text{op}}$ should be equal. Somewhat surprisingly they are, however, isomorphic. In the next chapter we will see a semantic argument why this is the case.

An alternative proof is to give the projections and the corresponding pairing directly which make $L + M$ into the product of L and M . The projections are generated by the rules:

$$\frac{\Gamma \Vdash_L \Delta}{\{0\} \times \Gamma \vdash_{\pi_0} \Delta} \qquad \frac{\Gamma \Vdash_M \Delta}{\{1\} \times \Gamma \vdash_{\pi_1} \Delta}$$

$$\frac{}{\{1\} \times \Gamma \vdash_{\pi_0} (\Gamma \neq \emptyset)} \qquad \frac{}{\{0\} \times \Gamma \vdash_{\pi_1} (\Gamma \neq \emptyset)}$$

For compatible consequence relations $\vdash_f: N \rightarrow L$ and $\vdash_g: N \rightarrow M$ the mediating morphism is generated by the single rule

$$\frac{\Gamma \vdash_f \Delta \quad \Gamma \vdash_g \Theta}{\Gamma \vdash_{\langle f, g \rangle} \{0\} \times \Delta, \{1\} \times \Theta.}$$

The proof that the rules generate compatible consequence relations and that they satisfy the conditions defining products uses the same techniques as the one for coproducts, and we do not consider it in any detail. To show that $(\vdash_{\langle f, g \rangle})^+$ is unique we need an auxiliary observation similar to Proposition 2.2.24, namely that the rule

$$\frac{\Gamma (\vdash_{\pi_0})^+ \Delta \quad \Gamma (\vdash_{\pi_1})^+ \Theta}{\Gamma (\Vdash_{L+M})^+ \{0\} \times \Delta, \{1\} \times \Theta}$$

is admissible.

Chapter 3

Semantics

In this chapter we make the connection between the topological spaces we studied in Chapter 1 and the syntactic objects of the last chapter. As mentioned before, continuous sequent calculi correspond to stably compact spaces and compatible consequence relations to ‘closed’ relations between them. A continuous functions between two spaces can be considered to be a special instance of such a relation, and we can characterise the consequence relations that give rise to functions in purely syntactic, i.e. logical, terms.

We first establish this link between syntax and semantics in an abstract way that says that the respective categories are equivalent. Then we investigate how we can determine whether a concrete continuous sequent calculus actually represents a certain stably compact space. Once we have these tools available, we can use them to perform domain constructions in logical form. We consider several examples in some detail. Function spaces are notoriously hard and we look at some of the problems that arise.

The space of relations, on the other hand, is much easier to handle. Certain relations can be understood to correspond to functions and we study this problem with logical, topological and categorical methods.

3.1 Logic and Topological Spaces

We begin by looking at theories, sets of formulae that are closed under internal reasoning \Vdash . This is essentially still a proof-theoretic concept, but a closer scrutiny of the poset of all theories on a continuous sequent calculus reveals that they are arithmetic locals, the Stone duals of stably compact spaces. Hence, these considerations lead directly to the semantics of continuous sequent calculi.

There is a close resemblance between theories, or filters, for continuous sequent calculi and those for *strong proximity lattices* which were introduced in [JS96, JS98]. The latter can in fact be understood as Lindenbaum alge-

bras for our logic. We make this connection precise and finally we give a topological semantics for both, continuous sequent calculi and compatible consequence relations.

3.1.1 Theories and Models

Let \vdash be a compatible consequence relation from L to M . For $X \subseteq L$ and $Y \subseteq M$, we define

$$\begin{aligned} X[\vdash] &:= \{ \phi \in M \mid (\exists \Gamma \subseteq_{\text{fin}} X) \Gamma \vdash \phi \} \\ [\vdash]Y &:= \{ \phi \in L \mid (\exists \Delta \subseteq_{\text{fin}} Y) \phi \vdash \Delta \} \end{aligned}$$

which we can think of as the \vdash -consequences of X and the dual for Y . As usual, in the case of singleton sets we write $\phi[\vdash]$ for $\{\phi\}[\vdash]$, and $[\vdash]\phi$ for $[\vdash]\{\phi\}$.

A *filter* of L is a set $F \subseteq L$ such that $F = F[\Vdash_L]$; an *ideal* of L is a set $I \subseteq L$ such that $I = [\Vdash_L]I$. Let $\text{filt}(L)$ and $\text{idl}(L)$ denote the partial orders of filters and ideals, respectively, both ordered by inclusion.

Consider the role that filters and ideals play in logic. Roughly, a filter corresponds to a theory. One typically says that a theory is consistent if it is not the entire language, and one formulation of soundness and completeness is that a theory is consistent if and only if it has a model. The latter means essentially that the corresponding filter is contained in a prime filter. We will return to this point later. First we prove some properties about filters and the poset $\text{filt}(L)$.

Proposition 3.1.1. *A subset F of a continuous sequent calculus is a filter if and only if*

1. $\top \in F$;
2. $\phi, \psi \in F$ implies $\phi \wedge \psi \in F$; and
3. $\phi \in F$ if and only if for some $\psi \in F$, $\psi \Vdash \phi$.

Furthermore, in any filter F , if $\phi \in F$ (or $\psi \in F$) then $\phi \vee \psi \in F$, and $\phi \wedge \psi \in F$ implies $\phi, \psi \in F$.

Proof. If F is a filter then the first two conditions are immediate from the rules (RT) and (R \wedge). The third one follows from Proposition 2.1.12.

Conversely, the last condition implies $F \subseteq F[\Vdash]$. For the other inclusion, suppose $\phi \in F[\Vdash]$, i.e. $\Gamma \Vdash \phi$ for $\Gamma \subseteq_{\text{fin}} F$. This implies $\bigwedge \Gamma \Vdash \phi$ by (L \wedge) and $\bigwedge \Gamma \in F$ because of the second condition; weakening and the first one take care of the trivial case $\Gamma = \emptyset$. So we have $\phi \in F$ by the third condition.

Weakening and (RV) show that for a filter F , if $\phi \in F$ or $\psi \in F$ then the element $\phi \vee \psi$ must also lie in F . For the last statement suppose $\phi \wedge \psi$ is an

element of a filter F . From condition (3) we get a $\chi \in F$ such that $\chi \Vdash \phi \wedge \psi$. So backwards $(R\wedge)$ implies $\chi \Vdash \phi$ and $\chi \Vdash \psi$ thus proving $\phi, \psi \in F$. \square

Condition (3) of the proposition is usually called *roundness* of filters. We will use it, and the other properties listed in the proposition, quite frequently in the remainder of this chapter.

For filters and ideals we can restrict \vdash to the singletons it relates:

Corollary 3.1.2. *Let $F \subseteq L$ be a filter and $\vdash: L \rightarrow M$ a compatible consequence relation. Then we have*

$$F[\vdash] = \{\phi \in M \mid (\exists \psi \in F) \psi \vdash \phi\},$$

and dually for ideals.

Proof. Suppose we have $\Gamma \subseteq_{\text{fin}} F$ and a sequent $\Gamma \vdash \phi$. By $(L\wedge)$ we get $\bigwedge \Gamma \vdash \phi$ and from the previous proposition we infer $\bigwedge \Gamma \in F$. This shows one inclusion, the other one is trivial. \square

We can use this observation in many proofs and as it is very elementary we will do so tacitly. Next, we observe that there is always a wealth of filters and ideals, namely the ones generated as ‘consequences’ of an arbitrary set.

Lemma 3.1.3. *$X[\vdash]$ is always a filter and $[\vdash]Y$ is always an ideal.*

Proof. Given any X the set $X[\vdash]$ clearly satisfies the first two conditions of Proposition 3.1.1. Since we made the general assumption that \vdash is compatible, the third condition follows from Proposition 2.1.12.

The result for $[\vdash]Y$ is for free by duality. The same goes for the remaining arguments in this section; we hence omit pointing it out every time. \square

We now embark on a first investigation of the internal structure of $\text{filt}(L)$ and $\text{idl}(L)$, namely the part that is inherited from the lattice structure of the power set of L .

Lemma 3.1.4. *The posets $\text{filt}(L)$ and $\text{idl}(L)$ are closed under directed unions and finite intersections.*

Proof. Since sequents are relations between finite sets it is clear that a directed union of filters is again a filter.

For finite meets we first observe that L itself is always a filter. Given two filters F and G we verify that $F \cap G$ is again a filter using Proposition 3.1.1. The first two conditions and the implication $F \cap G \ni \psi \Vdash \phi \implies \phi \in F \cap G$ are immediate consequences of F and G being filters. For $\phi \in F \cap G$ we get elements $\psi_F \in F$ and $\psi_G \in G$ such that $\psi_F \Vdash \phi$ and $\psi_G \Vdash \phi$ since F and G are round. This means $\psi_F \vee \psi_G \Vdash \phi$ by (LV) , and again from Proposition 3.1.1 we know that $\psi_F \vee \psi_G \in F \cap G$. \square

These posets $\text{filt}(L)$ and $\text{idl}(L)$ are dcpo's and \wedge -semilattices. In fact, they have much more structure and we come back to this points when we make the connection to topological spaces in Section 3.1.3. The reasons to emphasise finite infima and directed suprema at this point are the following: Firstly, these operations are very simple because they are just ordinary intersection and union. Secondly, $\text{filt}(L)$ and $\text{idl}(L)$ being \wedge -semilattices is a cue that it may be worthwhile to characterise the \wedge -prime elements. Directed suprema are needed later to show that there are 'enough' prime ideals and filters.

But before we tackle the primes we list some properties of the map $(\cdot)[\Vdash]$ which will be useful later:

Lemma 3.1.5. *The following properties hold:*

1. $X \subseteq Y$ implies $X[\Vdash_L] \subseteq Y[\Vdash_L]$;
2. $(\{\phi\} \cup X)[\Vdash_L] \cap (\{\psi\} \cup X)[\Vdash_L] = (\{\phi \vee \psi\} \cup X)[\Vdash_L]$;
3. if $\phi \in X[\Vdash_L]$ then $(\{\phi\} \cup Y)[\Vdash_L] \subseteq (X \cup Y)[\Vdash_L]$;
4. $\phi \in X[\Vdash_L]$ implies $(\{\phi\} \cup X)[\Vdash_L] = X[\Vdash_L]$; and
5. $(\{\phi\} \cup X)[\Vdash] = \bigcup_{\psi \in \phi[\Vdash]}^{\uparrow} (\{\psi\} \cup X)[\Vdash]$.

Proof. The first two properties are obvious from the definition and the rule (LV). The next one follows from (Cut') which is admissible by Proposition 2.1.8, and the fourth is an immediate consequence of (1) and (3).

For the last property one containment, namely

$$\bigcup_{\psi \in \phi[\Vdash]}^{\uparrow} (\{\psi\} \cup X)[\Vdash] \subseteq (\{\phi\} \cup X)[\Vdash]$$

follows directly from (3). For $\Gamma \subseteq_{\text{fin}} X$ and $\Gamma, \phi \Vdash \theta$ we interpolate to get $\phi \Vdash \psi$ and $\Gamma, \psi \Vdash \theta$ showing the other containment. We still have to prove that the union is directed. Given $\psi, \chi \in \phi[\Vdash]$ we find interpolants $\psi' \Vdash \psi$ and $\chi' \Vdash \chi$ such that $\phi \Vdash \psi'$ and $\phi \Vdash \chi'$. By (W), (L \wedge) and (R \wedge) we get $\psi' \wedge \chi' \Vdash \psi$ and $\psi' \wedge \chi' \Vdash \chi$, as well as $\phi \Vdash \psi' \wedge \chi'$. Because of (3), this implies

$$(\{\psi\} \cup X)[\Vdash], (\{\chi\} \cup X)[\Vdash] \subseteq (\{\psi' \wedge \chi'\} \cup X)[\Vdash]$$

and thus directedness. □

Proposition 3.1.6. *For a filter $F \subseteq L$ the following are equivalent:*

1. F is a \wedge -prime element of $\text{filt}(L)$;
2. F is inaccessible by finite disjunctions, i.e. $\phi \vee \psi \in F$ implies $\phi \in F$ or $\psi \in F$, and $\perp \notin F$;

3. if $\Gamma \Vdash_L \Delta$ for some $\Gamma \subseteq_{\text{fin}} F$ then $F \cap \Delta \neq \emptyset$.

The dual conditions characterise the \wedge -prime elements of $\text{idl}(L)$.

Proof. “(1) \Rightarrow (3)”: Suppose $\Gamma \Vdash \delta_1, \dots, \delta_n$ for some $\Gamma \subseteq_{\text{fin}} F$. Then we can find interpolants $\delta'_i \Vdash \delta_i$ such that $\Gamma \Vdash \delta'_1, \dots, \delta'_n$ and iterated applications of (RV) yields $\Gamma \Vdash \bigvee \Delta'$, where $\Delta' = \{\delta'_1, \dots, \delta'_n\}$. Using the previous lemma and the fact that F is a filter we get

$$F = F[\Vdash] \supseteq \Gamma[\Vdash] = (\Gamma \cup \{\bigvee \Delta'\})[\Vdash_L] = \bigcap_{i=1}^n ((\Gamma \cup \{\delta'_i\})[\Vdash]).$$

Since F is \wedge -prime there is a δ'_i such that $\delta'_i[\Vdash] \subseteq (\Gamma \cup \{\delta'_i\})[\Vdash] \subseteq F$. Because of $\delta'_i \Vdash \delta_i$ this shows $F \cap \Delta \neq \emptyset$.

“(3) \Rightarrow (2)”: Evidently, $\perp \notin F$, for otherwise we have $\perp \Vdash \emptyset$. If $\phi \vee \psi \in F$ then there is a set $\Gamma \subseteq_{\text{fin}} F$ such that $\Gamma \Vdash \phi \vee \psi$ which implies $\Gamma \Vdash \phi, \psi$ by backwards (RV). This means that either ϕ or ψ is in F .

“(2) \Rightarrow (1)”: If G and H are filters not contained in F then we can pick $\phi \in G \setminus F$ and $\psi \in H \setminus F$. Thus, we have $\phi \vee \psi \in (G \cap H) \setminus F$ by (2) and Proposition 3.1.1. This says that the intersection $G \cap H$ is also not contained in F which is hence \wedge -prime. \square

Remark. The converse of the third condition is equivalent to $F \subseteq F[\Vdash]$. Hence, we can directly characterise the sets that are prime filters: They are precisely those $F \subseteq L$ such that for all $\Delta \subseteq_{\text{fin}} L$ we have $F \cap \Delta \neq \emptyset$ if and only if there is a subset $\Gamma \subseteq_{\text{fin}} F$ such that $\Gamma \Vdash \Delta$.

Also note that to get to (2) we needed backwards rules for the first time.

A set satisfying any of these equivalent conditions is called a *prime filter*. A set satisfying the dual conditions is called a *prime ideal*. The following lemma is a useful tool to construct prime filters.

Lemma 3.1.7. If I is an ideal and F a filter, maximal with the property $F \cap I = \emptyset$, then F is prime.

Proof. Assume F not to be prime, then by the previous proposition there are sets $\Gamma \subseteq_{\text{fin}} F$ and $\Delta = \{\delta_1, \dots, \delta_n\}$ such that $\Gamma \Vdash \Delta$ but $F \cap \Delta = \emptyset$. Using interpolation we can make sure that for all δ_i we have $\delta_i[\Vdash] \setminus F \neq \emptyset$, i.e. the filters $(F \cup \{\delta_i\})[\Vdash]$ properly contain F . Because of the maximality of F this implies that for each $i = 1, \dots, n$ there is a $\phi_i \in I$ and a set $\Gamma_i \subseteq_{\text{fin}} F$ such that $\Gamma_i, \delta_i \Vdash \phi_i$. We put this together to get

$$\begin{array}{c} \Gamma_1, \delta_1 \Vdash \phi_1 \\ \vdots \\ \Gamma_n, \delta_n \Vdash \phi_n \\ \hline \Gamma \Vdash \Delta \quad \Gamma_n, \delta_n \Vdash \phi_n \quad (\text{Cut}')^* \\ \Gamma, \Gamma_1, \dots, \Gamma_n \Vdash \phi_1, \dots, \phi_n \\ \hline \Gamma, \Gamma_1, \dots, \Gamma_n \Vdash \phi_1 \vee \dots \vee \phi_n. \quad (\text{RV})^* \end{array}$$

We have $\Gamma \cup \Gamma_1 \cup \dots \cup \Gamma_n \subseteq F$, and as F is a filter and I is an ideal, $\phi_1 \vee \dots \vee \phi_n$ lies in F and I . Note that the case $\Delta = \emptyset$ is trivial as it implies $\perp \in F \cap I$. In either case we get a contradiction to $F \cap I = \emptyset$, thus F must be a prime filter. \square

Prime filters and prime ideals are almost complements:

Proposition 3.1.8. *If P is a prime filter then $[\Vdash](L \setminus P)$ is a prime ideal. Conversely, for a prime ideal I the set $(L \setminus I)[\Vdash]$ is a prime filter, and these two maps are mutually inverse order anti-isomorphisms.*

Proof. The set $I := [\Vdash](L \setminus P)$ is an ideal by Lemma 3.1.3. It does not meet P since $P \ni \phi \Vdash \Delta \subseteq (L \setminus P)$ implies $\phi \Vdash \bigvee \Delta \in (L \setminus P)$ as P is prime, and the latter contradicts P being a filter. As I is clearly the largest ideal satisfying $I \cap P = \emptyset$ it is prime by the previous lemma.

Because of duality all that remains to show is $P = (L \setminus I)[\Vdash]$. Obviously, P is contained in the right hand side as the latter is the largest filter that has an empty intersection with I . To see the converse, suppose $\phi \in ((L \setminus I)[\Vdash]) \setminus P$. We get a set $\Gamma \subseteq (L \setminus I)$ such that $\Gamma \Vdash \phi$ and hence $\bigwedge \Gamma \Vdash \phi$. This implies $\bigwedge \Gamma \in I$ by the construction of I which contradicts the primeness of I . \square

Consistency

As discussed at the beginning of this section a filter $F \subseteq L$ is *consistent* if it is a proper subset of L , and completeness can be expressed as each such filter being contained in a prime filter. But closer inspection of the proofs of completeness theorems, say for Gentzen's system \mathbf{K} [Gen34], shows that more is proved. In particular, we have nearly complete freedom to choose, apart from the formulae in F , what formulae are *not* to be satisfied in a particular model.

Say that a pair of sets (X, Y) for $X \subseteq L$ and $Y \subseteq M$ is *\vdash -consistent* provided that for all $\Gamma \subseteq_{\text{fin}} X$ and $\Delta \subseteq_{\text{fin}} Y$, it is the case that $\Gamma \not\vdash \Delta$. The idea here is to understand X as a set of formulae that 'hold' in L and Y as a set of formulae that do not hold in M . So the least we should expect is that \vdash does not contradict this understanding.

Among other things, the next proposition shows that consistency has to do essentially with filters and ideals.

Proposition 3.1.9. *For every compatible consequence relation \vdash from L to M the following are equivalent:*

1. (X, Y) is \vdash -consistent;
2. $(X, [\Vdash_M]Y)$ is \vdash -consistent;
3. $(X[\Vdash_L], Y)$ is \vdash -consistent;

4. $(X[\vdash], Y)$ is \Vdash_M -consistent;
5. $(X, [\vdash]Y)$ is \Vdash_L -consistent;
6. $X[\vdash] \cap [\Vdash_M]Y = \emptyset$;
7. $X[\Vdash_L] \cap [\vdash]Y = \emptyset$;
8. (X, I) is \vdash -consistent for some prime ideal $I \supseteq [\Vdash_M]Y$; and
9. (F, Y) is \vdash -consistent for some prime filter $F \supseteq X[\Vdash_L]$.

Proof. “(1) \Rightarrow (3)”: Suppose the sets $\Delta = \{\delta_1, \dots, \delta_n\} \subseteq X[\vdash]$ and $\Theta \subseteq_{\text{fin}} Y$ are such that $\Delta \vdash \Theta$. Then for each $\delta_i \in \Delta$ there is a $\Gamma_i \subseteq_{\text{fin}} X$ such that $\Gamma_i \Vdash \delta_i$. Multiple application of (L-Cut) yields $\Gamma_1, \dots, \Gamma_n \vdash \Theta$ which shows that (X, Y) is \vdash -inconsistent.

“(3) \Rightarrow (5)”: Given sequents $\delta_1 \vdash \Theta_1; \dots; \delta_n \vdash \Theta_n$ and $\Gamma \Vdash \Delta$, where $\Delta = \{\delta_1, \dots, \delta_n\}$, $\Theta_i \subseteq_{\text{fin}} Y$ and $\Gamma \subseteq_{\text{fin}} X$, we get $\bigvee \Delta \vdash \Theta_1, \dots, \Theta_n$ and $\Gamma \Vdash \bigvee \Delta$ by iterated application of (LV) and (RV), respectively. The element $\bigvee \Delta$ clearly lies in $X[\vdash]$ and because of $\Theta_1 \cup \dots \cup \Theta_n \subseteq Y$ the pair $(X[\vdash], Y)$ is \vdash -inconsistent.

“(5) \Rightarrow (7)”: The contraposition of this implication is obvious.

“(7) \Rightarrow (1)”: Let $\Gamma = \{\gamma_1, \dots, \gamma_n\} \subseteq X$ and $\Delta \subseteq_{\text{fin}} Y$ be sets such that $\Gamma \vdash \Delta$. By interpolation in the form of (L-Int') we get sequents $\gamma_i \Vdash \gamma'_i$ and $\Gamma' = \{\gamma'_1, \dots, \gamma'_n\} \vdash \Delta$. We construct

$$\frac{\begin{array}{c} \gamma_1 \Vdash \gamma'_1 \\ \vdots \\ \gamma_n \Vdash \gamma'_n \end{array}}{\Gamma \Vdash \bigwedge \Gamma'} (\text{R}\wedge)^* \quad \text{and} \quad \frac{\Gamma' \vdash \Delta}{\bigwedge \Gamma' \vdash \Delta} (\text{L}\wedge)^*$$

which shows $\bigwedge \Gamma' \in X[\Vdash] \cap [\vdash]Y$, thus contradicting (7).

We have shown the equivalence of the conditions (1), (3), (5) and (7). We proceed by proving (9) equivalent. To this end we assume that the filter $X[\Vdash]$ does not meet the ideal $[\vdash]Y$. From Lemma 3.1.4 we know that a directed supremum of filters is simply the union and hence if such filters have an empty intersection with any given ideal then so does their supremum. Thus we can apply Zorn's Lemma and get a maximal filter $F \supseteq X[\Vdash]$ such that $F \cap [\vdash]Y = \emptyset$. By Lemma 3.1.7 this F is a prime filter and since (7) is equivalent to (1) we also have that (F, Y) is \vdash -consistent.

Conversely, if (F, Y) is \vdash -consistent for a prime filter $F \supseteq X[\Vdash]$ then clearly $(X[\Vdash], Y)$ is also \vdash -consistent.

By duality (1), (2), (4), (6) and (8), and hence all conditions of the proposition are equivalent. \square

This proposition is quite central to the rest of the development of the theory. The equivalent conditions (8) and (9) in particular have many applications. For one thing, they show that there is a wealth of prime ideals and prime filters. They also correspond to completeness, as mentioned above. This will be made more precise in Section 3.1.3 where we discuss in which sense a prime filter can be considered to be a model. In topological terms this equivalence can also be understood as the Hofmann-Mislove Theorem in disguise.

Remark. *Consistency provides the following bridge between filt and idl : A morphism $\vdash: L \rightarrow M$ induces a function $(\cdot)[\vdash]: \text{filt}(L) \rightarrow \text{filt}(M)$ because of Lemma 3.1.3. As we will see in Lemma 3.1.23 this can be used to turn filt into a functor, but at the moment we are not interested in composition. Analogously, we get a function $[\vdash](\cdot): \text{idl}(M) \rightarrow \text{idl}(L)$ which makes idl a contravariant functor.*

As consistency is a relation, let us, for the moment, regard them both as covariant functors to Rel , the category of sets and relations. Now, consider the diagram

$$\begin{array}{ccc}
 \text{filt}(L) & \xrightarrow{(\cdot)[\vdash]} & \text{filt}(M) \\
 \downarrow c_L & & \downarrow c_M \\
 \text{idl}(L) & \xleftarrow{[\vdash](\cdot)} & \text{idl}(M)
 \end{array}$$

where ‘normal’ arrows indicate functional relations and the crossed arrows c_L and c_M the relations induced by \Vdash_L - and \Vdash_M -consistency. In these terms we see that the equivalence “(4) \iff (5)” says, in effect, that consistency acts as a natural transformation between filt and idl .

3.1.2 Algebraisation of the Logic

Before we go on to the topological semantics of continuous sequent calculi and compatible consequence relations we approach the issue from a different angle. We want to find the analogue of a Lindenbaum algebra for our logic. To this end we state the basic definitions of the paper [JS96].

Definition 3.1.10. A *strong proximity lattice* $(A; \vee, \wedge, \perp, \top)$ is a distributive lattice together with a binary transitive relation \prec satisfying $\prec \circ \prec = \prec$; the algebraic structure given by the lattice and the approximation structure are

connected by the following four axioms:

$$\begin{aligned}
 (\vee - \prec) \quad & (\forall a \in A, M \subseteq_{\text{fin}} A) M \prec a \iff \bigvee M \prec a \\
 (\prec - \wedge) \quad & (\forall a \in A, M \subseteq_{\text{fin}} A) a \prec M \iff a \prec \bigwedge M \\
 (\prec - \vee) \quad & (\forall a, x, y) a \prec x \vee y \implies (\exists x', y') x' \prec x, y' \prec y \text{ and } a \prec x' \vee y' \\
 (\wedge - \prec) \quad & (\forall a, x, y) x \wedge y \prec a \implies (\exists x', y') x \prec x', y \prec y' \text{ and } x' \wedge y' \prec a
 \end{aligned}$$

We use $a \prec M$ to mean $a \prec m$, for all $m \in M$, and analogously for $M \prec a$.

Mappings between proximity lattices are certain relations.

Definition 3.1.11. A relation $G \subseteq A \times B$ between strong proximity lattices A and B is called *approximable* if it satisfies the following conditions:

$$\begin{aligned}
 (G - \prec) \quad & G \circ \prec_B = G \\
 (\prec - G) \quad & \prec_A \circ G = G \\
 (\vee - G) \quad & (\forall M \subseteq_{\text{fin}} A, b \in B) M (G) b \iff \bigvee M (G) b \\
 (G - \wedge) \quad & (\forall a \in A, M \subseteq_{\text{fin}} B) a (G) M \iff a (G) \bigwedge M \\
 (G - \vee) \quad & (\forall a \in A, M \subseteq_{\text{fin}} B) a (G) \bigvee M \implies \\
 & (\exists N \subseteq_{\text{fin}} A) a \prec_A \bigvee N \text{ and } (\forall n \in N)(\exists m \in M) n (G) m
 \end{aligned}$$

If a relation satisfies all conditions but $(G - \vee)$ we call it a *weak approximable relation*.

Note that strong proximity lattices and weak approximable relations are self-dual notions. This means that as in Chapter 2 we usually get away with proving only half of our assertions.

Sometimes it is useful to have an alternative formulation of $(G - \vee)$. But first we prove a technical lemma that we use several times in the following; it appears as Lemma 7 in [JS96].

Lemma 3.1.12. *If G is a weak approximable relation, then:*

1. $x (G) y \implies x (G) y \vee y'$,
2. $x (G) y, \quad x' (G) y' \implies x \vee x' (G) y \vee y'$.

Proof. In every lattice the equation $y = y \wedge (y \vee y')$ holds. Hence, the axiom $(G - \wedge)$ implies (1).

To prove (2) suppose $x (G) y$ and $x' (G) y'$. From (1) we get $x (G) y \vee y'$ and $y (G) y \vee y'$ which implies $x \vee y (G) y \vee y'$ because of $(\vee - G)$. \square

Lemma 3.1.13. *For a weak approximable relation G from A to B the condition $(G - \vee)$ is equivalent to the conjunction of the two implications*

$$\begin{aligned} a (G) \perp &\implies a \prec_A \perp \quad \text{and} \\ a (G) x \vee y &\implies (\exists x', y') (x' (G) x, y' (G) y \text{ and } a \prec_A x' \vee y'). \end{aligned}$$

Proof. Let us begin by observing that the first implication is an instance of $(G - \vee)$ where $M = \emptyset$.

Now suppose $(G - \vee)$ holds and we are given $a (G) x \vee y$. We find a set N as in the condition $(G - \vee)$ and define $N_x := \{n \in N \mid n (G) x\}$ and $N_y := \{n \in N \mid n (G) y\}$. We clearly have $N = N_x \cup N_y$ and thus can infer

$$a \prec \bigvee N = \bigvee (N_x \cup N_y) = (\bigvee N_x) \vee (\bigvee N_y).$$

Using $(\vee - G)$ we also get $\bigvee N_x (G) x$ and $\bigvee N_y (G) y$.

We show the converse by induction on $|M|$. The case $M = \emptyset$ has already been taken care of, and if M is a singleton then $(\prec - G)$ allows us to return a singleton as N .

For the induction step assume we have $a (G) \bigvee M$ and M can be written as $M = M_1 \cup M_2$ where M_1 and M_2 are proper subsets. Because of $a (G) \bigvee M = (\bigvee M_1) \vee (\bigvee M_2)$ we find elements m_1 and m_2 such that $a \prec m_1 \vee m_2$, $m_1 (G) \bigvee M_1$ and $m_2 (G) \bigvee M_2$.

The induction hypothesis yields sets N_1 and N_2 according to $(G - \vee)$; in particular we have $m_1 \prec \bigvee N_1$ and $m_2 \prec \bigvee N_2$. By the previous lemma this implies $m_1 \vee m_2 \prec (\bigvee N_1) \vee (\bigvee N_2) = \bigvee (N_1 \cup N_2)$ and because of $a \prec m_1 \vee m_2$ also $a \prec \bigvee (N_1 \cup N_2)$. The only thing remaining to be checked is that for all $n \in N_1 \cup N_2$ there is an $m \in M = M_1 \cup M_2$ such that $n (G) m$. But this is clear since the analogue condition relating N_1 and M_1 as well as N_2 and M_2 are satisfied by construction. \square

An immediate corollary of the lemma is that \prec is an approximable relation itself. Hence, strong proximity lattices together with approximable relations as well as with weak approximable relations form categories SPL and SPL_w , respectively. Composition is given by relational product \circ , and the orders of approximation \prec act as identities.

Now, we want to compare this with continuous sequent calculi and consequence relations. For a weak approximable relation G we set

$$\Gamma \vdash_G \Delta \quad :\iff \quad \bigwedge \Gamma (G) \bigvee \Delta$$

which, as we will see in a moment, defines a compatible consequence relation.

Lemma 3.1.14. *For weak approximable relations G and H and the order of approximation \prec on a strong proximity lattice the following hold:*

1. \vdash_G is a consequence relation;
2. $\vdash_{G \circ H} = \vdash_G \circ \vdash_H$; and
3. \vdash_{\prec} has interpolants.

Proof. The only rules for consequence relations that do not follow immediately from the respective axioms for approximable relations are (LV), (R \wedge) and (W). Let us consider (LV): As strong proximity lattices are in particular distributive we have $(\phi \wedge \bigwedge \Gamma) \vee (\phi \wedge \bigwedge \Gamma) = (\phi \vee \psi) \wedge \bigwedge \Gamma$ which together with $(\vee - G)$ shows both directions of (LV). Weakening follows from Lemma 3.1.12 and its dual.

Using the cut rule (Cut) the second claim of the proposition is trivial.

To see that \vdash_{\prec} has interpolants take a 'sequent' $\bigwedge \Gamma \prec \bigvee \Delta \vee \phi$. From $(\prec - \vee)$ we get ϕ' and δ' such that $\phi' \prec \phi$, $\delta' \prec \bigvee \Delta$ and $\bigwedge \Gamma \prec \delta' \vee \phi'$. Interpolating again we find a ϕ'' with $\phi' \prec \phi'' \prec \phi$, and then Lemma 3.1.12 shows $\bigwedge \Gamma \prec \delta' \vee \phi'' \prec \bigvee \Delta \vee \phi''$. \square

Corollary 3.1.15. *The translation $\vdash_{(\cdot)}$ defines a functor from $\text{SPL}_{\mathbf{w}}$ to MLS .*

Proof. By the previous lemma \vdash_{\prec} has interpolants, is closed under (Cut) and hence is a continuous sequent calculus. The compatibility of consequence relations \vdash_G that arise from weak approximable relations is again a consequence of $\vdash_{(\cdot)}$ preserving composition. \square

Conversely, we may wonder whether we can turn continuous sequent calculi into strong proximity lattices and compatible consequence relations into weak approximable relations. We begin by turning a continuous sequent calculus into one where the underlying algebra is a distributive lattice. To this end we factor a given $(2, 2, 0, 0)$ -algebra L by the equations of such a lattice, and it turns out that \Vdash is invariant with respect to these equations. To make this precise, let \equiv denote the least congruence such that L/\equiv is a distributive lattice and write $[\phi] := \{\psi \mid \phi \equiv \psi\}$ for its congruence classes. We now claim that if $\gamma_1 \equiv \gamma'_1, \dots, \gamma_n \equiv \gamma'_n$ and $\delta_1 \equiv \delta'_1, \dots, \delta_m \equiv \delta'_m$ then

$$\frac{\gamma_1, \dots, \gamma_n \Vdash \delta_1, \dots, \delta_m}{\gamma'_1, \dots, \gamma'_n \Vdash \delta'_1, \dots, \delta'_m}.$$

This is proved as follows. We take the rules that say that \equiv is a congruence, i.e. symmetry, transitivity etc. (see for example [MT92, Definition 5.2.7]), plus the ones for distributive lattices and then show by induction over the derivation of $\phi \equiv \phi'$ that

$$\frac{\phi, \Gamma \Vdash \Delta}{\phi', \Gamma \Vdash \Delta} \quad \text{and} \quad \frac{\Gamma \Vdash \Delta, \phi}{\Gamma \Vdash \Delta, \phi'}.$$

These verifications are a bit tedious but all straight-forward. Hence we restrict ourselves to the example of the distributive law to give the flavour:

$$\frac{\frac{\frac{\phi \wedge (\psi \vee \chi), \Gamma \Vdash \Delta}{\phi, \psi \vee \chi, \Gamma \Vdash \Delta} (\mathbf{L}\wedge)}{\phi, \psi, \Gamma \Vdash \Delta} (\mathbf{L}\wedge)}{\phi \wedge \psi, \Gamma \Vdash \Delta} (\mathbf{L}\wedge) \quad \frac{\frac{\phi, \chi, \Gamma \Vdash \Delta}{\phi \wedge \chi, \Gamma \Vdash \Delta} (\mathbf{L}\wedge)}{\phi \wedge \chi, \Gamma \Vdash \Delta} (\mathbf{L}\wedge)}{(\phi \wedge \psi) \vee (\phi \wedge \chi), \Gamma \Vdash \Delta} (\mathbf{L}\vee)$$

Note, that for this to work it is essential that the logical rules can be used in both directions. It is worth pointing out that it is not distributivity in particular but already the lattice laws where the backwards rules are needed.

These considerations allow us to define a relation \Vdash_{\equiv} on the quotient algebra L/\equiv unambiguously by setting:

$$[\phi_1], \dots, [\phi_n] \Vdash_{\equiv} [\psi_1], \dots, [\psi_m] \quad :\iff \quad \phi_1, \dots, \phi_n \Vdash \psi_1, \dots, \psi_m.$$

The axioms for a continuous sequent calculus are readily checked as they are directly inherited from \Vdash . This argument is very similar to the one in Lemma 2.2.4.

We now have most of the information needed to support the following:

Theorem 3.1.16. *The categories MLS and $\text{SPL}_{\mathbf{w}}$ are equivalent.*

Proof. From Lemma 3.1.14 we already know that there is a functor from $\text{SPL}_{\mathbf{w}}$ to MLS . For singletons we note that $\phi (G) \psi$ is equivalent to $\phi \vdash_G \psi$ which already implies that the functor is faithful.

This observation also tells us how to prove fullness. Given an arbitrary compatible consequence relation \vdash between objects that are in the image of the functor we define

$$a G_{\vdash} b \iff a \vdash b.$$

The relation G_{\vdash} satisfies $(G - \prec)$ and $(\prec - G)$ because of $\vdash = \vdash \circ \Vdash_{\prec}$ and $\Vdash_{\prec} \circ \vdash = \vdash$. Similarly, $(\vee - G)$ and $(G - \wedge)$ are direct consequences of $(\mathbf{L}\vee)$ and $(\mathbf{R}\wedge)$. Hence, G_{\vdash} is a weak approximable relation and we have $\Gamma \vdash_{G_{\vdash}} \Delta \iff \bigwedge \Gamma G_{\vdash} \bigvee \Delta \iff \bigwedge \Gamma \vdash \bigvee \Delta \iff \Gamma \vdash \Delta$.

To show that we have an equivalence of categories we have to prove that every isomorphism class of objects of MLS is hit by this functor. Here we use the considerations preceding this proposition: Given a continuous sequent calculus L we look at L/\equiv with \Vdash_{\equiv} as given there. To make sure that this object is in the image it suffices that \Vdash_{\equiv} restricted to singletons is a \prec -relation: The conditions $(\vee - \prec)$, $(\prec - \wedge)$ and $\prec \circ \prec = \prec$ follow immediately from $(\mathbf{L}\vee)$, $(\mathbf{L}\perp)$, $(\mathbf{R}\wedge)$, $(\mathbf{R}\top)$ and $\Vdash_{\equiv} = \Vdash_{\equiv} \circ \Vdash_{\equiv}$. For $(\prec - \vee)$

consider $\phi \Vdash \psi \vee \chi$. We apply backwards (RV) and get $\phi \Vdash \psi, \chi$. Interpolation now yields ψ', χ' such that $\psi' \Vdash \psi, \chi' \Vdash \chi$ and $\phi \Vdash \psi', \chi'$ which gives us $\phi \Vdash \psi' \vee \chi'$ by (RV). The remaining condition follows from duality.

We also have to prove that L and L/\equiv are isomorphic in MLS, a situation that is reminiscent of Proposition 2.2.5. We define maps between L and L/\equiv by

$$\begin{aligned} \Gamma \vdash [\delta_1], \dots, [\delta_n] & : \iff \Gamma \Vdash \delta_1, \dots, \delta_n \\ [\delta_1], \dots, [\delta_n] \vdash' \Gamma & : \iff \delta_1, \dots, \delta_n \Vdash \Gamma. \end{aligned}$$

Clearly, \vdash and \vdash' are well-defined compatible consequence relations, and we have $\vdash \circ \vdash' = \Vdash$ and $\vdash' \circ \vdash = \Vdash_{\equiv}$ since the original relation \Vdash satisfies $\Vdash \circ \Vdash = \Vdash$. \square

The proof also shows that MLS is equivalent to its full subcategory whose objects are distributive lattices.

In the light of this proposition we can say that compatible consequence relations are the appropriate extension of approximable relations from individual elements to finite sets. More importantly, the cut rule (Cut*) is the corresponding purely structural generalisation of the relational product.

Filters

The close connection between continuous sequent calculi and strong proximity lattices, in particular the translation giving rise to the previous theorem, can be extended to filters: If F is a filter in a continuous sequent calculus L we write $[F] := \{[\phi] \mid \phi \in F\}$ for its image under the quotient map to the distributive lattice L/\equiv . The next proposition says, among other things, that $[F]$ is a filter in the sense of [JS96]. There a filter is defined to be a subset X such that $X = \uparrow X$ and $M \subseteq_{\text{fin}} X \implies \bigwedge M \in X$, where \uparrow refers to the order of approximation \prec which in our case is merely \Vdash_{\equiv} restricted to singletons.

Proposition 3.1.17. *For $F \in \text{filt}(L)$ the following hold:*

1. F is closed under \equiv , i.e. it is a (disjoint) union of \equiv -equivalence classes;
2. $[F]$ is a filter in the strong proximity lattice L/\equiv ; and
3. the map $[\cdot]$ is an order isomorphism between $\text{filt}(L)$ and the strong proximity lattice filters in L/\equiv .

Proof. If $\psi \equiv \phi \in F$ we know from roundness of filters (Proposition 3.1.1) that there is a $\chi \in F$ such that $\chi \Vdash \phi$. From the discussion preceding Theorem 3.1.16 we know that this implies $\chi \Vdash \psi$ and hence $\psi \in F$.

As a consequence of this and Proposition 3.1.1 we get that $[F]$ is a strong proximity lattice filter.

Given such a filter \mathcal{F} in L/\equiv we set $F_{\mathcal{F}} := [\cdot]^{-1}[\mathcal{F}] = \{\phi \mid [\phi] \in \mathcal{F}\}$. Again by Proposition 3.1.1 it is clear that this defines an element of $\text{filt}(L)$. This map as well as $[\cdot]$ are clearly monotone with respect to subset inclusion, so it remains to show that they are mutually inverse. On the one hand we have $[\phi] \in [F_{\mathcal{F}}] \iff \phi \in F_{\mathcal{F}} \iff [\phi] \in \mathcal{F}$ and on the other $\phi \in F_{[F]} \iff [\phi] \in [F] \iff \phi \in F$, where the last equivalence is true because of (1). \square

As $[\cdot]$ is an order isomorphism it takes prime filters to the \wedge -prime filters in the corresponding strong proximity lattice. In the distributive lattice L/\equiv ‘ordinary’ filters and those of strong proximity lattices coincide because of the previous proposition. Hence, Proposition 3.1.6 tells us that these \wedge -primes are exactly the prime filters as defined in [JS96].

We could thus use this observation to transfer the topological results from that paper to continuous sequent calculi. However, in the interest of keeping this thesis reasonably self-contained we prove the relevant results directly in the next section. Actually, this does not take much longer than quoting the results from [JS96] by translating them via Proposition 3.1.17: We want to give explicit constructions in terms of continuous sequent calculi anyway, and moreover we are studying a larger class of morphisms.

3.1.3 Topological Semantics

We now embark on a first investigation of the structure of the posets $\text{filt}(L)$ and $\text{idl}(L)$. From Lemma 3.1.4 we already know that they are dcpo’s and \wedge -semilattices.

Lemma 3.1.18. *The assignments $X \mapsto X[\Vdash_L]$ and $X \mapsto [\Vdash_L]X$ are Scott-continuous retractions on the power set of the continuous sequent calculus L .*

Proof. Both assignments are obviously monotone. Scott-continuity follows immediately from the fact that sequents are relations between *finite* sets, and idempotence is a consequence of the Lemma 3.1.3. \square

This shows that much more structure is inherited from $\mathfrak{P}(L)$:

Corollary 3.1.19. *The posets $\text{filt}(L)$ and $\text{idl}(L)$ are continuous lattices.*

Proof. The previous lemma shows that $\text{filt}(L)$ and $\text{idl}(L)$ are continuous retracts [AJ94, Proposition 3.1.7.1] of the algebraic lattice $(\mathfrak{P}(L), \subseteq)$. Hence they are continuous lattices [AJ94, Proposition 3.1.2, Theorem 3.1.4]. \square

As the next step we give an internal characterisation of the order of approximation in these lattices. Any filter F in a continuous sequent calculus is round and hence can be written as

$$F = \bigcup \{ \phi[\Vdash] \mid \phi \in F \}.$$

We now claim that this union is directed: Given $\phi, \psi \in F$ we have $\phi \wedge \psi \in F$ by Proposition 3.1.1 and the containment $(\phi \wedge \psi)[\Vdash] \supseteq \phi[\Vdash], \psi[\Vdash]$ follows readily from (W) and (L \wedge). This allows us to characterise \ll for filters and ideals:

Lemma 3.1.20. *In the continuous lattice $\text{filt}(L)$ we have $F \ll G$ if and only if there is a $\phi \in G$ such that $F \subseteq \phi[\Vdash]$.*

Proof. We know from Lemma 3.1.4 that directed suprema are simply given by unions. Hence, the condition is clearly sufficient for $F \ll G$. That it is also necessary follows from our discussion preceding this lemma. \square

The lemma gives us an alternative proof for Corollary 3.1.19 that does not rely on facts about retractions: The discussion preceding the lemma shows that every filter is the directed union of filters way below it. Hence, $\text{filt}(L)$ is continuous. To see that it is a complete lattice note that $\emptyset[\Vdash]$ is clearly the smallest filter of L . As we know that $\text{filt}(L)$ has directed suprema we only have to construct binary suprema. We can read off the construction from Lemma 3.1.18, but let us state it explicitly for future reference:

Lemma 3.1.21. *The supremum of two filters F and G is given by $(F \cup G)[\Vdash]$.*

Proof. We already know that $(\cdot)[\Vdash]$ is monotone. As F and G are filters we have $F[\Vdash] = F$ and $G[\Vdash] = G$ which implies that $(F \cup G)[\Vdash]$ contains F and G . Also because of monotonicity it is the least such filter. \square

These two lemmata enable us to infer even more about the structure of the lattices of filters and ideals:

Proposition 3.1.22. *For a continuous sequent calculus L the lattices $\text{filt}(L)$ and $\text{idl}(L)$ are arithmetic.*

Proof. We begin by proving that $\text{filt}(L)$ is distributive. For this it suffices to verify $F \wedge (G \vee H) \subseteq (F \wedge G) \vee (F \wedge H)$ for filters F, G and H ; the other containment is true in any lattice. We can write the left hand side as $F \cap (G \cup H)[\Vdash]$ because of the previous lemma. So let us suppose we have $\phi_1, \dots, \phi_n \in G \cup H$ and $\phi_1, \dots, \phi_n \Vdash \psi$ for a $\psi \in F$. As F is a filter we can find a $\chi \in F$ such that $\chi \Vdash \psi$ using roundness of F . Applying (L \vee) multiple times yields $\phi_1 \vee \chi, \dots, \phi_n \vee \chi \Vdash \psi$ and by Proposition 3.1.1 we know that all the formulae $\phi_i \vee \chi$ come either from $(F \cap G)$ or $(F \cap H)$. This means that we have shown $\psi \in ((F \cap G) \cup (F \cap H))[\Vdash] = (F \wedge G) \vee (F \wedge H)$.

Suppose F, G and H are filters and $F \ll G, H$. By Lemma 3.1.20 this means that we can find $\phi \in G, \psi \in H$ such that $F \subseteq \phi[\Vdash], \psi[\Vdash]$. The infimum of G and H is given by the intersection because of Lemma 3.1.4, and by Proposition 3.1.1 we have $\phi \vee \psi \in G \cap H$. Finally, Lemma 3.1.5.(2) allows us to conclude

$$F \subseteq \phi[\Vdash] \cap \psi[\Vdash] = (\phi \vee \psi)[\Vdash]$$

which shows $F \ll G \cap H$ by the previous lemma.

The fact that the top element L is compact is an immediate consequence of $L = \perp[\vdash]$. \square

We know from Proposition 1.3.25 that arithmetic lattices are the Stone duals of stably compact spaces. Hence, the previous proposition is a first glimpse at the topological meaning of a continuous sequent calculus. At the moment we can describe the space as $\text{pt}(\text{idl}(L))$, but we will find a much more economical description of the space. Before we continue in this vein we study how we can turn idl and filt into functors since we are also interested in the meaning of compatible consequence relations.

Morphisms

Lemma 3.1.23. *Let L and M be continuous sequent calculi and $\vdash: L \rightarrow M$ a consequence relation. Then $(\cdot)[\vdash]$ is a Scott-continuous \wedge -semilattice homomorphism from $\text{filt}(L)$ to $\text{filt}(M)$.*

Moreover, this assignment is functorial, i.e. $(\cdot)[\vdash \circ \vdash'] = (\cdot)[\vdash'] \circ (\cdot)[\vdash]$.

Proof. We know from Lemma 3.1.3 that $(\cdot)[\vdash]$ takes filters to filters and the map is clearly monotone. By Lemma 3.1.4 the supremum of a directed set of filters is just the union, and as \vdash relates finite sets the function $(\cdot)[\vdash]$ is Scott-continuous.

We have $L[\vdash] = M$ since $\perp \in L$ and \vdash satisfies the rules (L \perp) and (W). For binary meets it is clear that $(F \cap G)[\vdash] \subseteq F[\vdash] \cap G[\vdash]$ by monotonicity. For the other direction take a $\phi \in F[\vdash] \cap G[\vdash]$, i.e. there are $\psi \in F$ and $\chi \in G$ such that $\psi \vdash \phi$ and $\chi \vdash \phi$. From this we get

$$\frac{\psi \vdash \phi \quad \chi \vdash \phi}{\phi \vee \chi \vdash \phi} \text{ (LV)}$$

and because of Proposition 3.1.1 we know that $\psi \vee \chi \in F \cap G$. This shows $\phi \in (F \cap G)[\vdash]$.

Finally, we have to show $\text{filt}(\vdash \circ \vdash') = \text{filt}(\vdash') \circ \text{filt}(\vdash)$ —note that relational composition is from left to right whereas concatenation of functions from right to left. We know that for any filter F the sets $F[\vdash]$, $(F[\vdash])[\vdash']$ and $F[\vdash \circ \vdash']$ are filters. Because of Corollary 3.1.2 and the original formulation of the cut rule (Cut) we can thus restrict our considerations to singletons:

$$\begin{aligned} \phi \in (F[\vdash])[\vdash'] &\iff (\exists \psi, \chi) \psi \in F \text{ and } \psi \vdash \chi \text{ and } \chi \vdash' \phi \\ &\iff (\exists \psi) \psi \in F \text{ and } \psi (\vdash \circ \vdash') \phi \\ &\iff \phi \in F[\vdash \circ \vdash'] \end{aligned}$$

\square

We introduce the auxiliary category ASL (arithmetic semilattices) whose objects are arithmetic lattices and whose morphisms are Scott-continuous semilattice morphisms. The category of arithmetic lattices is a non-full subcategory as not every continuous semilattice morphism is a frame morphism; we will come back to the exact connection later. ASL is a full subcategory of the category of complete semilattices that we have already seen in the remark following Corollary 1.3.12.

The previous lemma almost shows that filt can be considered as a functor from MLS to ASL. The only thing that we have not shown, yet, is that it preserves identities. Fortunately, this is trivial as $(\cdot)[\vdash]$ is by definition the identity of filters. Using some of the information we have gained so far, in particular Proposition 3.1.9, we can actually prove a much stronger result:

Theorem 3.1.24. *The functor $\text{filt}: \text{MLS} \rightarrow \text{ASL}$ is full and faithful. Dually, idl is a contravariant, full and faithful filter from MLS to ASL.*

Proof. Assume \vdash and \vdash' are two compatible consequence relations from L to M such that $F[\vdash] = F[\vdash']$ for all filters $F \subseteq L$. We start by observing that, because of weakening, for two *finite* sets Γ, Δ the condition $\Gamma \vdash \Delta$ is equivalent to (Γ, Δ) being \vdash -inconsistent. Hence, using Proposition 3.1.9, we can conclude

$$\begin{aligned} \Gamma \vdash \Delta &\iff (\Gamma, \Delta) \text{ is } \vdash\text{-inconsistent} \iff (\Gamma[\Vdash_L], \Delta) \text{ is } \vdash\text{-inconsistent} \\ &\iff (\Gamma[\Vdash_L][\vdash], \Delta) \text{ is } \Vdash_M\text{-inconsistent.} \end{aligned}$$

Now $\Gamma[\Vdash_L][\vdash]$ is equal to $\Gamma[\Vdash_L][\vdash']$ by assumption and turning the above argument upside down we see that $\Gamma \vdash \Delta$ is equivalent to $\Gamma \vdash' \Delta$. Hence the functor filt is faithful.

To show that it is also full take any Scott-continuous semilattice homomorphism $f: \text{filt}(L) \rightarrow \text{filt}(M)$. We have to find a consequence relation \vdash_f such that $\text{filt}(\vdash_f) = f$. Looking at the chain of equivalences of the previous paragraph we see that we are more or less forced to define

$$\Gamma \vdash_f \Delta \quad :\iff \quad (f(\Gamma[\Vdash_L]), \Delta) \Vdash_M\text{-inconsistent.}$$

To show that this is a compatible consequence relation we have to check the conditions of Corollary 2.1.13.

Weakening is obvious and so are the rules (L \perp), (R \top) concerning the constants. To check (LV) assume $\phi, \Gamma \vdash_f \Delta$ and $\psi, \Gamma \vdash_f \Delta$. Because of Proposition 3.1.9 we can find formulae $\alpha \in f((\phi, \Gamma)[\Vdash_L]) \cap [\Vdash_M]\Delta$ and $\beta \in f((\psi, \Gamma)[\Vdash_L]) \cap [\Vdash_M]\Delta$. By Proposition 3.1.1 the element $\alpha \vee \beta$ lies in $[\Vdash_M]\Delta$ as well as in

$$\begin{aligned} f((\phi, \Gamma)[\Vdash_L]) \cap f((\psi, \Gamma)[\Vdash_L]) &= f((\phi, \Gamma)[\Vdash_L] \cap (\psi, \Gamma)[\Vdash_L]) \\ &= f((\phi \vee \psi, \Gamma)[\Vdash_L]) \end{aligned}$$

where the last equality follows from Lemma 3.1.5.(2). Thus, we have shown $\phi \vee \psi, \Gamma \vdash_f \Delta$.

For (L-Int') suppose $\phi, \Gamma \vdash_f \Delta$ which again means that we can find an element $\theta \in f((\phi, \Gamma)[\Vdash_L]) \cap [\Vdash_M]\Delta$. By Lemma 3.1.5.(5) we get

$$(\phi, \Gamma)[\Vdash_L] = \bigcup_{\psi \in \phi[\Vdash_L]}^{\uparrow} (\psi, \Gamma)[\Vdash_L]$$

and as f is continuous there is a ψ such that $\phi \Vdash_L \psi$ and $\theta \in f((\psi, \Gamma)[\Vdash_L])$ which shows $\psi, \Gamma \vdash_f \Delta$.

The admissibility of (L-Cut) is a consequence Lemma 3.1.5.(3). The right rules (R \wedge), (R-Int') and (R-Cut) are proved using the same techniques; they follow from (R \wedge), (R-Int) and (Cut') for \Vdash_M , respectively.

It remains to verify that $\text{filt}(\vdash_f) = f$. For a filter $F \subseteq L$ we calculate

$$\begin{aligned} F[\vdash_f] &= \{ \phi \mid (\exists \Gamma \subseteq_{\text{fin}} F) \Gamma \vdash_f \phi \} \\ &= \{ \phi \mid (\exists \Gamma \subseteq_{\text{fin}} F) (f(\Gamma[\Vdash]), \phi) \Vdash_M \text{-inconsistent} \} \\ &= \{ \phi \mid (\exists \Gamma \subseteq_{\text{fin}} F) f(\Gamma[\Vdash]) \cap [\Vdash_M]\phi \neq \emptyset \} \\ &= \{ \phi \mid (\exists \Gamma \subseteq_{\text{fin}} F) \phi \in f(\Gamma[\Vdash]) \} \\ &= \bigcup_{\Gamma \subseteq_{\text{fin}} F}^{\uparrow} f(\Gamma[\Vdash]) \\ &= f\left(\bigcup_{\Gamma \subseteq_{\text{fin}} F}^{\uparrow} \Gamma[\Vdash]\right) \\ &= f(F) \end{aligned}$$

using Proposition 3.1.9 to get from the second line to the third and the fact that $f(\Gamma[\Vdash])$ is a filter to get to the next. \square

The theorem implies that the image of filt is equivalent to MLS, and as we shall later see the image is itself equivalent to all of ASL. Whereas the category MLS is clearly self-dual, this property is not obvious for ASL. It was originally discovered by Jimmie Lawson, [Law79].

Compact Saturated Sets

We return to the topological spaces that arise from a continuous sequent calculus L . We know that both $\text{idl}(L)$ and $\text{filt}(L)$ are arithmetic lattices and hence correspond to stably compact spaces. If we consider the category ASL, we see that some of the morphisms there are actually frame morphisms and thus can be thought of as continuous functions via Stone duality. This means that we should think of ASL as a category of frames rather than locales. Because of this it is more natural to take the contravariant functor idl to define the open sets of the space corresponding to a continuous sequent calculus.

If we make this choice, what is the meaning of $\text{filt}(L)$? In the light of Corollary 1.3.12 and the surrounding discussion it is not far fetched to surmise that it is the co-compact topology, or equivalently, the compact saturated sets of the topology given by $\text{idl}(L)$.

Coming back to Theorem 1.2.12 we see that a compact saturated set is essentially the same thing as a Scott-open filter in the topology. Given a filter F we consider

$$\mathcal{F} := \{I \in \text{idl}(L) \mid I \cap F \neq \emptyset\}.$$

It is clear from the definition that \mathcal{F} is a Scott-open subset of $\text{idl}(L)$ as directed suprema are simply directed unions. To see that it is filtered suppose $I, J \in \mathcal{F}$ which is witnessed by $\phi \in F \cap I$ and $\psi \in F \cap J$. Because of Proposition 3.1.1 we get $\phi \wedge \psi \in F \cap I \cap J$ which shows that $I \cap J \in \mathcal{F}$.

Now, let us assume that we are given a Scott-open filter $\mathcal{F} \subseteq \text{idl}(L)$. We define

$$F := \{\phi \mid [\Vdash]\phi \in \mathcal{F}\}$$

and verify that F is a filter by checking the conditions of Proposition 3.1.1: We trivially have $\top \in F$ as filters must not be empty and this in turn implies $L = [\Vdash]\top \in \mathcal{F}$. For $[\Vdash]\phi, [\Vdash]\psi \in \mathcal{F}$ we get $[\Vdash](\phi \wedge \psi) = [\Vdash]\phi \cap [\Vdash]\psi$ by the dual of Lemma 3.1.5.(2), and this intersection lies in the filter \mathcal{F} . Given $[\Vdash]\phi \in \mathcal{F}$ and $\phi \Vdash \psi$ we get $[\Vdash]\phi \subseteq [\Vdash]\psi$ and hence $[\Vdash]\psi \in \mathcal{F}$. Finally, take $[\Vdash]\phi \in \mathcal{F}$. From the dual of Lemma 3.1.5.(5) we know that we can write it as $[\Vdash]\phi = \bigvee^\uparrow \{[\Vdash]\psi \mid \psi \Vdash \phi\}$ which implies the existence of a $\psi \Vdash \phi$ such that $\psi \in F$ since we assumed that \mathcal{F} is Scott-open.

Proposition 3.1.25. *For a continuous sequent calculus L the Scott-open filters of $\text{idl}(L)$ are order-isomorphic to $\text{filt}(L)$.*

Proof. The two translations that we have just discussed are clearly monotone. So, all that we have to show is that they are mutually inverse: Consider

$$F \mapsto \{I \in \text{idl}(L) \mid I \cap F \neq \emptyset\} \mapsto \{\phi \mid [\Vdash]\phi \cap F \neq \emptyset\}.$$

Because of the roundness of filters the result is simply F .

Now let \mathcal{F} be a Scott-open filter. We construct

$$\mathcal{F} \mapsto \{\phi \mid [\Vdash]\phi \in \mathcal{F}\} \mapsto \{I \mid (\exists \phi \in I) [\Vdash]\phi \in \mathcal{F}\}$$

and claim that the result is \mathcal{F} . Suppose we have $[\Vdash]\phi \in \mathcal{F}$ for a $\phi \in I$, then $[\Vdash]\phi \subseteq I$ and hence $I \in \mathcal{F}$. For the other containment assume $I \in \mathcal{F}$. By the argument preceding Lemma 3.1.20 we can write I as a directed union of ideals $[\Vdash]\phi$ for $\phi \in I$. As \mathcal{F} is Scott-open this implies that there is a $\phi \in I$ such that $[\Vdash]\phi \in \mathcal{F}$. \square

The Spectrum

We can use this result to get an alternative descriptions of the ‘points’ of $\text{idl}(L)$, i.e. the completely prime filters. Every supremum can be written as a directed supremum of finite suprema. So, a Scott-open filter is completely prime if and only if it is inaccessible by finite suprema.

Theorem 3.1.26. *Let L be a continuous sequent calculus. The order-isomorphism of Proposition 3.1.25 restricts and co-restricts to prime filters in L and completely prime filters in $\text{idl}(L)$.*

Proof. The prime filters of L are precisely the \cap -prime elements in $\text{filt}(L)$. We already have an order-isomorphism between $\text{filt}(L)$ and the Scott-open filters of $\text{idl}(L)$. Hence, all that we have to show is that such a filter is completely prime if and only if it is prime with respect to finite intersections of Scott-open filters.

By Lemma 1.2.15 a completely prime filter is even prime with respect to arbitrary intersections of upper sets.

Suppose conversely that $\mathcal{P} \subseteq \text{idl}(L)$ is a Scott-open filter that is \cap -prime. Our considerations preceding this proposition show that we only have to check that \mathcal{P} is inaccessible by binary suprema. So let us assume $I \vee J \in \mathcal{P}$. As $\text{idl}(L)$ is a continuous lattice we can write this as a directed supremum of elements $I' \vee J'$ where $I' \ll I$ and $J' \ll J$. Now, \mathcal{P} is Scott-open which implies that there are such ideals I' and J' that satisfy $I' \vee J' \in \mathcal{P}$. We have $\uparrow I' \cap \uparrow J' = \uparrow(I' \vee J') \subseteq \mathcal{P}$. Since $\text{idl}(L)$ is arithmetic the Scott-open sets $\uparrow I'$ and $\uparrow J'$ are also filters and we can infer that either of them must be contained in \mathcal{P} . This shows that either I or J is an element of \mathcal{P} and hence that \mathcal{P} is completely prime. \square

The following lemma is a consequence of the equivalent formulations of consistency given in Section 3.1.1. It can be seen as a reformulation of the Hofmann-Mislove Theorem, in the form of the key Lemma 1.2.11, and has many applications.

Lemma 3.1.27. *If I is an ideal which does not meet a filter F in a continuous sequent calculus, then there is a prime filter $P \supseteq F$ such that $P \cap I = \emptyset$. Consequently, every filter is the intersection of the prime filters that contain it.*

Proof. The first statement is an immediate consequence of Proposition 3.1.9 as $F \cap I = \emptyset$ is equivalent to (F, I) being \Vdash -consistent.

For the second claim take an arbitrary $\phi \notin F$. This implies $\Vdash \phi \cap F = \emptyset$ because of $F = F[\Vdash]$. Hence there is a prime filter $P \supseteq F$ such that $\Vdash \phi \cap P = \emptyset$. As $P = P[\Vdash]$ it cannot contain ϕ , and thus $F = \bigcap_{F \subseteq P} P$. \square

By Proposition 3.1.25 a filter $F \subseteq L$ corresponds to the Scott-open filter $\{J \in \text{idl}(L) \mid J \cap F \neq \emptyset\}$ in $\text{idl}(L)$. Hence, the lemma says, in effect, that

if an ideal I does not lie in this Scott-open filter then there is a prime filter $P \supseteq F$ such that I is not an element of the Scott-open filter corresponding to P , either. This is exactly the statement of Lemma 1.2.11, the key ingredient for the proof of the Hofmann-Mislove theorem.

We define the *spectrum* of a continuous sequent calculus L to be

$$\text{spec}(L) := \{P \in \text{filt}(L) \mid P \text{ prime}\}.$$

Theorem 3.1.26 tells that this is essentially the topological space $\text{pt}(\text{idl}(L))$. To complete the description of these spaces in terms of prime filters of L we also have to translate the topology: In $\text{pt}(\text{idl}(L))$ an open set is given by an ideal $I \in \text{idl}(L)$ and contains the ‘points’ $\{\mathcal{P} \in \text{pt}(\text{idl}(L)) \mid I \in \mathcal{P}\}$. We apply the translations given before Proposition 3.1.25 and get:

Corollary 3.1.28. *An ideal $I \in \text{idl}(L)$ corresponds to the open set of ‘points’*

$$\mathcal{O}_I := \{P \in \text{spec}(L) \mid P \cap I \neq \emptyset\}.$$

Proof. A prime filter P corresponds to $\mathcal{P} := \{J \in \text{idl}(L) \mid J \cap P \neq \emptyset\}$. We can read off immediately that $I \in \mathcal{P}$ is equivalent to $I \cap P \neq \emptyset$. This shows that \mathcal{O}_I corresponds to $\{\mathcal{P} \in \text{pt}(\text{idl}(L)) \mid I \in \mathcal{P}\}$. \square

We can also describe compact saturated sets in this way. A filter F corresponds to the Scott-open filter $\{I \in \text{idl}(L) \mid I \cap F \neq \emptyset\}$. By the Hofmann-Mislove Theorem (1.2.12) this in turn corresponds to

$$\begin{aligned} & \left\{ \mathcal{P} \in \text{pt}(\text{idl}(L)) \mid \{I \in \text{idl}(L) \mid I \cap F \neq \emptyset\} \subseteq \mathcal{P} \right\} \\ & = \left\{ \mathcal{P} \in \text{pt}(\text{idl}(L)) \mid (\forall I \in \text{idl}(L)) F \cap I \neq \emptyset \implies I \in \mathcal{P} \right\}. \end{aligned}$$

Corollary 3.1.29. *A filter $F \in \text{filt}(L)$ describes the following compact saturated set of ‘points’:*

$$\mathcal{K}_F := \{P \in \text{spec}(L) \mid P \supseteq F\}$$

Proof. As in the previous proof we consider a prime filter P and its translation \mathcal{P} . If we have $F \subseteq P$ and $F \cap I \neq \emptyset$, then we get $P \cap I \neq \emptyset$ and hence $I \in \mathcal{P}$.

For the other containment suppose $F \not\subseteq P$. Then we can find a $\phi \in F \setminus P$ and by roundness another $\psi \in F$ such that $\psi \Vdash \phi$. This shows that the ideal $[\Vdash]\phi$ meets F , but by construction it does not meet P . We conclude $[\Vdash]\phi \notin \mathcal{P}$ which implies that \mathcal{P} is not a member of the set of completely prime filters corresponding to F . \square

In Section 1.3 we saw that for a stably compact space X the co-compact topology X_κ is also stably compact, and that the topology for the latter is given by the complements of compact saturated sets from X . We can now express this in terms of the underlying continuous sequent calculus:

As we observed several times in Chapter 2 the setup of our logical system is symmetric: We can go from a continuous sequent calculus L to the dual structure L^{op} by turning around the \Vdash , and interchanging \wedge and \vee as well as \top and \perp . Furthermore, this process turns filters into ideals and vice versa. In addition Proposition 3.1.8 tells us that L^{op} has essentially the same ‘points’ as L . The connection to the topological dualisation of taking the co-compact topology is given by the following:

Theorem 3.1.30. *Let L be a continuous sequent calculus. Then we have*

$$(\text{spec}(L))_\kappa \cong \text{spec}(L^{\text{op}})$$

and this homeomorphism is the function given by the pseudo-complementation of Proposition 3.1.8.

Proof. Because of $\text{idl}(L^{\text{op}}) = \text{filt}(L)$ the first statement is an immediate consequence of Proposition 3.1.25 and Corollary 1.3.12.

For the second statement we only have to show the following: For a filter F pseudo-complementation translates \mathcal{K}_F into the complement of \mathcal{O}_F , considering F as an ideal in L^{op} . Take $F \in \text{filt}(L)$, $P \in \text{spec}(L)$ and define $P' := [\Vdash](L \setminus P)$ according to Proposition 3.1.8. As this P' is the largest ideal that is disjoint from P , we have $F \subseteq P \iff F \cap P' = \emptyset$. This shows the claim. \square

Logic

A formula ϕ gives rise to an ideal $[\Vdash]\phi$ and a filter $\phi[\Vdash]$. So we can ask about the *open* and the *compact* reading of such a formula. We use \mathcal{O}_ϕ and \mathcal{K}_ϕ as abbreviations for $\mathcal{O}_{[\Vdash]\phi}$ and $\mathcal{K}_{\phi[\Vdash]}$, respectively. Since filters are round we express \mathcal{O}_ϕ more economically as

$$\mathcal{O}_\phi = \{P \in \text{spec}(L) \mid P \cap [\Vdash]\phi \neq \emptyset\} = \{P \in \text{spec}(L) \mid \phi \in P\}.$$

We can interpret this in the following way: As in propositional logic, a prime filter on a continuous sequent calculus represents a model. The spectrum $\text{spec}(L)$, then, is the space of all models, and every formula ϕ of L defines a subset \mathcal{O}_ϕ of models, namely, those in which ϕ is true. The definition of the topology on $\text{spec}(L)$ is such that all these extents of formulae are open. We will see shortly that the \mathcal{O}_ϕ are in fact a basis of this topology. In the classical setting of Boolean algebras and Stone spaces the extents are also compact. This is not the case here. However, every formula

has the canonical compact subset \mathcal{K}_ϕ associated with it. The logic is translated into set-theoretic operations both through the open and the compact interpretation:

Proposition 3.1.31. *The following are true for a continuous sequent calculus L :*

$$\begin{aligned} \mathcal{O}_{\phi \wedge \psi} &= \mathcal{O}_\phi \cap \mathcal{O}_\psi & \mathcal{O}_{\phi \vee \psi} &= \mathcal{O}_\phi \cup \mathcal{O}_\psi \\ \mathcal{K}_{\phi \wedge \psi} &= \mathcal{K}_\phi \cap \mathcal{K}_\psi & \mathcal{K}_{\phi \vee \psi} &= \mathcal{K}_\phi \cup \mathcal{K}_\psi \end{aligned}$$

Proof. The two equations for \mathcal{O} follow from Proposition 3.1.1 and 3.1.6, and the remaining claims are then a consequence of Theorem 3.1.30. \square

We summarise the different concepts in a table which complements the one given in Section 1.2.2:

Logic	Spectrum	Dual Spectrum
formula	open and compact reading	open and compact reading
prime filter/ideal	point	point
ideal	open set	compact saturated set
filter	compact saturated set	open set

There are many more connections between the logic and the topological interpretation:

Proposition 3.1.32. *Let L be a continuous sequent calculus and $\phi, \psi \in L$ formulae. Then the following properties hold:*

1. $\mathcal{O}_\phi \subseteq \mathcal{K}_\phi$;
2. $\phi \Vdash \psi \iff \mathcal{K}_\phi \subseteq \mathcal{O}_\psi$; and
3. for a compact saturated set $K \subseteq \text{spec}(L)$ and an open $O \subseteq \text{spec}(L)$ satisfying $K \subseteq O$ there is a formula $\phi \in L$ such that $K \subseteq \mathcal{O}_\phi \subseteq \mathcal{K}_\phi \subseteq O$.

Proof. Property (1) and one direction of (2) can be read off directly from the description of \mathcal{O} and \mathcal{K} . If we have $\mathcal{K}_\phi \subseteq \mathcal{O}_\psi$, this means that all prime filters $P \supseteq \phi[\Vdash]$ meet the ideal $[\Vdash]\psi$. By Lemma 3.1.27 this implies $\phi[\Vdash] \cap [\Vdash]\psi \neq \emptyset$ and hence $\phi \Vdash \psi$.

As in the previous argument, we can translate $K \subseteq O$ into $F \cap I \neq \emptyset$ for the ideal I and the filter F corresponding to K and O , respectively. We can pick a $\phi \in F \cap I$ and get $K = \mathcal{K}_F \subseteq \mathcal{O}_\phi \subseteq \mathcal{K}_\phi \subseteq \mathcal{O}_I = O$; the last inclusion follows from the roundness of I . \square

The proposition shows in particular that the sequent $\phi \Vdash \psi$ implies $\mathcal{O}_\phi \ll \mathcal{O}_\psi$ and $\mathcal{K}_\psi \ll \mathcal{K}_\phi$ (see Lemma 1.2.17).

Remark. *The condition $\phi \Vdash \psi$ is strictly stronger than the conjunction of $[\Vdash]\phi \ll [\Vdash]\psi$ and $\psi[\Vdash] \ll \phi[\Vdash]$ —a counterexample is easily found in the continuous sequent calculus constructed in Proposition 3.1.34. This*

matches up with the theory of abstract bases for domains [AJ94, Definition 2.2.20]: If (B, \prec) is such a basis then $x \prec y$ implies $\downarrow x \ll \downarrow y$, but not vice versa. In fact, \Vdash restricted to singletons is an abstract basis, and the continuous sequent calculus ideals are precisely the ideals for this abstract basis. Unlike the general case, the relation \Vdash has a concrete meaning for individual tokens of the basis as given by the preceding proposition.

The third condition of the proposition implies that there are ‘enough’ formulae to generate the topology:

Corollary 3.1.33. *For a continuous sequent calculus L the sets \mathcal{O}_ϕ form a basis for the topology on $\text{spec}(L)$.*

Proof. The third property of the previous proposition applies in particular in the situation $\uparrow x \subseteq O$. \square

From Spaces to Continuous Sequent Calculi

There is one obvious question that we have not tackled, yet:

Which stably compact spaces arise as spectra of continuous sequent calculi?

As we will see, up to isomorphism, the answer is: all of them.

Given a stably compact space X we know that $\Omega(X)$ and $\mathcal{K}(X)$ are arithmetic lattices. Hence, it is quite natural to define a strong proximity lattice from them first:

Proposition 3.1.34. *For a stably compact space X the following defines a strong proximity lattice:*

- $L_X := \{ \langle O, K \rangle \in \Omega(X) \times \mathcal{K}(X) \mid O \subseteq K \}$
- $\langle O, K \rangle \vee \langle O', K' \rangle := \langle O \cup O', K \cup K' \rangle$
- $\langle O, K \rangle \wedge \langle O', K' \rangle := \langle O \cap O', K \cap K' \rangle$
- $\perp := \langle \emptyset, \emptyset \rangle, \top := \langle X, X \rangle$
- $\langle O, K \rangle \prec \langle O', K' \rangle \quad :\iff \quad K \subseteq O'$

Proof. It is obvious that L_X is a distributive lattice and that $\prec \circ \prec \subseteq \prec$. For the converse conclusion assume $K \subseteq O'$. As X is in particular locally compact, the neighbourhood filter of K has a basis of compact saturated sets which means that there is such a set L that satisfies $K \subseteq \text{int}(L) \subseteq L \subseteq O'$. This shows that \prec is interpolative, i.e. $\prec \subseteq \prec \circ \prec$.

Verifying the two axioms $(\vee - \prec)$ and $(\prec - \wedge)$ is trivial since all involved operations are set theoretic.

For $(\prec - \vee)$ let K be a subset of $O_1 \cup O_2$. Since X is locally compact this implies that for every point $x \in K$ we can find a compact neighbourhood K_x such that either $K_x \subseteq O_1$ or $K_x \subseteq O_2$. The interiors of the K_x cover K and by compactness we find a finite sub-cover. We collect them in two finite unions depending on which of the two open sets they are a subset of. The two resulting compact sets and their interiors give us the required interpolating elements.

If we replace X by X_κ we get an L_{X_κ} which is just the dual of L_X . This means that we can think of the axiom $(\wedge - \prec)$ as $(\prec - \vee)$ for the co-compact topology which implies that it also holds. \square

Section 3.1.1, and in particular the discussion of consistency, helps to explain the meaning of the open and compact sets O and K making up the tokens $\langle O, K \rangle$ in this proposition. An open set O can be seen to represent positive information, and a compact set K to represent the negative information $X \setminus K$. The constraint $O \subseteq K$ avoids self-contradiction of tokens.

The results from Section 3.1.2 show that by slight abuse of notation we can think of L_X as a continuous sequent calculus where \prec is just \Vdash restricted to singletons. We now have to check that L_X actually represents X :

Theorem 3.1.35. *If X is a stably compact space, then we have $\text{spec}(L_X) \cong X$.*

Proof. We map an ideal $I \in \text{idl}(L_X)$ to the set

$$\bigcup \{O \mid (\exists K) \langle O, K \rangle \in I\}$$

and as a union of open sets this set is clearly open. Conversely, we take an open set O to

$$\{\langle O', K' \rangle \in L_X \mid K' \subseteq O\}.$$

We verify that this is an ideal using Proposition 3.1.1: The first two conditions are trivial, and the argument for the third one is the same as for $\prec \subseteq \prec \circ \prec$ in the proof of the previous Proposition.

The next step is to check that these two mappings are mutually inverse. Beginning with an open set we see that we get it back because X is locally compact. Starting with an ideal I we get

$$\begin{aligned} I &\mapsto \bigcup \{O \mid (\exists K) \langle O, K \rangle \in I\} \\ &\mapsto \left\{ \langle O', K' \rangle \mid K' \subseteq \bigcup \{O \mid (\exists K) \langle O, K \rangle \in I\} \right\}. \end{aligned}$$

For every element in the original filter $\langle O, K \rangle \in I$ we can find a token $\langle O', K' \rangle \in I$ that satisfies $\langle O, K \rangle \prec \langle O', K' \rangle$, or in other words $K \subseteq O'$. This proves that $\langle O, K \rangle$ appears in the resulting ideal.

For the other containment take $O' \subseteq K' \subseteq \bigcup \{O \mid (\exists K) \langle O, K \rangle \in I\}$. Since K' is compact and I as an ideal is closed under finite suprema there must be a $\langle O, K \rangle \in I$ such that $K' \subseteq O$. This implies $\langle O', K' \rangle \prec \langle O, K \rangle$ and hence $\langle O', K' \rangle \in I$.

The two translations are clearly monotone which shows $\text{idl}(L) \cong \Omega(X)$. From this we immediately get $\text{spec}(L) \cong \text{pt}(\text{idl}(L)) \cong \text{pt}(\Omega(X)) \cong X$. \square

A token $\langle O, K \rangle$ has an open and a compact reading. We might expect that they are essentially O and K which is indeed the case: The token gives rise to the ideal

$$[\Vdash]\langle O, K \rangle = \{\langle O', K' \rangle \in L_X \mid K' \subseteq O\},$$

precisely the ideal corresponding to O in the above proof.

It also gives rise to a filter $\langle O, K \rangle[\Vdash]$ in L_X which corresponds to the Scott-open filter $\{I \in \text{idl}(L_X) \mid \langle O, K \rangle \in I\}$ by roundness. Every ideal I in it stands for a concrete open set $\llbracket I \rrbracket := \bigcup \{O' \mid (\exists K') \langle O', K' \rangle \in I\}$. By the Hofmann-Mislove Theorem it suffices to show the equality of the following two Scott-open filters:

$$\{O' \in \Omega(X) \mid O' \supseteq K\} \stackrel{!}{=} \{\llbracket I \rrbracket \mid \langle O, K \rangle \in I\}$$

If $O' \supseteq K$, then we have $\langle O, K \rangle \in [\Vdash]\langle O', X \rangle$, and as we already know $\llbracket [\Vdash]\langle O', X \rangle \rrbracket = O'$ we get $O' \in \{\llbracket I \rrbracket \mid \langle O, K \rangle \in I\}$. For $\langle O, K \rangle \in I$ we get a token $\langle O', K' \rangle \in I$ satisfying $K \subseteq O'$ which implies $\llbracket I \rrbracket \supseteq K$.

As an immediate consequence of the previous theorem we get that the functor $\text{filt}: \text{MLS} \rightarrow \text{ASL}$ reaches all objects up to isomorphism.

Corollary 3.1.36. *The categories MLS and ASL are equivalent.*

Proof. This follows directly from the previous theorem, Theorem 3.1.24 and the fact that the stably compact spaces are precisely the Stone duals of the arithmetic lattices (Proposition 1.3.25). \square

3.1.4 Semantics of Morphisms

Given a compatible consequence relation \vdash between continuous sequent calculi L and M it is natural to consider the function

$$(\cdot)[\vdash]: P \mapsto P[\vdash].$$

By Lemma 3.1.3 the result is a filter, whatever the P . But even if P is a prime filter there is no reason why $P[\vdash]$ should be prime.

As an example consider the relation that comprises all sequents. This is always a compatible consequence relation and for all $P \subseteq L$ we get $P[\vdash] = M$ which is never prime as it contains \perp .

Morphisms as Multi-Functions

We can, however, consider $(\cdot)[\vdash]$ as a function from $\text{spec}(L)$ to $\text{filt}(M)$. Filters in M correspond to compact saturated subsets of $\text{spec}(M)$, but if we want to give a functional interpretation of morphisms we have to turn $\text{filt}(M)$ into a topological space in its own right. The poset $(\text{filt}(M), \subseteq) \cong (\mathcal{K}(\text{spec}(M)), \supseteq)$ is an arithmetic lattice. Endowed with the Scott topology it is the so called *Smyth* or *upper power domain*, though usually the top element \emptyset is omitted. We can directly describe its topology from the continuous sequent calculus:

Lemma 3.1.37. *Let L be a continuous sequent calculus. Then the sets*

$$\hat{\mathcal{O}}_\phi := \{F \in \text{filt}(L) \mid \phi \in F\}$$

form a basis for the Scott topology on $(\text{filt}(L), \subseteq)$.

Proof. By Lemma 3.1.4 directed suprema are unions and hence the sets $\hat{\mathcal{O}}_\phi$ are Scott-open.

Assume that a filter F lies in a basic open set $\uparrow F'$. Then we can find an interpolating filter F'' satisfying $F' \ll F'' \ll F$ and by Lemma 3.1.20 there is a $\phi \in F$ such that $F'' \subseteq \phi[\vdash]$. This implies $F \in \hat{\mathcal{O}}_\phi \subseteq \uparrow F'' \subseteq \uparrow F'$. \square

Using the terminology of Proposition 1.2.19 we can also understand the elements of this basis as $\hat{\mathcal{O}}_\phi = \square \mathcal{O}_\phi$.

Before we tackle the correspondence of consequence relations and functions $\text{spec}(L) \rightarrow \text{filt}(M)$ we provide some useful facts for the proof of the following proposition:

Lemma 3.1.38. *For a continuous sequent calculus L and formulae $\phi, \psi \in L$ the following hold:*

1. $\phi \Vdash \psi \implies \mathcal{K}_\phi \subseteq \mathcal{K}_\psi, \hat{\mathcal{O}}_\phi \subseteq \hat{\mathcal{O}}_\psi;$
2. $\mathcal{K}_{\phi \wedge \psi} \subseteq \mathcal{K}_\phi, \hat{\mathcal{O}}_\phi \subseteq \hat{\mathcal{O}}_{\phi \vee \psi};$ *and*
3. $\hat{\mathcal{O}}_\phi = \bigcup^{\uparrow}_{\psi \Vdash \phi} \hat{\mathcal{O}}_\psi.$

Proof. The first two properties follow immediately from Proposition 3.1.31, 3.1.32 and 3.1.1. The union in (3) is directed because of (LV) and (2). Finally, we get the equality of the two expressions from the fact that filters are round. \square

Proposition 3.1.39. *Let L and M be continuous sequent calculi. The compatible consequence relations from L to M are in a bijection with the continuous maps from $\text{spec}(L)$ to $\mathcal{K}(\text{spec}(M)) \cong \text{filt}(M)$.*

Proof. Let us begin by checking that the function $(\cdot)[\vdash]$ is continuous. To this end we look at the preimage of a basic open set

$$\begin{aligned} ((\cdot)[\vdash])^{-1}[\hat{O}_\psi] &= \{P \in \mathbf{spec}(L) \mid P[\vdash] \in \hat{O}_\psi\} \\ &= \{P \in \mathbf{spec}(L) \mid \psi \in P[\vdash]\} \\ &= \{P \in \mathbf{spec}(L) \mid (\exists \phi \in P) \phi \vdash \psi\} \quad (\text{Corollary 3.1.2}) \\ &= \bigcup_{\phi \vdash \psi} \mathcal{O}_\phi, \end{aligned}$$

which is obviously open.

Next we show that the map $(\vdash \mapsto (\cdot)[\vdash])$ is injective. Suppose $\vdash \neq \vdash'$, then by Theorem 3.1.24 we have a filter $F \subseteq L$ such that $F[\vdash] \neq F[\vdash']$. Without loss of generality, we can assume that there is a $\psi \in F[\vdash] \setminus F[\vdash']$. This implies that F does not meet the ideal $[\vdash']\psi$ and because of Lemma 3.1.27 there is a prime filter $P \supseteq F$ such that $P \cap [\vdash']\psi = \emptyset$. Because of Corollary 3.1.2 this means $\psi \notin P[\vdash']$, but note that we do have $\psi \in P[\vdash]$. So, $(\cdot)[\vdash]$ and $(\cdot)[\vdash']$ are different maps even when we restrict the domain to prime filters.

For surjectivity take a continuous function $f: \mathbf{spec}(L) \rightarrow \mathbf{filt}(M)$. We set

$$\Gamma \vdash_f \Delta \quad :\iff \quad \mathcal{K}_{\wedge \Gamma} \subseteq f^{-1}[\hat{O}_{\vee \Delta}]$$

and show that it is a consequence relation. Permutation of Γ and Δ —the implicit structural rule—follows from Proposition 3.1.17.(1), weakening from Lemma 3.1.38.(2).

The only non-trivial logical rules are (L \vee) and (R \wedge). The former follows from Proposition 3.1.31 since we can calculate

$$\begin{aligned} \mathcal{K}_{(\phi \vee \psi) \wedge \Gamma} &= \mathcal{K}_{\phi \vee \psi} \cap \mathcal{K}_{\wedge \Gamma} = (\mathcal{K}_\phi \cup \mathcal{K}_\psi) \cap \mathcal{K}_{\wedge \Gamma} \\ &= (\mathcal{K}_\phi \cap \mathcal{K}_{\wedge \Gamma}) \cup (\mathcal{K}_\psi \cap \mathcal{K}_{\wedge \Gamma}) = \mathcal{K}_{\phi \wedge \Gamma} \cup \mathcal{K}_{\psi \wedge \Gamma}. \end{aligned}$$

For (R \wedge) we have to be more subtle since $\hat{O}_{\phi \wedge \psi} = \hat{O}_\phi \cap \hat{O}_\psi$ holds, as a consequence of Proposition 3.1.1, but $\hat{O}_{\phi \vee \psi} = \hat{O}_\phi \cup \hat{O}_\psi$ fails as \hat{O} does not refer to prime filters only. Nonetheless, if F is a filter we know from Proposition 3.1.17.(1) that $\phi \vee (\psi \wedge \chi) \in F$ is equivalent to $(\phi \vee \psi) \wedge (\phi \vee \chi) \in F$ because of $\phi \vee (\psi \wedge \chi) \equiv (\phi \vee \psi) \wedge (\phi \vee \chi)$. This suffices to verify that $\mathcal{K}_{\wedge \Gamma} \subseteq f^{-1}[\hat{O}_{\vee \Delta \vee \phi}], f^{-1}[\hat{O}_{\vee \Delta \vee \psi}]$ if and only if $\mathcal{K}_{\wedge \Gamma} \subseteq f^{-1}[\hat{O}_{\vee \Delta \vee (\phi \wedge \psi)}]$.

Having shown that \vdash_f is a consequence relation we now check that it is compatible: From property (1) of the previous lemma it is clear that we have $\Vdash_L \circ \vdash_f \circ \Vdash_M \subseteq \vdash_f$.

If, conversely, we have $\mathcal{K}_{\wedge \Gamma} \subseteq f^{-1}[\hat{O}_{\vee \Delta}]$, then we apply Proposition 3.1.32.(3) and find a formula $\phi \in L$ such that

$$\mathcal{K}_{\wedge \Gamma} \subseteq \mathcal{O}_\phi \subseteq \mathcal{K}_\phi \subseteq f^{-1}[\hat{O}_{\vee \Delta}].$$

This implies $\Gamma \Vdash \phi$ by Proposition 3.1.32.(2) and backwards (L \wedge).

We also want to interpolate on the right. To this end we observe that by property (3) of the previous lemma we have $\hat{O}_{\vee \Delta} = \bigcup^{\uparrow}_{\psi \Vdash \vee \Delta} \hat{O}_{\psi}$. Taking preimages preserves directed suprema and as \mathcal{K}_{ϕ} is compact, one of these ψ satisfies $\mathcal{K}_{\phi} \subseteq f^{-1}[\hat{O}_{\psi}]$. This shows $\Gamma \Vdash_L \phi \vdash_f \psi \Vdash_M \Delta$ and hence $\vdash_f = \Vdash_L \circ \vdash_f \circ \Vdash_M$ which is clearly compatible.

The last thing that remains to be done is to verify that \vdash_f is mapped to f . For a prime filter P we calculate:

$$\begin{aligned}
 P[\vdash_f] &= \{\psi \mid (\exists \phi \in P) \phi \vdash_f \psi\} && \text{(Corollary 3.1.2)} \\
 &= \{\psi \mid (\exists \phi \in P) \mathcal{K}_{\phi} \subseteq f^{-1}[\hat{O}_{\psi}]\} \\
 &= \{\psi \mid P \in f^{-1}[\hat{O}_{\psi}]\} && (\dagger) \\
 &= \{\psi \mid f(P) \in \hat{O}_{\psi}\} \\
 &= \{\psi \mid \psi \in f(P)\} \\
 &= f(P)
 \end{aligned}$$

The step to (\dagger) follows from the continuity of f and Proposition 3.1.32.(3). \square

We can reformulate this further: For a stably compact space X the Smyth power domain $(\mathcal{K}(X), \supseteq)$ is an arithmetic lattice. Every continuous lattice is an FS domain when equipped with the Scott topology. Hence, we can apply Proposition 1.3.19 and see that $\mathcal{K}(X)$ is again stably compact. A continuous function between stably compact spaces lifts to a mapping between the compact set lattices: The lifted function takes a compact saturated set to the saturation of its image, a Scott-continuous mapping by Lemma 1.3.16. This lifting is functorial which can be easily seen from the localic description of Proposition 1.2.16. Consequently, \mathcal{K} defines an endofunctor on StCp , the category of stably compact spaces.

This functor is also part of a monad: Its ‘unit’ takes a point x to $\uparrow x$, its upper closure with respect to the specialisation order. The ‘multiplication’ of the monad takes a Scott-compact family of compact saturated sets $\mathcal{C} \in \mathcal{K}(\mathcal{K}(X))$ to $\bigcup \mathcal{C} \in \mathcal{K}(X)$ which is again compact by Lemma 1.2.20. For the full details that this defines a monad see [Sch93, Proposition 7.21]. We will show in a moment that the category of compatible theories is exactly the Kleisli category of this monad.

Some aspects of the following argument are more transparent if we start with semilattices rather than continuous sequent calculi. For this reason we relate the previous proposition and Lemma 3.1.23. The latter tells us that a compatible consequence relation $\vdash: L \rightarrow M$ gives rise to a Scott-continuous semilattice homomorphism from $[\vdash](\cdot): \text{idl}(M) \rightarrow \text{idl}(L)$. We claim that taking preimages under this map corresponds to the function

$(\cdot)[\vdash]: \text{spec}(L) \rightarrow \text{filt}(M)$ via Proposition 3.1.25. A prime filter $P \subseteq L$ corresponds to the completely prime filter $\{I \in \text{idl}(L) \mid I \cap P \neq \emptyset\}$ and its preimage under $[\vdash](\cdot)$ is

$$\{I \in \text{idl}(M) \mid P \cap [\vdash]I \neq \emptyset\} = \{I \in \text{idl}(M) \mid P[\vdash] \cap I \neq \emptyset\},$$

where the equality follows from Proposition 3.1.9. The resulting Scott-open filter is just the translation of $P[\vdash]$.

Theorem 3.1.40. *The categories MLS , SPL_w and ASL are equivalent to the Kleisli category $\text{StCp}_{\mathcal{K}}$ of the Smyth power monad $(\mathcal{K}, \uparrow, \cup)$.*

Proof. Because of Theorem 3.1.35 we already know that we reach all objects up to isomorphism. Moreover, most of the work for the morphism has already been done in Proposition 3.1.39. We only have to show that the translation appearing in the proof there is functorial. Then we know that the functor is full, faithful and hence part of an equivalence.

First, we check identities: $P[\uparrow] = P$, but now this prime filter, seen as a compact saturated set rather than a point of the spectrum, corresponds to the saturation $\{Q \in \text{spec}(L) \mid P \subseteq Q\}$. This means that the identity \uparrow is mapped to the unit of the monad which is the identity in the Kleisli category.

For composition we consider the composition of two semilattice homomorphisms ρ and σ :

$$A \xrightarrow{\rho} B \xrightarrow{\sigma} C$$

$$\begin{array}{ccccc}
 & & \text{pt}(A) & & \text{pt}(B) & & \text{pt}(C) \\
 & & \swarrow r & & \downarrow \eta & & \swarrow s \\
 & & \mathcal{K}(\text{pt}(A)) & \leftarrow \dots \leftarrow & \mathcal{K}(\text{pt}(B)) & & \\
 & & \uparrow \mu & & \swarrow \mathcal{K}(r) & & \\
 & & \mathcal{K}(\mathcal{K}(\text{pt}(A))) & & & &
 \end{array}$$

In the diagram the continuous functions r and s correspond to ρ and σ , respectively, and from the preceding discussion we know that they are essentially given by taking preimages under ρ and σ . The Kleisli composition

of r and s is given by $\mu \circ \mathcal{K}(r) \circ s$, where the ‘multiplication’ μ of the monad is union.

We think of $\mathcal{K}(\text{pt}(B))$ and $\mathcal{K}(\text{pt}(A))$ as spaces of Scott-open filters of B and A by Theorem 1.2.12. The space $\mathcal{K}(\mathcal{K}(\text{pt}(A)))$, however, we prefer to consist of the compact saturated subsets of $\mathcal{K}(\text{pt}(A))$. In these terms we can describe $\mathcal{K}(r)$ as follows: It takes a Scott-open Filter $F \subseteq B$, which corresponds to $\{P \in \text{pt}(B) \mid P \supseteq F\}$, to the saturation of

$$\{\rho^{-1}[P] \mid P \in \text{pt}(B) \text{ and } P \supseteq F\}.$$

Normally, the multiplication μ is taking the union of compact saturated sets. In terms of Scott-open filters this translates into intersection, as an argument very similar to the proof of Corollary 1.2.14 shows. If we compose this with $\mathcal{K}(r)$ we can disregard the saturation of the compact set as this does not add anything new to this intersection. Using Lemma 3.1.27 we can calculate

$$\mu\left(\mathcal{K}(r)(F)\right) = \bigcap \{\rho^{-1}[P] \mid P \in \text{pt}(B) \text{ and } P \supseteq F\} = \rho^{-1}[F].$$

From this observation it follows that for all $P \in \text{pt}(C)$ we have

$$(\sigma \circ \rho)^{-1}[P] = \rho^{-1}[\sigma^{-1}[P]] = (\mu \circ \mathcal{K}(r) \circ s)(P)$$

which concludes the proof of the theorem. \square

Morphisms as Relations

We can alternatively describe these morphisms as certain relations between stably compact spaces. For any set theoretic function $f: X \rightarrow \mathfrak{P}(Y)$ we can define a binary relation

$$R_f := \{\langle x, y \rangle \mid y \in f(x)\} \subseteq X \times Y.$$

Conversely, every relation $R \subseteq X \times Y$ determines a function

$$f_R(x) := \{y \mid x R y\}: X \rightarrow \mathfrak{P}(Y),$$

and these translations are mutually inverse. Categorically speaking, the category of sets and relations is the Kleisli category for the power set monad on Set . We now tackle the analogue in our topological setting.

Proposition 3.1.41. *Let X and Y be stably compact spaces. For a continuous function $f: X \rightarrow \mathcal{K}(Y) \subseteq \mathfrak{P}(Y)$ the relation R_f is closed in $X \times Y_\kappa$.*

Proof. Suppose $x \not R_f y$ which is equivalent to $y \notin f(x)$. The set $f(x) \subseteq Y$ is compact saturated, and as Y is locally compact its neighbourhood filter has a basis of compact saturated sets. So there is such a set K satisfying

$$y \notin K \supseteq \text{int}(K) \supseteq f(x).$$

Since f is continuous and the set $\{L \in \mathcal{K}(Y) \mid L \subseteq \text{int}(K)\}$ is open by Lemma 1.2.17 there is an open neighbourhood V of x such that for all $x' \in V$ we have $f(x') \subseteq \text{int}(K) \subseteq K$. This implies $(V \times (Y \setminus K)) \cap R_f = \emptyset$. The set $V \times (Y \setminus K)$ is open and contains $\langle x, y \rangle$ thus showing that R is closed. \square

Unfortunately, the converse is slightly more involved. We will make use of the fact that we can freely translate between a stably compact space X , the stably compact space arising from the co-compact topology X_κ and the compact ordered space X_π as discussed in Section 1.3.2.

Proposition 3.1.42. *Let X and Y be stably compact and $R \subseteq X \times Y_\kappa$ a closed relation. Then f_R is a continuous function from X to $\mathcal{K}(Y)$.*

Proof. First, we show that for all $x \in X$ the set $f_R(x) = \{y \mid x R y\}$ is compact saturated: The specialisation order of the product topology is the product of the two orders, i.e. point-wise. Now R is closed and hence a lower set, and we see that $f_R(x)$ is a lower subset of Y_κ . By Proposition 1.1.2 this is equivalent to $f_R(x)$ being an upper, or equivalently saturated, subset of Y .

The topology on $X_\pi \times Y_\pi$ is finer than that on $X \times Y_\kappa$. Thus, R is a closed subset of the compact Hausdorff space $X_\pi \times Y_\pi$. Moreover, the subspace $f_R(x) \subseteq Y_\pi$ is homeomorphic to the closed set $(\{x\} \times Y) \cap R \subseteq X_\pi \times Y_\pi$. We conclude that it is compact in the patch topology and consequently also in the coarser original topology. Putting the two observations together we see that $f_R(x)$ is compact saturated.

Now assume that f_R is not continuous. Because of Lemma 1.2.17 this implies that there is an $x \in X$ and an open set $V \subseteq Y$ such that $f_R(x) \not\subseteq V$, but for all neighbourhoods U of x there is an element $x' \in U$ for which $f_R(x') \subseteq V$. We can reformulate this as $K_U := (U \times (Y \setminus V)) \cap R \neq \emptyset$. As X is locally compact we can restrict ourselves to compact saturated neighbourhoods U of x . The sets $Y \setminus V$ are patch-compact which implies that all the sets K_U are patch-compact, as well. They are also clearly filtered which means that they form a filter basis of compact non-empty sets in the compact Hausdorff space $X_\pi \times Y_\pi$. This allows us to find an element $\langle \tilde{x}, \tilde{y} \rangle \in \bigcap K_U \neq \emptyset$. By construction we have $\tilde{x} R \tilde{y}$ and $x \sqsubseteq \tilde{x}$ which implies $x R \tilde{y}$ because R as a closed set is also a lower set. But we also have $\tilde{y} \in Y \setminus V$ which contradicts $f_R(x) \subseteq V$. This shows that f_R is continuous. \square

The last two propositions already characterise the relations that arise from MLS-morphisms. The only thing that remains to be done is to investigate how composition behaves under this translation.

Note that for the function $(\cdot)[\vdash]$ we can write $R_{\vdash} := R_{(\cdot)[\vdash]}$ as

$$P R_{\vdash} Q \iff P[\vdash] \subseteq Q$$

using the correspondence of Corollary 3.1.29.

Theorem 3.1.43. *The category MLS is equivalent to the category of stably compact spaces with closed relations $R \subseteq X \times Y_{\kappa}$. The composition is given by the ordinary relational product.*

Proof. Suppose $\vdash: L \rightarrow M$ and $\vdash': M \rightarrow N$ are compatible consequence relations. We need to show:

$$R_{\vdash \circ \vdash'} = R_{\vdash} \circ R_{\vdash'}$$

Suppose $O \subseteq L$ and $Q \subseteq N$ are two prime filters such that $O[\vdash \circ \vdash'] = O[\vdash][\vdash'] \subseteq Q$, the equality being a consequence of Lemma 3.1.23. Then we must have $O[\vdash] \cap [\vdash'](N \setminus Q) = \emptyset$ and using Proposition 3.1.9 we find a prime filter $P \subseteq M$ such that $O[\vdash] \subseteq P$ and $P \cap [\vdash'](N \setminus Q) = \emptyset$. As P is a filter, Corollary 3.1.2 applies and thus we have $P[\vdash] \subseteq Q$.

For the other containment assume $O[\vdash] \subseteq P$ and $P[\vdash'] \subseteq Q$. This immediately implies $O[\vdash \circ \vdash'] = O[\vdash][\vdash'] \subseteq P[\vdash'] \subseteq Q$. \square

We now have several descriptions of essentially the same category. Syntactically we introduced it as MLS and showed that it is equivalent to strong proximity lattices $\text{SPL}_{\mathbf{w}}$, a formulation that is very close to Abramsky's prelocales. From Corollary 3.1.36 we know that we can also see it as arithmetic semilattices with Scott-continuous semilattice morphisms; this is the localic viewpoint. Topologically we can understand it either as the Kleisli category $\text{StCp}_{\mathcal{K}}$ or as stably compact spaces with closed (in the appropriate sense) relations, the category that we will call StCp^* . Note that in this category the identities are given by the specialisation order.

Having several different ways of seeing this category has the advantage that for any construction we want to perform in it, we can choose the most convenient point of view for that particular purpose.

Remark. *It is instructive to consider the self-duality of this category in its different manifestations. We already discussed the case MLS in Section 2.1.1: It flips a consequence relation $\vdash: L \rightarrow M$ around and interchanges the connectives \vee and \perp with their duals \wedge and \top . The same procedure takes care of strong proximity lattices.*

For arithmetic semilattices the self-duality is given by exponentiation $\mathbf{2}^{(\cdot)}$. The elements of $\mathbf{2}^A$ correspond precisely to the Scott-open filters in A and hence to the compact saturated subsets of the Stone dual by the Hofmann-Mislove Theorem. Because of Corollary 1.3.12 this gives rise to an involution on ASL. The key ingredients for the proof that it does indeed correspond to the dualisation in MLS are contained in the discussion preceding Theorem 3.1.40.

The hardest case is that of the Kleisli category. Given $f: X \rightarrow \mathcal{K}(Y)$ we claim that the dual map $f^: Y \rightarrow \mathcal{K}(X)$ is given by*

$$f^*(y) := \{x \in X \mid y \in f(x)\}.$$

For this to work it suffices to verify that for a compatible consequence relation $\vdash: L \rightarrow M$ and $P \in \text{spec}(M^{\text{op}})$ we have

$$\{Q \in \text{spec}(L^{\text{op}}) \mid Q \supseteq [\vdash]P\} \stackrel{!}{=} \{Q \in \text{spec}(L^{\text{op}}) \mid Q^*[\vdash] \subseteq P^*\},$$

where $(\cdot)^*$ denotes the pseudo-complementation given in Proposition 3.1.8 and Theorem 3.1.30. To this end let $Q \supseteq [\vdash]P$, $\phi \in Q^*$ and $\phi \vdash \psi$. This implies $\psi \notin P$ since $Q \cap Q^* = \emptyset$, and because of interpolation on the right we get $\psi \in P^*$. The argument for the other containment is almost identical.

As an immediate consequence we get $y \in f^*(x) \iff x \in f(y)$ which means that the dual of a closed relation $R \subseteq X \times Y_\kappa$ is simply R again.

Functions

In general, consequence relations correspond to relations between stably compact spaces, but some of them are really functions in disguise: If X is a stably compact space, then we can think of it and $\mathcal{K}(X)$ as $\text{spec}(L)$ and $\text{filt}(L)$ for the same continuous sequent calculus. We see from the bases of the respective topologies given in Corollary 3.1.33 and Lemma 3.1.37 that X is a subspace of $\mathcal{K}(X)$. Consequently, we can compose a continuous function from $\text{spec}(L)$ to $\text{spec}(M)$ with the embedding $\text{spec}(M) \hookrightarrow \text{filt}(M)$. Conversely, if a function $f: \text{spec}(L) \rightarrow \text{filt}(M)$ can be co-restricted to the subspace $\text{spec}(M)$ then this is also a continuous function.

This is readily rephrased in terms of closed relations. Given a continuous function $f: X \rightarrow Y$ between stably compact spaces we can consider the *hyper-graph*

$$\{\langle x, y \rangle \in X \times Y \mid f(x) \sqsubseteq y\}.$$

If $R \subseteq X \times Y_\kappa$ is a closed relation with the property that for all x there is a least y such that $x R y$ we can consider the function that maps x to y . These translations correspond directly to the ones given in the previous paragraph using Proposition 3.1.41, 3.1.42 and Corollary 3.1.29. Hence, it is *a priori* clear that the hypergraph is a closed relation, that the other translation yields a continuous function and that they are mutually inverse.

For completeness sake we also look at the situation in ASL. The discussion before Theorem 3.1.40 can be summed up as follows: A Scott-continuous semilattice morphism $\rho: A \rightarrow B$ corresponds to a continuous function $r: \text{pt}(B) \rightarrow \mathcal{K}(\text{pt}(A))$. It is given by taking preimages if we identify $\mathcal{K}(\text{pt}(A))$ with the Scott-open filters of A . Hence, we can certainly think of a ρ as a ‘function’ if it is in fact a frame morphism.

If it is not, then there are $x, y \in A$ such that $\rho(x) \vee \rho(y) < \rho(x \vee y)$. The Scott topology on any continuous domain has a basis of Scott-open filters (see [AJ94, Lemma 2.3.8] or the proof of Proposition 1.2.18). This implies,

together with Lemma 1.2.11, that there is a completely prime filter $P \subseteq B$ such that $\rho(x \vee y) \in P$, but $\rho(x) \vee \rho(y) \notin P$. The preimage $\rho^{-1}[P]$ contains $x \vee y$ but neither x nor y and hence is not a completely prime filter. We have shown: A Scott-continuous semilattice homomorphism corresponds to a ‘function’ if and only if it is a frame morphism.

The problem we are really interested in is to give a syntactic characterisation of these semantic notions of *being a function*:

Theorem 3.1.44. *A compatible consequence relation $\vdash: L \rightarrow M$ corresponds to a ‘function’ in the above sense if and only if for all $\Gamma \vdash \Delta$ there is a set $\Theta \subseteq_{\text{fin}} L$ such that $\Gamma \Vdash_L \Theta$ and for all $\phi \in \Theta$ there is a $\psi \in \Delta$ with $\phi \vdash \psi$.*

Proof. For a prime filter $P \subseteq L$ we check that $P[\vdash]$ is also prime using Proposition 3.1.6. Assume that $P[\vdash]$ contains \perp . We get $\phi \in P$ such that $\phi \vdash \perp$ and hence $\phi \vdash$. The condition of the theorem implies $\phi \Vdash$ which contradicts the primeness of P .

Now let us suppose $P \ni \phi \vdash \psi \vee \chi$. We take the right hand side apart to get $\phi \vdash \psi, \chi$ which implies that there is a set $\Delta \subseteq_{\text{fin}} L$ such that $\phi \Vdash \Delta$ and for all $\delta \in \Delta$ we have $\delta \vdash \psi$ or $\delta \vdash \chi$. Because of $P \ni \phi \Vdash \Delta$ we see that $\bigvee \Delta$ lies in $P = P[\Vdash]$ and as P is prime this implies that there is a $\delta \in \Delta \cap P$. Hence, we either have $\psi \in P[\vdash]$ or $\chi \in P[\vdash]$.

For the converse we study the relation \vdash_f that arises from a continuous function $f: \text{spec}(L) \rightarrow \text{filt}(M)$ that can actually be co-restricted to $\text{spec}(M)$. According to the proof of Proposition 3.1.39 $\Gamma \vdash_f \Delta$ is given by $\mathcal{K}_{\wedge \Gamma} \subseteq f^{-1}[\hat{\mathcal{O}}_{\bigvee \Delta}]$. Since we can actually consider f as a map to the subspace $\text{spec}(M)$ we can write this as $\mathcal{K}_{\wedge \Gamma} \subseteq f^{-1}[\mathcal{O}_{\bigvee \Delta}]$. The advantage of this is that $\mathcal{O}_{(\cdot)}$ preserves \bigvee whereas $\hat{\mathcal{O}}_{(\cdot)}$ does not.

Given $\Gamma \vdash_f \delta_1, \dots, \delta_n$ we get

$$\mathcal{K}_{\wedge \Gamma} \subseteq f^{-1}[\mathcal{O}_{\delta_1 \vee \dots \vee \delta_n}] = f^{-1}\left[\bigcup_{i=1}^n \mathcal{O}_{\delta_i}\right] = \bigcup_{i=1}^n f^{-1}[\mathcal{O}_{\delta_i}]$$

and by Proposition 3.1.32.(3) this means that for all $x \in \mathcal{K}_{\wedge \Gamma}$ we can find an $i \in \{1, \dots, n\}$ and a formula $\phi \in L$ such that $x \in \mathcal{O}_{\phi} \subseteq \mathcal{K}_{\phi} \subseteq f^{-1}[\mathcal{O}_{\delta_i}]$. As $\mathcal{K}_{\wedge \Gamma}$ is compact, finitely many of these \mathcal{O}_{ϕ} cover it. Taking the disjunctions of the formulae that correspond to each δ_i we get $\phi_1, \dots, \phi_n \in L$ such that $\mathcal{K}_{\wedge \Gamma} \subseteq \mathcal{O}_{\phi_1} \cup \dots \cup \mathcal{O}_{\phi_n} = \mathcal{O}_{\phi_1 \vee \dots \vee \phi_n}$ and $\mathcal{K}_{\phi_i} \subseteq f^{-1}[\mathcal{O}_{\delta_i}]$, for $i = 1, \dots, n$. This is just a translation of the condition stated in the theorem. \square

We now have explicit descriptions of ‘functions’ in all of the equivalent categories MLS, ASL, $\text{StCp}_{\mathcal{X}}$ and StCp with closed relations. Later we will see how we can characterise the subcategory of functions in purely categorical terms.

Note that the identities are ‘functions’. For the semantic categories this is obvious, and hence it is also true for the syntactic category MLS where it is a consequence of interpolation.

3.2 Domain Constructions

We are now in a position where we can perform domain constructions in logical form. Our first step is to improve on our representation theorem for stably compact spaces by characterising all continuous sequent calculi that correspond to a given stably compact space. This allows us to make the connection between syntactic and topological constructions. We can also turn the argument in the characterisation around to construct continuous sequent calculi.

Then we look at several domain constructions to show how this technique can be applied. In particular we consider lifting, sums, products, the Smyth power domain, the function and the relation space. There are only partial results about the functions space construction. We contrast this with the space of relations which has a very nice logical description.

3.2.1 Representing Stably Compact Spaces

Theorem 3.1.35 tells us that for every stably compact space there is a continuous sequent calculus that represents it. As its tokens are pairs of open and compact sets this is far from being a syntactic representation. Normally, we want to construct a continuous sequent calculus syntactically and then prove that it corresponds to a certain stably compact space.

We have already done most of the work in Section 3.1.3 under the heading “Logic”. The next theorem says in effect that Proposition 3.1.32 identifies the key properties.

Theorem 3.2.1. *Let X be a stably compact space and L a continuous sequent calculus. Then $\text{spec}(L)$ and X are homeomorphic if and only if there are maps $o: L \rightarrow \Omega(X)$ and $k: L \rightarrow \mathcal{K}(X)$ satisfying the following properties:*

1. $(\forall \phi \in L) o(\phi) \subseteq k(\phi)$;
2. $(\forall \phi, \psi \in L) \phi \Vdash \psi \iff k(\phi) \subseteq o(\psi)$; and
3. for $K \subseteq O \subseteq X$, where K is compact and O open, there is a formula ϕ such that $K \subseteq o(\phi) \subseteq k(\phi) \subseteq O$.

In this case o corresponds to $\mathcal{O}_{(\cdot)}$, k to $\mathcal{K}_{(\cdot)}$, and both maps translate finite conjunction and disjunction into intersection and union, respectively.

Proof. Let us start by making some observations: The first two conditions, together with Lemma 1.2.17, imply that o is a monotone map from (L, \Vdash) to $(\Omega(X), \ll)$ and that k is antitone from (L, \Vdash) to $(\mathcal{K}(X), \ll)$. This together with (3) ensures that both maps take \perp to the empty set and \top to X .

If $X = \text{spec}(L)$ we use $\mathcal{O}_{(\cdot)}$ and $\mathcal{K}_{(\cdot)}$ as our maps o and k . The conditions of the theorem are then verbatim the ones of Proposition 3.1.32, and

it is clear that they are invariant under going to a homeomorphic space $X \cong \text{spec}(L)$.

Conversely, assume we are given o and k satisfying the three conditions of the theorem. The outline of the proof is to extend o as indicated in the diagram

$$\begin{array}{ccc} \text{idl}(L) & \xrightarrow{O} & \Omega(X) \\ & \swarrow \text{[]} & \nearrow o \\ & L & \end{array}$$

and to prove that O is an isomorphism. We define the extension by

$$O(I) := \bigcup^{\uparrow} \{o(\phi) \mid \phi \in I\}.$$

To see that the union is directed take $\phi, \psi \in I$. We get $\phi \vee \psi \in I$ and by roundness we find a $\chi \in I$ such that $\phi \vee \psi \Vdash \chi$. This implies $\phi \Vdash \chi$ and $\psi \Vdash \chi$ and as o is monotone with respect to \Vdash we get $o(\phi), o(\psi) \subseteq o(\chi)$. The map O is clearly monotone and we claim that it also reflects the order. Given $O(I) \subseteq O(J)$ and $\phi \in I$ we find an element $\psi \in I$ satisfying $\phi \Vdash \psi$, by roundness of I , and we get

$$o(\phi) \subseteq k(\phi) \subseteq o(\psi) \subseteq O(I) \subseteq O(J) = \bigcup^{\uparrow} \{o(\chi) \mid \chi \in J\}.$$

As $k(\phi)$ is compact, there is a $\chi \in J$ such that $k(\phi) \subseteq o(\chi)$ which implies $\phi \Vdash \chi \in J$ and thus $\phi \in J$. So, O reflects the order and is, in particular, injective.

To prove that it is an isomorphism it remains to show that it is surjective. Given $U \in \Omega(X)$ we define

$$I_U := \{\phi \in L \mid k(\phi) \subseteq U\}$$

and verify that this is a filter using Proposition 3.1.1. Closure under \vee is the only property that is not an immediate consequence of (2) and (3). So, let us suppose $k(\phi), k(\psi) \subseteq U$. By (3) we find a $\chi \in L$ such that

$$k(\phi), k(\psi) \subseteq k(\phi) \cup k(\psi) \subseteq o(\chi) \subseteq k(\chi) \subseteq U.$$

We get $\phi \Vdash \chi$ and $\psi \Vdash \chi$ which entails $\phi \vee \psi \Vdash \chi$ and thus $\phi \vee \psi \in I_U$.

Obviously, $O(I_U) \subseteq U$ holds since for all ϕ we have $o(\phi) \subseteq k(\phi)$. For the converse inclusion, take an $x \in U$. This is equivalent to $\{x\} \subseteq U$ and the third condition ensures the existence of a formula $\phi \in L$ such that $\{x\} \subseteq \uparrow x \subseteq o(\phi) \subseteq k(\phi) \subseteq U$. Thus, we get in particular $\phi \in I_U$ and $x \in o(\phi)$ which entails $U = O(I_U)$. So, O is a monotone, order reflecting and surjective map and hence an isomorphism.

At this point we already know that L corresponds to X :

$$X \cong \text{pt}(\Omega(X)) \cong \text{pt}(\text{idl}(L)) \cong \text{spec}(L)$$

Now, we prove that O is an extension of o . Because of the monotonicity of o we have $O([\Vdash]\phi) = \bigcup^\uparrow \{o(\psi) \mid \psi \Vdash \phi\} \subseteq o(\phi)$. For $x \in o(\phi)$ we again find a ψ such that $\{x\} \subseteq o(\psi) \subseteq k(\psi) \subseteq o(\phi)$ which implies $\psi \Vdash \phi$ and thus $x \in O([\Vdash]\phi)$. This proves that o corresponds to $\mathcal{O}_{(\cdot)}$.

For k we define the map $K: \text{filt}(L) \rightarrow \mathcal{K}(X)$ by

$$K(F) := \bigcap_{\downarrow} \{k(\phi) \mid \phi \in F\}.$$

The proof that it extends k is analogous to the argument for O . We claim that this map is also the one induced by O via the Hofmann-Mislove Theorem. We know from Proposition 3.1.25 that the isomorphism between Scott-open filters on $\text{idl}(L)$ and $\text{filt}(L)$ is given by the map

$$\psi(F) := \{I \in \text{idl}(L) \mid I \cap F \neq \emptyset\}.$$

So, we have to chase the diagram

$$\begin{array}{ccc} \text{filt}(L) & \xrightarrow{K} & \mathcal{K}(X) \\ \downarrow \psi & & \downarrow \text{HMT} \\ \Sigma\text{OFilt}(\text{idl}(L)) & \xleftarrow{O^{-1}[\cdot]} & \Sigma\text{OFilt}(\Omega(X)) \end{array}$$

i.e. we need to check the equality

$$\begin{aligned} F \mapsto K(F) \mapsto \{U \in \Omega(X) \mid K(F) \subseteq U\} &\mapsto \{I \in \text{idl}(L) \mid K(F) \subseteq O(I)\} \\ &\stackrel{!}{=} \{I \in \text{idl}(L) \mid I \cap F \neq \emptyset\}. \end{aligned}$$

Expanding the condition $K(F) \subseteq O(I)$ yields $\bigcap_{\downarrow} \{k(\phi) \mid \phi \in F\} \subseteq O(I)$. By Corollary 1.2.14 such a filtered intersection of compact saturated sets is contained in an open set if and only if one of the members of the family already is. So, the condition is equivalent to

$$(\exists \phi \in F) k(\phi) \subseteq O(I) = \bigcup^\uparrow \{o(\psi) \mid \psi \in I\}$$

and this, in turn, is equivalent to $(\exists \phi \in F, \psi \in I) k(\phi) \subseteq o(\psi)$. Now, $k(\phi) \subseteq o(\psi)$ is tantamount to $\phi \Vdash \psi$ and because of roundness of filters and ideals we see that $K(F) \subseteq O(I)$ is equivalent to $I \cap F \neq \emptyset$. This means that k corresponds to $\mathcal{K}_{(\cdot)}$ and, because of Proposition 3.1.31 and our considerations at the beginning of this proof, we can infer that o and k respect \top , \perp , \wedge and \vee . \square

The first condition of the theorem is usually straightforward to check; it is just a *sanity* requirement about o and k . The two implications of the second condition can be understood as *soundness* and *completeness*. Finally, the third one means that the interpretations of token are *dense* with respect to open and compact saturated sets. Density implies that the open sets of the form $o(\phi)$ are a basis of the topology and that the complements of the $k(\phi)$ are a basis for the co-compact topology. The first statement is the content of Corollary 3.1.33, the latter is an immediate consequence of Theorem 3.1.30.

We can turn the above characterisation around to construct continuous sequent calculi:

Theorem 3.2.2. *Let X be a stably compact space, L a $(2, 2, 0, 0)$ -algebra, and $o: L \rightarrow \Omega(X)$ and $k: L \rightarrow \mathcal{K}(X)$ two maps that satisfy*

1. $(\forall \phi \in L) o(\phi) \subseteq k(\phi)$;
2. *for $K \subseteq O \subseteq X$, where K is compact and O open, there is a formula ϕ such that $K \subseteq o(\phi) \subseteq k(\phi) \subseteq O$;*

and that translate finite conjunction and disjunction into intersection and union. Then

$$\Gamma \Vdash \Delta \quad :\iff \quad k(\bigwedge \Gamma) \subseteq o(\bigvee \Delta)$$

defines a continuous sequent calculus that represents X .

Proof. The rules for consequence relations are all easy to check. Let us do (L \vee) as an example: Suppose we have $\phi, \Gamma \Vdash \Delta$ and $\psi, \Gamma \Vdash \Delta$, i.e. $k(\phi) \cap k(\bigwedge \Gamma), k(\psi) \cap k(\bigwedge \Gamma) \subseteq o(\bigvee \Delta)$. From this we immediately get

$$\begin{aligned} \left(k(\phi) \cap k(\bigwedge \Gamma) \right) \cup \left(k(\psi) \cap k(\bigwedge \Gamma) \right) &= (k(\phi) \cup k(\psi)) \cap (\bigwedge \Gamma) \\ &= k(\phi \vee \psi) \cap (\bigwedge \Gamma) \subseteq o(\bigvee \Delta) \end{aligned}$$

and hence $\phi \vee \psi, \Gamma \Vdash \Delta$.

To see that \Vdash is a continuous sequent calculus it suffices to check that it is closed under (Cut) and has interpolation. The former is trivial, and for (L-Int) consider $\phi, \Gamma \Vdash \Delta$, or in other words $k(\phi) \cap k(\bigwedge \Gamma) \subseteq o(\bigvee \Delta)$. As a saturated set $k(\phi)$ is the filtered intersection of its open neighbourhoods and thus the density condition (2) implies that we can write it as a filtered intersection

$$k(\phi) = \bigcap_{\downarrow} \{k(\psi) \mid k(\phi) \subseteq o(\psi) \subseteq k(\psi)\}.$$

Thus we can also write $k(\phi) \cap k(\bigwedge \Gamma)$ as a filtered intersection of terms $k(\psi) \cap k(\bigwedge \Gamma)$, with ψ as above. By Corollary 1.2.14 this implies that there

is a ψ such that $k(\phi) \subseteq o(\psi)$ and $k(\psi) \cap k(\bigwedge \Gamma) \subseteq o(\bigvee \Delta)$. Translating these subset containments yields $\phi \Vdash \psi$ and $\psi, \Gamma \Vdash \Delta$.

Having shown that \Vdash is a continuous sequent calculus it is clear that o and k satisfy the conditions of Theorem 3.2.1 by definition and consequently that L represents X . \square

Examples

To show how these results can be used we construct continuous sequent calculi that represent two well-known examples of stably compact spaces:

Example 3.2.3 (The Extended Positive Reals). We begin with a very simple space, the extended positive real numbers:

$$\overline{\mathbb{R}_0^+} := \{r \in \mathbb{R} \mid 0 \leq r\} \cup \{\infty\}$$

The topology is the Scott topology for the canonical order on \mathbb{R}_0^+ and the new top element ∞ .

We take L to be the term algebra over the rational numbers \mathbb{Q}_0^+ and set

$$\begin{aligned} o(q) &:=]q, \infty] \\ k(q) &:= [q, \infty] \end{aligned}$$

There is a unique way to extend o and k to L such that they respect the algebraic structure, in particular we have to set $o(\perp) = k(\perp) = \emptyset$ and $o(\top) = k(\top) = \overline{\mathbb{R}_0^+}$.

It is clear that o and k satisfy the conditions of Theorem 3.2.2 as the rational numbers are dense in the real numbers. So, the theorem tells us how we have to define \Vdash_L such that L represents $\overline{\mathbb{R}_0^+}$.

We now want to give syntactic rules that generate the consequence relation \Vdash_L . From Lemma 2.2.1 we know that we only have to cater for sequents made up exclusively from the generators of L . Since $\overline{\mathbb{R}_0^+}$ is linearly ordered, so are $\Omega(\overline{\mathbb{R}_0^+})$ and $\mathcal{K}(\overline{\mathbb{R}_0^+})$. This means that sequents relating generators arise from ones of the form $\phi \Vdash \psi$ by weakening. These are obviously generated by the single rule

$$\frac{q < r}{r \Vdash q}$$

where we can assume that, whatever our syntactic representation of rational numbers, we can decide the inequality $q < r$.

Example 3.2.4 (The Interval Domain). The interval domain can be seen as a model for computations that produce real numbers, in our case taken from the unit interval $[0, 1]$. It is given by

$$\mathbb{I} := \{[x, y] \mid 0 \leq x \leq y \leq 1\},$$

the compact sub-intervals of $[0, 1]$, ordered by reverse inclusion. This yields a continuous Scott domain and hence, in particular, a stably compact space. The order of approximation is given by

$$[x, y] \ll [x', y'] \iff [x', y'] \in \text{int}([x, y]).$$

Here, the interior is taken in $[0, 1]$, and not in \mathbb{R} .

We construct a continuous sequent calculus to represent it by refining the previous example. We take the atomic formulae $\lambda(q)$ and $\rho(q)$, where q is a rational number such that $0 < q < 1$, and call the term algebra generated by them L . The intended reading of the atomic formulae is as follows:

$$\begin{aligned} o(\lambda(q)) &:= \{I \in \mathbb{I} \mid I \subseteq [0, q[\} = \dagger[0, q] \\ o(\rho(q)) &:= \{I \in \mathbb{I} \mid I \subseteq]q, 1] \} = \dagger[q, 1] \\ k(\lambda(q)) &:= \{I \in \mathbb{I} \mid I \subseteq [0, q] \} = \uparrow[0, q] \\ k(\rho(q)) &:= \{I \in \mathbb{I} \mid I \subseteq [q, 1] \} = \uparrow[q, 1] \end{aligned}$$

As before there is a unique way of extending o and k to L that translates logical connectives into the corresponding set theoretic operations.

The only condition of Theorem 3.2.2 that is non-trivial is density: First, note that we get $o(\lambda(r) \wedge \rho(q)) = \dagger[q, r]$ and $k(\lambda(r) \wedge \rho(q)) = \uparrow[q, r]$, i.e. we have tokens for all $\dagger I$ and $\uparrow I$ where I is an interval with rational endpoints. Now, let us suppose we are given $O \in \Omega(\mathbb{I})$ and $K \in \mathcal{K}(\mathbb{I})$ such that $K \subseteq O$. As \mathbb{I} is a continuous domain, the sets of the form $\dagger I$ are a basis of the Scott topology, and since \mathbb{Q} is dense in \mathbb{R} we can restrict ourselves to intervals with rational endpoints. By the compactness of K we get finitely many intervals $I_1, \dots, I_n \in O$ such that $K \subseteq \dagger I_1 \cup \dots \cup \dagger I_n$ and by our previous observations L contains the tokens to guarantee density.

One task remains to be done, namely to come up with syntactic rules for the continuous sequent calculus \Vdash that we can extract via Theorem 3.2.2. As in the previous example we only have to worry about sequents containing atomic formulae. Furthermore, as $[0, 1]$ is linearly ordered, we see that $\lambda(q) \cap \lambda(r)$ must be either $\lambda(q)$ or $\lambda(r)$, and similarly for $\lambda(q) \cup \lambda(r)$ and the corresponding terms involving ρ . This implies that any sequent that contains more than one λ - or more than one ρ -atom on either side of the turnstile is derivable from a simpler sequent by weakening.

We can go even further: The meaning of a sequent

$$\rho(q), \lambda(r) \Vdash \rho(q'), \lambda(r')$$

is

$$\uparrow[q, r] \subseteq \dagger[q', 1] \cup \dagger[0, r']$$

which is equivalent to the disjunction of $[q, r] \subseteq]q', 1]$ and $[q, r] \subseteq [0, r'[$. This shows that we can restrict ourselves to singletons on the right and the left, unless the right hand side is empty.

We sum up the discussion with the three sound and complete rules for our continuous sequent calculus L that represents the interval domain \mathbb{I} :

$$\frac{q < r}{\lambda(q) \Vdash \lambda(r)} \quad \frac{q < r}{\rho(r) \Vdash \rho(q)} \quad \frac{q < r}{\lambda(q), \rho(r) \Vdash}$$

3.2.2 Individual Constructions

Our next topic are domain constructions in logical form. We discuss a number of them in some detail to show how the techniques we have developed so far can be applied.

From now on we use the following convention: If the correspondence between a continuous sequent calculus L and a stably compact space X is understood, in particular if we say that L represents X , then we assume that we have already constructed the functions required by Theorem 3.2.1 and we refer to them as $\mathbf{O}[\cdot]$ and $\mathbf{K}[\cdot]$.

Let us begin with the most elementary domain construction:

Lifting

For any topological space X its lifting $X_\perp := X \cup \{\perp\}$ has one new point \perp which is the new bottom element. We can ensure this topologically by adding just one open set to the topology:

$$\Omega(X_\perp) := \Omega(X) \cup \{X_\perp\}$$

Now suppose that L is a continuous sequent calculus that represents X . We let L_\perp be the algebra freely generated by L , but where we identify the old and the new *falsum* $\perp = \perp'$. Note that we distinguish between the old and the new *verum* $\top \neq \top'$. We generate the logic for L from the one rule

$$\frac{\Gamma \Vdash_L \Delta}{\Gamma \Vdash_{L_\perp} \Delta}$$

and set $o(\phi) := \mathbf{O}[\phi]$ and $k(\phi) := \mathbf{K}[\phi]$ for $\phi \in L$. As in our previous examples we can extend o and k to all of L_\perp and we then have to check that the conditions of Theorem 3.2.2 are satisfied. But this is the case because of $X_\perp = o(\top') = k(\top')$ and the fact that for all compact $K \subseteq X$ we have

$$K \subseteq \mathbf{O}[\top] = \mathbf{K}[\top] = X \subseteq X_\perp.$$

This completes the proof that L_\perp represents X_\perp .

Sums

The sum or coproduct of topological spaces is simply the disjoint union with the topology that is generated by the disjoint union of the respective topologies.

This construction is also very easy because we have already done all the necessary work in Section 2.2.3: In MLS the coproduct $L + M$ can be constructed from L and M using the rules given on page 81.

From Theorem 3.1.40 we know that MLS and the Kleisli category $\text{StCp}_{\mathcal{K}}$ are equivalent. Furthermore, the usual left adjoint from StCp to $\text{StCp}_{\mathcal{K}}$ (see [Mac71, Chapter VI, Theorem 1]) is the identity on objects. This means in particular that coproducts in $\text{StCp}_{\mathcal{K}}$ and StCp are the same since left adjoints preserve them. Thus, we get $\text{spec}(L + M) \cong \text{spec}(L) + \text{spec}(M)$, where the latter coproduct is taken in StCp .

It is also not hard to verify that the functions o and k are the extensions of $\mathbb{O}[\cdot]$ and $\mathbb{K}[\cdot]$ for L and M .

Remark. *Semantic considerations along these lines can be used to see that slight changes to the rules given in Section 2.2.3 produce related sum constructions: Without the rule*

$$\frac{\Vdash_L \Gamma \quad \Vdash_M \Delta}{\Vdash_{L+M} \{0\} \times \Gamma, \{1\} \times \Delta}$$

we get the lifted sum, and if we omit the side condition $\Gamma \neq \emptyset$ for the first two rules given there we get the coalesced sum.

At the end of Section 2.2.3 we have seen that we can endow the coproduct with projections that also make it into a product. Hence, in MLS and all the equivalent categories finite products and finite coproducts are isomorphic. Semantically, the reason for this is easy to see if we consider how disjoint unions and the co-compact topology interact:

Lemma 3.2.5. *If X and Y are stably compact spaces, then $X_{\kappa} + Y_{\kappa} = (X + Y)_{\kappa}$ holds.*

Proof. This follows immediately from the following two observations: If $K \subseteq X$ is compact saturated then $K \cup Y$ is a compact saturated subset of $X + Y$. Starting with a compact saturated $K \subseteq X + Y$, the sets $K \cap X$ and $K \cap Y$ are again compact saturated. \square

If L corresponds to X and M to Y , then we know from Theorem 3.1.30 that $(L^{\text{op}} + M^{\text{op}})^{\text{op}}$ corresponds to $(X_{\kappa} + Y_{\kappa})_{\kappa}$. This space is simply $X + Y$ because of the previous lemma and the fact that taking the co-compact topology is an involution (Corollary 1.3.12).

Remark. *There is another, more direct semantic argument that in StCp^* finite products and coproducts are isomorphic. The injections from X and*

Y to $X + Y$ are given by $x I_0 x' :\Leftrightarrow x \sqsubseteq_X x'$ and $y I_1 y' :\Leftrightarrow y \sqsubseteq_Y y'$. In the light of our discussion at the end of the previous chapter, we can also see them as the relations corresponding to the continuous functions given by the two subset inclusions. Given two closed relations $F: X \multimap Z$ and $G: Y \multimap Z$, the unique mediating morphism D is given by

$$a D z \quad :\Leftrightarrow \quad \begin{cases} a F z & \text{for } a \in X \\ a G z & \text{for } a \in Y \end{cases}$$

where we assume without loss of generality that X and Y are disjoint.

These morphism correspond directly to the MLS-morphisms \vdash_{ι_0} , \vdash_{ι_1} and $\vdash_{f,g}$ constructed in Section 2.2.3, but it is also straightforward to check that they satisfy the universal property.

The projections from the product are very similar to the injections:

$$\begin{aligned} x P_0 x' & \quad :\Leftrightarrow \quad x \sqsubseteq_X x' & \text{and} \\ y P_1 y' & \quad :\Leftrightarrow \quad y \sqsubseteq_Y y' \end{aligned}$$

Note that these relations almost correspond to functions; they are however partial. It is readily checked that P_0 and P_1 are closed subsets of $(X + Y) \times X_\kappa$ and $(X + Y) \times Y_\kappa$, respectively, and hence morphisms in StCp^* . For two closed relations $F: Z \multimap X$ and $G: Z \multimap Y$ we define:

$$z \langle F, G \rangle a \quad :\Leftrightarrow \quad \begin{cases} z F a & \text{for } a \in X \\ z G a & \text{for } a \in Y \end{cases}$$

This is clearly a morphism and uniquely determined by F and G , thus showing that $X + Y$ is a product of X and Y .

Products

We know from Proposition 1.3.4 that the product of two stably compact spaces X and Y is again stably compact. Let the continuous sequent calculi L and M represent X and Y , respectively. Similarly to sums we freely generate $L \otimes M$ from $\{0\} \times L \cup \{1\} \times M$, the disjoint union of L and M . The intended reading of the tokens is:

$$\begin{aligned} o(0, \phi) & := \mathbf{O}[\phi] \times Y & o(1, \psi) & := X \times \mathbf{O}[\psi] \\ k(0, \phi) & := \mathbf{K}[\phi] \times Y & k(1, \psi) & := X \times \mathbf{K}[\psi] \end{aligned}$$

The extensions of o and k to $L \otimes M$ clearly satisfy the first condition of Theorem 3.2.2. To show density suppose we have $K \subseteq O \subseteq X \times Y$, where K is compact saturated and O open. For every point $\langle x, y \rangle \in K$ there are open sets $U \in \mathcal{N}(x)$ and $V \in \mathcal{N}(y)$ such that $U \times V \subseteq O$. Hence, there are

tokens $\phi_{xy} \in L$ and $\psi_{xy} \in M$ satisfying $x \in \mathbf{O}[\phi_{xy}] \subseteq \mathbf{K}[\phi_{xy}] \subseteq U$ and $y \in \mathbf{O}[\psi_{xy}] \subseteq \mathbf{K}[\psi_{xy}] \subseteq V$. As K is compact, finitely many of the sets

$$\mathbf{O}[\phi_{xy}] \times \mathbf{O}[\psi_{xy}] = o(\langle 0, \phi_{xy} \rangle \wedge \langle 1, \psi_{xy} \rangle)$$

cover it. This means that we can find a token

$$\chi := (\langle 0, \phi_{(xy)_1} \rangle \wedge \langle 1, \psi_{(xy)_1} \rangle) \vee \cdots \vee (\langle 0, \phi_{(xy)_n} \rangle \wedge \langle 1, \psi_{(xy)_n} \rangle)$$

such that $K \subseteq o(\chi) \subseteq k(\chi) \subseteq O$.

As we have seen several times before, we can now extract $\Vdash_{L \otimes M}$ and only have to come up with rules that generate enough sequents relating atomic formulae. To this end we consider a generic situation where a compact saturated set

$$K = (K_1 \times Y) \cap \cdots \cap (K_l \times Y) \cap (X \times C_1) \cap \cdots \cap (X \times C_m)$$

is contained in an open set of the form

$$(U_1 \times Y) \cup \cdots \cup (U_n \times Y) \cup (X \times V_1) \cup \cdots \cup (X \times V_p).$$

The set K is of the form $K' \times C'$ and hence has to be a subset of either $(U_1 \times Y) \cup \cdots \cup (U_n \times Y)$ or $(X \times V_1) \cup \cdots \cup (X \times V_p)$; let us call it O . Depending on which is the case we get that either $(K_1 \times Y) \cap \cdots \cap (K_l \times Y)$ is already a subset of O , or $(X \times C_1) \cap \cdots \cap (X \times C_m)$ is. This means that because of weakening we get all the sequents that we are interested in from ones relating atomic formulae of the same kind. Hence, we can generate our logic by the two rules:

$$\frac{\Gamma \Vdash_L \Delta}{\{0\} \times \Gamma \Vdash_{L \otimes M} \{0\} \times \Delta} \quad \frac{\Gamma \Vdash_M \Delta}{\{1\} \times \Gamma \Vdash_{L \otimes M} \{1\} \times \Delta}$$

Note that this also takes care of extreme cases like $\Gamma \Vdash$ which might be a problem because of $\emptyset \times Y = \emptyset$.

The Smyth Power Domain

For a space X this power domain is simply $\mathcal{K}(X)$, i.e. the set of compact saturated subsets ordered by reverse inclusion, though usually without the compact top element \emptyset .

It has a very simple description which follows almost immediately from some of the results in Section 3.1.4: Let us suppose that L represents a stably compact space X . Then Lemma 3.1.37 tells us that the sets $\hat{\mathcal{O}}_\phi = \square \mathcal{O}_\phi$, $\phi \in L$, form a basis for the Scott topology on $\mathcal{K}(X)$. From Lemma 1.1.5 we know that every compact saturated set in $\mathcal{K}(X)$ can be written as an intersection of finite unions of sets $\uparrow K$.

This suggests that we can essentially reuse the tokens from L . We freely generate $\mathcal{P}_S(L)$ from atomic formulae $\Box\phi$, where $\phi \in L$, and set:

$$\begin{aligned} o(\Box\phi) &:= \Box\mathbf{O}[\phi] = \{K \in \mathcal{K}(X) \mid K \subseteq \mathbf{O}[\phi]\} \\ k(\Box\phi) &:= \Box\mathbf{K}[\phi] = \uparrow\mathbf{K}[\phi] = \{K \in \mathcal{K}(X) \mid K \subseteq \mathbf{K}[\phi]\} \end{aligned}$$

We clearly have $o(\Box\phi) \subseteq k(\Box\phi)$, and hence the first condition of Theorem 3.2.2 is satisfied for the extensions of o and k . For density we have to combine the observations from the previous paragraph. To do this we have to refine the proof of Lemma 3.1.37 slightly: Let $\mathcal{C} \subseteq \mathcal{K}(X)$ be compact saturated and $\mathcal{O} \supseteq \mathcal{C}$ Scott-open in $\mathcal{K}(X)$. For every $K \in \mathcal{C}$ we find $L, L' \in \mathcal{O}$ such that $L \ll L' \ll K$. By Lemma 1.2.17 this means that there is an open set O such that $K \subseteq O \subseteq L'$ and thus we can find a token $\phi \in L$ satisfying $K \subseteq \mathbf{O}[\phi] \subseteq \mathbf{K}[\phi] \subseteq L'$. This implies $K \in o(\Box\phi)$ and $k(\Box\phi) \subseteq \uparrow L' \subseteq \uparrow L \subseteq \mathcal{O}$. As \mathcal{C} is compact, finitely many of the sets $o(\Box\phi)$ suffice to cover it and hence we get

$$\mathcal{C} \subseteq o(\Box\phi_1 \vee \dots \vee \Box\phi_n) \subseteq k(\Box\phi_1 \vee \dots \vee \Box\phi_n) \subseteq \mathcal{O}.$$

As before we can extract $\Vdash_{\mathcal{P}_S(L)}$ via o and k , and we only have to characterise the sequents relating atomic formulae and come up with rules that generate enough such sequents. To begin with we observe that

$$k(\Box\phi_1) \cap \dots \cap k(\Box\phi_n) = \Box\mathbf{K}[\phi_1] \cap \dots \cap \Box\mathbf{K}[\phi_n] = \Box(\mathbf{K}[\phi_1] \cap \dots \cap \mathbf{K}[\phi_n]).$$

This is a principal filter in $\mathcal{K}(X)$ and as such contained in a union of Scott-open sets $o(\Box\psi) = \Box\mathbf{O}[\psi]$ if and only if it is contained in one of them already. We infer that up to weakening we can restrict ourselves to singletons on the right, disregarding for the moment the case of an empty sequent on the right. Moreover, a containment $\Box K = \uparrow K \subseteq \Box O$, for $K \subseteq X$ compact saturated and $O \subseteq X$ open, holds if and only if $K \subseteq O$. This means that we need only one rule:

$$\frac{\phi_1, \dots, \phi_n \Vdash_L \psi}{\Box\phi_1, \dots, \Box\phi_n \Vdash_{\mathcal{P}_S(L)} \Box\psi}$$

If we consider the lattice $\mathcal{K}(X)$ we never have $\Box\phi_1, \dots, \Box\phi_n \Vdash_{\mathcal{P}_S(L)}$ since the empty set is contained in every $\Box K$. If we want to describe the usual Smyth power domain $\mathcal{P}_S(X) = \mathcal{K}(X) \setminus \{\emptyset\}$ we need one extra rule that expresses $\Box\emptyset = \emptyset$:

$$\frac{\phi_1, \dots, \phi_n \Vdash_L}{\Box\phi_1, \dots, \Box\phi_n \Vdash_{\mathcal{P}_S(L)}}$$

3.2.3 Function Spaces

We begin with a brief exposition of the central results about topological function spaces, very much in the spirit of Chapter 1. It also contains the prerequisites for the next section on relation spaces. Most of the material follows the exposition in [Law88].

More Basics

Given two topological spaces X and Y we can endow Y^X , the set of all continuous functions from X to Y with the *compact-open topology*. It has the sets

$$N(K, O) := \{f: X \rightarrow Y \mid f[K] \subseteq O\},$$

for $K \subseteq X$ compact and $O \subseteq Y$ open, as a subbasis. In general this does not make Y^X into an exponential in Top. This is, however, the case if X is locally compact:

Proposition 3.2.6. *For a locally compact space X the function spaces Y^X , endowed with the compact-open topology, is an exponential of X and Y .*

Proof. We have to check a number of details that appear in the following diagram that defines exponentials:

$$\begin{array}{ccc}
 Y^X & & Y^X \times X \xrightarrow{\epsilon} Y \\
 \uparrow \exists! \hat{f} & & \uparrow \hat{f} \times X \\
 Z & & Z \times X
 \end{array}$$

The evaluation ϵ is given by $\epsilon(f, x) := f(x)$. Let us prove that it is continuous: For $f(x) \in O$ open, we find a compact neighbourhood K of x such that $f[K] \subseteq O$ since f is continuous and X locally compact. This immediately implies $\epsilon[N(K, O) \times K] \subseteq O$ and $\langle f, x \rangle \in N(K, O) \times K$.

It is clear that we must define \hat{f} by $\hat{f}(z) := \lambda x. f(z, x)$. So, we only have to show that for all z the function $\hat{f}(z)$ is continuous and, moreover, that \hat{f} itself is continuous. The former is easy as we can write $\hat{f}(z)$ as a composition of continuous functions

$$X \cong X \times \{z\} \hookrightarrow X \times Z \xrightarrow{f} Y.$$

For the latter suppose $\hat{f}(z) \in N(K, O)$. This implies that for all $x \in K$ we have $f(z, x) \in O$ and by continuity of f we find neighbourhoods $U_x \in \mathcal{N}(z)$ and $V_x \in \mathcal{N}(x)$ such that $f[U_x \times V_x] \subseteq O$. As the set K is compact, finitely many V_{x_1}, \dots, V_{x_n} cover it and for the corresponding neighbourhoods U_{x_i} in Z we get $\hat{f}[U_{x_1} \cap \dots \cap U_{x_n}] \subseteq N(K, O)$. \square

Remark. One can generalise this result to core compact spaces, i.e. spaces whose sobrification is locally compact, using the Isbell topology. The proposition also has a converse: The compact-open topology (or the Isbell topology for the non-sober case) is the only candidate for a topology on the exponential, i.e. if it does not yield one then the two spaces do not have an exponential in Top .

Unfortunately, the compact-open topology is in general not locally compact. This is one reason why it is hard to find cartesian closed categories of topological spaces. Continuous domains are always locally compact and so the question arises how the exponentials taken in Top and those taken in DCPO are related. Furthermore, to construct a CCC of domains one has to come up with a class of domains such that exponentiation does not destroy continuity. We mentioned several such classes in Section 1.1.2.

Our focus are sober spaces and we know from Proposition 1.2.9 that they form dcpo 's with respect to their specialisation orders. The following lemma shows that this order is well-behaved with respect to taking exponentials.

Lemma 3.2.7. *The specialisation order on Y^X is the extensional order of functions, i.e. $f \sqsubseteq g \iff (\forall x \in X) f(x) \sqsubseteq g(x)$.*

Proof. Let us call the extensional order \sqsubseteq_e and the specialisation order with respect to the compact-open topology \sqsubseteq_s . If we have $f \sqsubseteq_e g$ and $f[K] \subseteq O$, where K is compact and O is open, we get $g[K] \subseteq \uparrow f[K] \subseteq O$ and thus $f \sqsubseteq_s g$.

Conversely, suppose $f \sqsubseteq_s g$ and $f(x)$ lies in an open set O . We get $f[\uparrow x] \subseteq O$ and hence $f \in N(\uparrow x, O)$ which implies $g \in N(\uparrow x, O)$ and thus $g(x) \in O$. This proves $f \sqsubseteq_e g$. \square

If Y is a continuous Scott domain with the Scott topology then Y^X gets most of its structure from that space. In particular, we will see that it is again a continuous Scott domain and that the Scott topology on Y^X agrees with the compact-open topology.

The most important tool to understanding such an exponential as a domain are *step functions*. Given an open set $O \subseteq X$ and a point $y \in Y$ we define:

$$(O \searrow y)(x) := \begin{cases} y & \text{if } x \in O \\ \perp & \text{otherwise} \end{cases}$$

Lemma 3.2.8. *Let f be a continuous function from a space X to a pointed continuous domain Y equipped with the Scott topology, $O \in \Omega(X)$ and $y \in Y$. If $O \ll f^{-1}[\uparrow y]$ then $(O \searrow y) \ll f$.*

Proof. Suppose we have $f \sqsubseteq \bigsqcup^\uparrow g_i$. From Lemma 1.2.10 and the comment following it we know that the preimage of a Scott-open set under a directed

supremum of continuous functions is just the directed union of the preimages of the individual functions. This means that we get

$$O \ll f^{-1}[\uparrow y] \subseteq (\bigsqcup^{\uparrow} g_i)^{-1}[\uparrow y] = \bigcup^{\uparrow} (g_i^{-1}[\uparrow y])$$

and thus $O \subseteq g_i^{-1}[\uparrow y]$ for an index i . This implies $(O \searrow y) \sqsubseteq g_i$ and thus $(O \searrow y) \ll f$. \square

Proposition 3.2.9. *For a locally compact space X and a pointed continuous domain Y with the Scott topology the following hold:*

1. *If Y is a Scott domain then so is Y^X .*
2. *If the exponential Y^X is a continuous stably compact domain, then the Scott topology and the compact open topology agree.*

Proof. We begin by showing that every continuous function $f: X \rightarrow Y$ is the supremum of step functions $(U \searrow z)$, where $U \ll f^{-1}[\uparrow z]$. For any $x \in X$ and $y \ll f(x)$ we have $x \in f^{-1}[\uparrow y]$, and as X is locally compact there is an open set O such that $x \in O \ll f^{-1}[\uparrow y]$. This implies that

$$\bigsqcup \{(U \searrow z)(x) \mid U \ll f^{-1}[\uparrow z]\}$$

is larger than y , and as Y is continuous that the supremum is in fact $f(x)$. Hence, f is the supremum of such step functions and by the previous lemma we know that these step functions approximate f . In general, however, this supremum is far from directed.

Next, we check that Y^X has bounded suprema if Y does. The constant bottom function is the least element of the exponential, and for f and g bounded we construct the point-wise supremum $f \sqcup g$ as follows: We add a compact top element to Y to get the continuous lattice Y_{\top} and call the embedding $e: Y \hookrightarrow Y_{\top}$. Taking binary suprema on Y_{\top} is clearly Scott-continuous as directed suprema commute with binary suprema. Now, we consider

$$X \xrightarrow{\langle f, g \rangle} Y \times Y \xrightarrow{e \times e} Y_{\top} \times Y_{\top} \xrightarrow{\sqcup} Y_{\top}$$

and see that if f and g are bounded then this map co-restricts to Y . This yields precisely $f \sqcup g$ which, as a composition of continuous maps, is continuous and hence the supremum in Y^X .

In any dcpo we have that $x, y \ll z$ implies $x \sqcup y \ll z$, provided the supremum exists. This implies that for a Scott domain Y any function $f: X \rightarrow Y$ is the directed supremum of approximating functions, namely that of finite suprema of step functions as constructed above. This shows that Y^X is a continuous Scott domain.

It remains to prove that under the conditions of the proposition the two topologies on the function space coincide. From Lemma 3.2.7 we infer that the subbasic open neighbourhoods $N(K, O)$ are upper sets. To show that they are Scott-open suppose $K \subseteq (\bigsqcup^\uparrow f_i)^{-1}[O] = \bigcup^\uparrow f_i^{-1}[O]$. Then compactness implies $K \subseteq f_i^{-1}[O]$ for an index i and hence $f_i \in N(K, O)$.

For the other direction note that the sets $\{f \in Y^X \mid O \ll f^{-1}[\dagger y]\}$ are Scott-open because of the interpolation property in $\Omega(X)$. We claim that moreover these sets form a subbasis of the Scott topology: Suppose f is in the Scott-open set $U \subseteq Y^X$. We know that f is the supremum of step functions $(O \searrow y)$, where $O \ll f^{-1}[\dagger y]$, or equivalently

$$\uparrow f = \bigcap \{\uparrow(O \searrow y) \mid O \ll f^{-1}[\dagger y]\} \subseteq U.$$

As Y^X is stably compact by assumption, we get from Corollary 1.2.14 that the intersection of finitely many such sets $\uparrow(O \searrow y)$ is already contained in U . The finite intersection of the corresponding subbasic open neighbourhoods $\{g \mid O \ll g^{-1}[\dagger y]\} \subseteq \dagger(O \searrow y) \subseteq \uparrow(O \searrow y)$ is a neighbourhood of f that is a subset of U . Hence, it suffices to consider such Scott-open sets. If f satisfies $O \ll f^{-1}[\dagger y]$, then local compactness allows us to find a compact saturated set K such that $O \subseteq K \subseteq f^{-1}[\dagger y]$ which implies $f \in N(K, \dagger y) \subseteq \dagger(O \searrow y)$: The function f clearly belongs to $N(K, \dagger y)$ and for any $g \in N(K, \dagger y)$ we get $O \subseteq K \subseteq g^{-1}[\dagger y]$ and thus $(O \searrow y) \ll g$ by the previous lemma. \square

A continuous Scott domain is stably compact. Thus, for a locally compact space X and a Scott domain Y the exponential Y^X is a Scott domain and carries the Scott topology. The reason to formulate the proposition in a slightly more general way is that this also covers the case of function spaces between pointed FS domains.

Remark. *We could relax the condition that Y has to be pointed by using the techniques described in [AJ94, Section 4.3.2] and [Jun89, Chapter 3]. As this extra generality is not necessary for the further development of the theory we avoid the technical complications that would be introduced by it.*

The Function Space Construction

Suppose L and M are continuous sequent calculi representing the stably compact spaces X and Y . The space X is locally compact by definition and hence the exponential Y^X in Top exists because of Proposition 3.2.6. Looking at the definition of the compact open topology gives us a hint for a candidate for the continuous sequent calculus corresponding to the function space.

We take terms $(\phi \rightarrow \psi)$, for $\phi \in L$ and $\psi \in M$, as basic tokens and let $[L \rightarrow M]$ be the algebra freely generated by it. The intended meaning of such a token is that it stands for the set of functions mapping the set represented by ϕ into that represented by ψ . However, each token χ denotes an open set $\mathbf{O}[\chi]$ and a compact saturated set $\mathbf{K}[\chi]$; on the other hand we have to provide an open and a compact saturated reading of the new token $(\phi \rightarrow \psi)$. The following interpretations of the new tokens suggest themselves:

$$\begin{aligned} o(\phi \rightarrow \psi) &:= \{f: X \rightarrow Y \mid f[\mathbf{K}[\phi]] \subseteq \mathbf{O}[\psi]\} = N(\mathbf{K}[\phi], \mathbf{O}[\psi]) \\ k(\phi \rightarrow \psi) &:= \{f: X \rightarrow Y \mid f[\mathbf{O}[\phi]] \subseteq \mathbf{K}[\psi]\} \end{aligned}$$

It is clear from the definition that the interpretation $o(\phi \rightarrow \psi)$ is open. Also note that an expression $N(A, B) := \{f: X \rightarrow Y \mid f[A] \subseteq B\}$ —we use this by slight abuse of notation even if A is not compact and B not open—is antitone in A and monotone in B with respect to subset inclusion. From this observation and the fact that $\mathbf{O}[\chi] \subseteq \mathbf{K}[\chi]$ holds for every formula χ in any continuous sequent calculus we get $o(\phi \rightarrow \psi) \subseteq k(\phi \rightarrow \psi)$, the sanity condition of Theorem 3.2.2.

The following proposition shows that from the open point of view we have defined the right continuous sequent calculus together with the right open semantics.

Proposition 3.2.10. *The sets $o(\phi \rightarrow \psi)$ with $\phi \in L$ and $\psi \in M$ form a subbasis of the compact open topology on Y^X .*

Proof. Suppose $f \in N(K, O)$ for a compact saturated set $K \subseteq X$ and an open set $O \subseteq Y$. Then we infer $K \subseteq f^{-1}[O]$ and by Theorem 3.2.1.(3) we can find a token $\phi \in L$ such that $K \subseteq \mathbf{O}[\phi] \subseteq \mathbf{K}[\phi] \subseteq f^{-1}[O]$. This implies $\uparrow f[\mathbf{K}[\phi]] \subseteq O$ and so we can apply the same theorem again to get a token $\psi \in M$ such that $\uparrow f[\mathbf{K}[\phi]] \subseteq \mathbf{O}[\psi] \subseteq \mathbf{K}[\psi] \subseteq O$. Putting it all together we get $f \in o(\phi \rightarrow \psi) \subseteq N(K, O)$, where the subset inclusion follows from the fact that $N(\cdot, \cdot)$ is antitone in the first and monotone in the second argument. \square

For a more restricted case we will prove a stronger property in Theorem 3.2.12, namely density in the sense of Theorem 3.2.2.(2). But prior to this we consider the compact reading of a token $(x \rightarrow y)$. Unfortunately, here the situation is not as nice as in the open case. The sets $k(\phi \rightarrow \psi)$ are clearly saturated, but in general they fail to be compact as the following example illustrates.

Take $I := [0, 1]$ to be the unit interval with the usual topology. It is a compact Hausdorff space and hence in particular stably compact. As I is both open and compact we have $N(I, I) = I^I$ and this corresponds to an

open and a compact reading. The functions

$$f_i(x) = \begin{cases} 0, & \text{if } x \geq 1/i \\ 1 - ix, & \text{otherwise} \end{cases}$$

converge point-wise to the constant zero function with the exception of the argument 0 where its value is 1. As the compact open topology is finer than the topology of point-wise convergence there cannot be a subnet that converges to a continuous function in I^I which thus cannot be compact.

In a way the counterexample is not very surprising: The fact that the exponentials of locally compact spaces are in general not locally compact is the reason why it is so hard to find cartesian closed subcategories of Top. It makes it clear, however, that here, for the first time, we need extra assumptions about the continuous sequent calculi or the spaces they represent.

In spite of this observation, we can prove the compactness of $k(\phi \rightarrow \psi)$ under additional assumptions. The following proposition is given in a general form to be applicable to two situations, the function space between domains and the relations space construction of the next section.

Theorem 3.2.11. *Let X be locally compact and Y a pointed continuous domain such that the exponential Y^X is stably compact. For an open $O \subseteq X$ and a compact saturated $K \subseteq Y$ the set*

$$N(O, K) = \{f: X \rightarrow Y \mid f[O] \subseteq K\}$$

is compact if either X is a continuous domain or K is of the form $K = \uparrow y$.

Proof. The idea of the proof is the following: We show that the set $N(O, K)$ is the intersection of finitely generated compact saturated sets, i.e. sets $\uparrow \mathcal{M}$ where $\mathcal{M} \subseteq_{\text{fin}} Y^X$. Because we assume that Y^X is stably compact this implies that $N(O, K)$ is itself compact.

To this end suppose $f \notin N(O, K)$ which implies that there is an $x \in O$ such that $f(x) \notin K$. If K is a principal filter $\uparrow y$, then we consider the step function $(O \searrow y)$ and instantly see $f \notin \uparrow(O \searrow y) \supseteq N(O, \uparrow y)$.

In the other case where K is an arbitrary compact saturated set but X is a domain we proceed as follows. By Lemma 1.1.5 the compact set K has a neighbourhood $\uparrow\{m_1, \dots, m_n\}$ that does not contain $f(x)$. Now, pick an $x' \in O$ satisfying $x' \ll x$ and consider the finite set of step functions $\mathcal{M} := \{(\uparrow x' \searrow m_1), \dots, (\uparrow x' \searrow m_n)\}$. We claim $f \notin \uparrow \mathcal{M} \supseteq N(O, K)$: The function f is clearly not in $\uparrow \mathcal{M}$ as it maps x to a point that is not above any of the m_i . For the second part of the claim suppose $g \in N(O, K)$. Then $g(x')$ must be larger than at least one m_i which implies $g \supseteq (\uparrow x' \searrow m_i)$.

In either case this implies that the intersection of the sets $\uparrow \mathcal{M} \supseteq N(O, K)$ with \mathcal{M} finite is exactly $N(O, K)$. \square

Note that the theorem applies in particular if X and Y are Scott domains or pointed FS domains. Moreover, as an immediate corollary we get that if the continuous sequent calculi L and M represent such spaces then the compact interpretation $k(\phi \rightarrow \psi)$ of any basic token $(\phi \rightarrow \psi)$ is indeed compact saturated.

Remark. *It is not clear at the moment under which exact conditions the saturated set $N(O, K)$ is compact. For the approach of the above proof to work we have to assume that Y^X is supersober. The domain structure of Y is only needed to guarantee that for every point y that does not lie in a compact set K there is a finitely generated compact set containing K but not y . This appears to be related to Y being a quasicontinuous domain (see [GLS83]). For X we seem to need the property that for an open neighbourhood U of x there is an $x' \in U$ such that $x \in \text{int}(\uparrow x')$ which, for example, is clearly not the case if X is Hausdorff.*

We now return to density for $o(\cdot)$ and $k(\cdot)$, i.e. property (2) of Theorem 3.2.2.

Theorem 3.2.12. *Let L and M be continuous sequent calculi representing stably compact spaces X and Y such that Y has a least element and Y^X is a stably compact continuous domain. For an open $O \subseteq Y^X$ and $K \in \mathcal{K}(Y^X)$ with $K \subseteq O$ there is a formula $\chi \in [L \rightarrow M]$ such that $K \subseteq o(\chi) \subseteq k(\chi) \subseteq O$.*

Proof. We begin by simplifying the problem. By Lemma 1.1.5 we can find finitely many functions f_1, \dots, f_n such that $K \subseteq \uparrow\{f_1, \dots, f_n\} \subseteq O$. Suppose we produce formulae χ_1, \dots, χ_n such that $\uparrow f_i \subseteq o(\chi_i) \subseteq k(\chi_i) \subseteq O$, for $i = 1, \dots, n$. Then we get

$$K \subseteq \uparrow f_1 \cup \dots \cup \uparrow f_n \subseteq o(\chi_1 \vee \dots \vee \chi_n) \subseteq k(\chi_1 \vee \dots \vee \chi_n) \subseteq O$$

which proves the theorem. Hence, it suffices to show that for any $f \in Y^X$ and any open set $O \subseteq Y^X$ containing f there is a formula χ such that $\uparrow f \subseteq o(\chi) \subseteq k(\chi) \subseteq O$.

Let us therefore assume we are given such a function $f \in O \in \Omega(Y^X)$. Then for every $x \in X$ and every $y \ll f(x)$ we can find a $\psi \in M$ satisfying

$$f(x) \in \mathbf{O}[\psi] \subseteq \mathbf{K}[\psi] \subseteq \uparrow y.$$

We get $x \in f^{-1}[\mathbf{O}[\psi]] \in \Omega(X)$ and hence another formula $\phi \in L$ such that

$$x \in \mathbf{O}[\phi] \subseteq \mathbf{K}[\phi] \subseteq f^{-1}[\mathbf{O}[\psi]].$$

We now claim:

$$f \in o(\phi \rightarrow \psi) \subseteq k(\phi \rightarrow \psi) \subseteq \uparrow(\mathbf{O}[\phi] \searrow y) \quad (3.1)$$

The only thing that is not obvious is the last containment. To check it note that $g \in k(\phi \rightarrow \psi)$ is by definition equivalent to $g[\mathbf{O}[\phi]] \subseteq \mathbf{K}[\psi] \subseteq \uparrow y$. This implies $y \ll g(x)$ for all $x \in \mathbf{O}[\phi]$ which proves $(\mathbf{O}[\phi] \searrow y) \sqsubseteq g$.

Repeating the argument from the proof of Proposition 3.2.9 we see that f is the supremum of such step functions $(\mathbf{O}[\phi_i] \searrow y_i)$, or equivalently $\bigcap_i \uparrow(\mathbf{O}[\phi_i] \searrow y_i) = \uparrow f \subseteq O$. Since we assumed that Y^X is stably compact we can invoke Corollary 1.2.14 and infer that there are finitely many such functions that already satisfy

$$f \in \uparrow(\mathbf{O}[\phi_1] \searrow y_1) \cap \cdots \cap \uparrow(\mathbf{O}[\phi_n] \searrow y_n) \subseteq O.$$

Together with the equation (3.1) this implies

$$f \in o((\phi_1 \rightarrow \psi_1) \wedge \cdots \wedge (\phi_n \rightarrow \psi_n)) \subseteq k((\phi_1 \rightarrow \psi_1) \wedge \cdots \wedge (\phi_n \rightarrow \psi_n)) \subseteq O$$

which concludes the proof. \square

For continuous sequent calculi L and M satisfying the premises of the theorem this is a strengthening of Proposition 3.2.10. It implies that the sets $o(\phi \rightarrow \psi)$ form a basis of the topology on Y^X , that the complements of the sets $k(\phi \rightarrow \psi)$ form a basis for the co-compact topology, and moreover that they do so in a joint way.

Putting the last two theorems together we see that under reasonably mild conditions, for example if L and M represent pointed FS domains, the algebra $[L \rightarrow M]$ satisfies the properties of Theorem 3.2.2. Hence, it can be equipped with a continuous consequence relation $\Vdash_{[L \rightarrow M]}$ such that it represents the exponential Y^X of the corresponding spaces.

Unfortunately, this does not mean that we have a function space construction in logical form. For this we lack a syntactic description of sequences $\Gamma \Vdash_{[L \rightarrow M]} \Delta$ relating atomic formulae. Even under more restrictive conditions, like the case of continuous Scott domains, it is not clear how to proceed.

It is instructive to compare this to the situation in Abramsky's original work [Abr91b]. It is at this point where the primality predicate C is needed. We always have:

$$\begin{aligned} (\phi \rightarrow (\psi \wedge \chi)) &\approx (\phi \rightarrow \psi) \wedge (\phi \rightarrow \chi) \quad \text{and} \\ ((\phi \vee \psi) \rightarrow \chi) &\approx (\phi \rightarrow \chi) \wedge (\psi \rightarrow \chi) \end{aligned}$$

For disjunction on the right and conjunction on the left we only get

$$\begin{aligned} (\phi \rightarrow (\psi \vee \chi)) &\gtrsim (\phi \rightarrow \psi) \vee (\phi \rightarrow \chi) \quad \text{and} \\ ((\phi \wedge \psi) \rightarrow \chi) &\gtrsim (\phi \rightarrow \chi) \vee (\psi \rightarrow \chi) \end{aligned}$$

which is a consequence of $(\cdot \rightarrow \cdot)$ being antitone in the first and monotone in the second argument. The first equation becomes an equivalence if $C(\phi)$ holds.

As we observed in the discussion leading up to Proposition 1.2.24 the compact open subsets in a continuous domain are precisely those of the form $\uparrow M$, where M is a finite set of compact elements. The \vee -irreducible elements in $\mathcal{K}\Omega(X)$ for a stably compact algebraic domain X are thus simply the principal filters $\uparrow x$. In effect, the C -predicate says that a formula corresponds to a point in the domain. Given that one of the slogans underlying locales and domains in logical form is to consider open sets, or properties, as primary objects and points as derived ones, this is a bit strange. For algebraic domains this is however justified by the fact that compactness of an element is an intrinsic property, unlike being a member of an arbitrarily chosen basis for a domain. Thus the compact elements give rise to a canonical subbasis of the Scott topology. Moreover, in applications they typically correspond to finite objects and thus, even from a constructivist point of view, the nature of their existence is not in doubt.

In the continuous case points are more problematic. It is also not clear to what extent formulae whose compact interpretation is a point and whose open reading is a Scott-open filter can be used as a substitute for formulae satisfying Abramsky's C -predicate.

Surprisingly, the situation for the relation space is much nicer. This is the topic of the next section.

3.2.4 Relation Spaces

Let us take two continuous sequent calculi L and M corresponding to the topological spaces X and Y and consider the hom-set $\text{MLS}(L, M)$, or equivalently $\text{StCp}^*(X, Y) \cong \text{StCp}_{\mathcal{K}}(X, Y) = \text{StCp}(X, \mathcal{K}(Y))$. From the last description we can see that we can endow it with the compact open topology and make it into a topological space again; in fact, this space $[X \Rightarrow Y]$ turns out to be stably compact. Unfortunately, this does not give rise to a cartesian closed structure nor to a symmetric monoidal closed one. Nonetheless, it is not just an arbitrary object and we will explore the universal property it satisfies. This *relation space* also has a nice description in syntactic terms.

Closure

We begin by showing that StCp^* is not cartesian closed. It has a zero object $\mathbf{0}$, namely the singleton $\{*\}$. Moreover, it is non-trivial, in the sense that it has morphisms that are different from the identities on the objects. The claim now follows immediately from the following well-known general observation:

Lemma 3.2.13. *In a cartesian closed category with a zero object all objects are isomorphic.*

Proof. The zero object $\mathbf{0}$ is both initial and terminal, and hence we have $X \times \mathbf{0} \cong X$ for all objects X . If the product functor $X \times -$ has a right adjoint, it preserves colimits and in particular $\mathbf{0}$. This implies that for all objects X we get $X \cong X \times \mathbf{0} \cong \mathbf{0}$. \square

A candidate for a symmetric monoidal structure suggests itself: the product taken in the category of continuous functions rather than in the category of relations. We have studied the construction of $L \otimes M$ in Section 3.2.2 and have seen that this continuous sequent calculus corresponds to $X \times Y$, where the latter is the cartesian product with the product topology. To avoid confusion with the categorical product, which is the same as the coproduct, we also refer to this space as $X \otimes Y$.

Similar to the cartesian product in Rel this tensor gives rise to a symmetric monoidal closed structure on StCp^* , albeit a rather trivial one. We have $\text{StCp}^*(X \otimes Y, Z) \cong \text{StCp}^*(X, Y_\kappa \otimes Z)$ as is readily checked.

In the following we are more interested in the relation space construction $[X \Rightarrow Y]$, although it cannot have a left adjoint: As we know, in StCp^* finite products and coproducts agree. We can easily construct finite counterexamples that show that $[X \Rightarrow -]$ does not preserve products, for example:

$$\left| [1 \Rightarrow 1 + 2] \right| = \left| \text{StCp}^*(1, 1 + 2) \right| = \left| \text{StCp}(1, \mathcal{K}(1 + 2)) \right| = 6$$

But on the other hand:

$$\left| [1 \Rightarrow 1] + [1 \Rightarrow 2] \right| = \left| [1 \Rightarrow 1] \right| + \left| [1 \Rightarrow 2] \right| = 2 + 3 = 5$$

Nonetheless, $X \otimes -$ and $[X \Rightarrow -]$ are almost adjoint, as we will see later. In fact, $[X \Rightarrow Y]$ is a weak exponential, in the sense that there is an evaluation relation, and for all other morphisms $R: Z \otimes X \multimap Y$ there is a closed relation $\hat{R}: Z \multimap [X \Rightarrow Y]$ —though not necessarily a unique one—making the relevant diagram commute. The relation \hat{R} turns out to be a function, and as a function it is unique. Note that this is precisely the same situation as for the relation space in Rel , the category of sets and relations. We will come back to this universal property once we have set up the necessary machinery.

The Topological Relation Space

As we have already discussed at the beginning of this section, the hom-set $\text{StCp}^*(X, Y)$ is isomorphic to $\text{StCp}(X, \mathcal{K}(Y))$ and hence we can equip it with the topology inherited from the compact open topology on $\mathcal{K}(Y)^X$. We call this space $[X \Rightarrow Y]$, and by Proposition 3.2.9 this is a continuous Scott domain with the Scott topology which implies, in particular, that it is stably compact because of Proposition 1.3.19.

Let us rephrase the topology of this space in terms of relations rather than functions to the Smyth power domain. If $O \subseteq Y$ is open, then the set $\square O = \{K \in \mathcal{K}(Y) \mid K \subseteq O\}$ is Scott-open by Lemma 1.2.17. Hence we get the open neighbourhoods

$$\begin{aligned} N(K, \square O) &= \{f: X \rightarrow \mathcal{K}(Y) \mid f[K] \subseteq \square O\} \\ &\cong \{R: X \multimap Y \mid (\forall x \in K) x R y \implies y \in O\} \\ &= \{R: X \multimap Y \mid R[K] \subseteq O\} \end{aligned}$$

and we will soon prove that they form a subbasis of the topology.

For an open set $O \subseteq X$ and $K \subseteq Y$ compact saturated we can also consider the sets

$$N(O, \square K) := \{R: X \multimap Y \mid R[O] \subseteq K\},$$

where we use $\square K$ to denote $\{L \in \mathcal{K}(Y) \mid L \subseteq K\}$ by slight abuse of notation. As this set $\square K$ is simply $\uparrow K$ taken in $\mathcal{K}(K)$, where the order is reverse inclusion, Theorem 3.2.11 applies and we see that $N(O, \square K)$ is compact.

Next we study the order structure on the relation space:

Proposition 3.2.14. *The specialisation order on the space $[X \Rightarrow Y]$ is reverse inclusion. Furthermore, arbitrary suprema and finite infima are given by intersection and union, respectively.*

Proof. By Lemma 3.2.7, the inequality $R \sqsubseteq S$ holds for relations from X to Y if and only if for the corresponding functions $f_R, f_S: X \rightarrow \mathcal{K}(Y)$ we have

$$(\forall x \in X) f_R(x) \supseteq f_S(x).$$

This is equivalent to $R \supseteq S$.

The second claim follows immediately from the fact that closed sets are closed under intersections and finite unions. \square

Before we can tackle the universal property of the relation space we first have to turn \otimes into a bifunctor. To this end let us take two closed relations $R: X \multimap X'$ and $S: Y \multimap Y'$. We define $R \otimes S: X \otimes Y \multimap X' \otimes Y'$ by:

$$\langle x, y \rangle (R \otimes S) \langle x', y' \rangle \quad :\iff \quad x R x' \text{ and } y S y'$$

This assignment is clearly functorial, the only question is whether it defines a closed relation. If we have $x \not R x'$ we find $O \in \Omega(X)$ and $K \in \mathcal{K}(X')$ such that $O \times (X' \setminus K) \cap R = \emptyset$. This implies

$$(O \times Y) \times \underbrace{((X' \setminus K) \times Y')}_{(X' \times Y') \setminus (K \times Y')} \cap (R \otimes S) = \emptyset.$$

The case $y \not\mathcal{S} y'$ is argued analogously.

As we have discussed before we can think of a continuous function $f: X \rightarrow Y$ as a relation, namely its hypergraph $F \subseteq X \times Y$, by setting $x F y :\Leftrightarrow f(x) \sqsubseteq y$. This defines a faithful functor $i: \text{StCp} \hookrightarrow \text{StCp}^*$. Using this embedding we can see that \otimes extends the product bifunctor on StCp :

For $f: X \rightarrow X'$ and $g: Y \rightarrow Y'$ we get $f \times g: X \times Y \rightarrow X' \times Y'$, and its embedding is given by:

$$\langle x, y \rangle (i(f \times g)) \langle x', y' \rangle \iff \langle f(x), g(y) \rangle \sqsubseteq \langle x', y' \rangle$$

We can also embed f and g first and then compute $i(f) \otimes i(g)$ which yields the exact same relation since the order on the product is component-wise. This shows that the following diagram of functors commutes:

$$\begin{array}{ccc} \text{StCp}^2 & \xrightarrow{- \times -} & \text{StCp} \\ \downarrow i^2 & & \downarrow i \\ (\text{StCp}^*)^2 & \xrightarrow{- \otimes -} & \text{StCp}^* \end{array}$$

As a final preliminary consideration we study the composition of a continuous function f , or more precisely of the hypergraph of f , with a closed relation R . In the following we will use crossed arrows for closed relations and normal ones for continuous functions considered as relations. We claim that

$$X \xrightarrow{f} Y \xrightarrow{\text{---} R} Z$$

is simply given by

$$x (f \circ R) z \iff f(x) R z$$

where as before we write the composition of relations as the usual relational product, i.e. from left to right. The composition certainly contains all these pairs, and for the other inclusion it suffices to observe that $f(x) \sqsubseteq y R z$ implies $f(x) R z$.

Theorem 3.2.15. *The relation space construction $[X \Rightarrow -]$ is the right adjoint to the functor $i \circ (X \times -) = (X \otimes -) \circ i$.*

Proof. We begin by defining the evaluation relation $\epsilon: [X \Rightarrow Y] \otimes X \text{---} \rightarrow Y$:

$$\langle R, x \rangle (\epsilon) y \quad :\iff \quad x R y$$

To see that this is indeed a morphism suppose $x \not R y$. The relation R is closed in $X \times Y_\kappa$ by assumption and these two spaces are locally compact. Hence, we find compact neighbourhoods $K \ni x$ and $L \ni y$ such that $(K \times L) \cap R = \emptyset$. Note that L is a neighbourhood with respect to the co-compact topology on Y . From this we immediately get

$$\left(N(K, \square(Y \setminus L)) \times K \times L \right) \cap (\epsilon) = \emptyset$$

and the product is a neighbourhood of $\langle R, x, y \rangle$. This proves that ϵ is closed in $[X \Rightarrow Y] \times X \times Y_\kappa$.

Now we have to verify the details of the universal property indicated in the following diagram:

$$\begin{array}{ccc}
 [X \Rightarrow Y] & & [X \Rightarrow Y] \otimes X \xrightarrow{\epsilon} Y \\
 \uparrow \exists! \hat{R} & & \uparrow \hat{R} \otimes X \\
 Z & & Z \otimes X
 \end{array}$$

$\swarrow \forall R$

We define the function \hat{R} by:

$$x (\hat{R}(z)) y \quad :\iff \quad \langle z, x \rangle R y$$

Rather than checking all the details ‘by hand’ we use the fact that certain exponentials exist in StCp : The relation R corresponds to a function $f_R: X \times Z \rightarrow \mathcal{K}(Y)$. Now, X is locally compact and hence exponentiable which means that f_R has an exponential transpose. This transpose is a continuous function from Z to $\mathcal{K}(Y)^X$, namely \hat{R} up to the isomorphism $\mathcal{K}(Y)^X \cong [X \Rightarrow Y]$. This shows both that for all $z \in Z$ the relation $\hat{R}(z)$ is closed in $X \times Y_\kappa$ and that \hat{R} is continuous.

It follows from the discussion preceding the theorem that the function \hat{R} is uniquely determined. \square

The Relation Space Construction

Now suppose L and M represent X and Y , respectively. We generate $[L \Rightarrow M]$ freely from the tokens $(\phi \Rightarrow \psi)$ for $\phi \in L$ and $\psi \in M$. We have already proved that the two functions

$$\begin{aligned}
 o(\phi \Rightarrow \psi) & := N(\mathbf{K}[\phi], \square \mathbf{O}[\psi]) = \{R: X \multimap Y \mid R[\mathbf{K}[\phi]] \subseteq \mathbf{O}[\psi]\} \\
 k(\phi \Rightarrow \psi) & := N(\mathbf{O}[\phi], \square \mathbf{K}[\psi]) = \{R: X \multimap Y \mid R[\mathbf{O}[\phi]] \subseteq \mathbf{K}[\psi]\}
 \end{aligned}$$

define an open and a compact saturated reading. It is also clear that o and k satisfy the sanity condition $o(\phi \Rightarrow \psi) \subseteq k(\phi \Rightarrow \psi)$.

The proof of density of o and k , in the sense of Theorem 3.2.2.(2), is essentially the same as that for the function space. We have to modify the proof of Theorem 3.2.12 only slightly: Take a relation seen as a multi-function $f \in \mathcal{O} \in \Omega(\mathcal{K}(Y)^X)$. For every $x \in X$ and open set O containing the compact saturated set $f(x)$ we find a token $\psi \in M$ such that

$$f(x) \subseteq \mathbf{O}[\psi] \subseteq \mathbf{K}[\psi] \subseteq O.$$

This implies $x \in f^{-1}[\mathbf{O}[\psi]] \in \Omega(X)$ and hence we find a formula $\phi \in L$ such that

$$x \in \mathbf{O}[\phi] \subseteq \mathbf{K}[\phi] \subseteq f^{-1}[\mathbf{O}[\psi]].$$

This implies

$$f \in o(\phi \Rightarrow \psi) \subseteq k(\phi \Rightarrow \psi) \subseteq \uparrow(\mathbf{O}[\phi] \searrow \mathbf{K}[\psi]),$$

where by slight abuse notation we consider $o(\phi \Rightarrow \psi)$ and $k(\phi \Rightarrow \psi)$ as subsets of $\mathcal{K}(Y)^X$ rather than of the isomorphic $[X \Rightarrow Y]$. The saturated set $f(x)$ is the intersection of sets $\mathbf{K}[\psi]$ as constructed above which implies that f is the supremum of such step functions, and we can conclude the proof as in the case of the function space.

To complete the syntactic relation space construction we have to come up with rules that generate sufficiently many $\Vdash_{[L \Rightarrow M]}$ -sequents between atomic tokens. First, note that there are no such sequents with an empty right hand side because of the empty relation which is a member of every $k(\phi \Rightarrow \psi)$. Next we show that we can restrict the right hand side to singletons: Take a situation

$$k(\phi_1 \Rightarrow \psi_1) \cap \dots \cap k(\phi_m \Rightarrow \psi_m) \subseteq o(\sigma_1 \Rightarrow \tau_1) \cup \dots \cup o(\sigma_n \Rightarrow \tau_n)$$

where we cannot leave out any of the $o(\sigma_i \Rightarrow \tau_i)$ on the right. If $n > 1$ we can find closed relations R_1, \dots, R_n that lie in the intersection of the compact saturated sets, but $R_i \notin o(\sigma_i \Rightarrow \tau_i)$. The union $\bigcup_i R_i$ still satisfies all the conditions on the left, but it is not a member of any of the $o(\sigma_i \Rightarrow \tau_i)$, a contradiction. This proves that in such an irreducible situation we have $n = 1$, or in other words that, modulo weakening, we can restrict ourselves to singletons on the right.

This means that we have to characterise the sequents:

$$(\phi_1 \Rightarrow \psi_1), \dots, (\phi_n \Rightarrow \psi_n) \Vdash_{[L \Rightarrow M]} (\sigma \Rightarrow \tau) \quad (3.2)$$

The key is the following relation. We construct the following multi-function

$$f := (\mathbf{O}[\phi_1] \searrow \mathbf{K}[\psi_1]) \vee \dots \vee (\mathbf{O}[\phi_n] \searrow \mathbf{K}[\psi_n])$$

from X to $\mathcal{K}(Y)$. A component in this disjunction $(\mathbf{O}[\phi_i] \searrow \mathbf{K}[\psi_i])$ is the smallest function such that the corresponding relation lies in $k(\phi_i \Rightarrow \psi_i)$.

Hence, the relation $R_f : X \multimap Y$ corresponding to the supremum f of these step functions is the largest relation—in terms of subset inclusion rather than the specialisation order on $[X \Rightarrow Y]$ —that satisfies the premise of the above sequent. If the sequent (3.2) holds, then we must also have $R_f[\mathbf{K}[\sigma]] \subseteq \mathbf{O}[\tau]$.

We now study what this means for the relative position of $\mathbf{K}[\sigma]$ and $\mathbf{O}[\phi_1], \dots, \mathbf{O}[\phi_n]$ on the one side and for $\mathbf{O}[\tau]$ and $\mathbf{K}[\psi_1], \dots, \mathbf{K}[\psi_n]$ on the other. We first consider a trivial case, although it will later be subsumed by the general one: Suppose $\mathbf{K}[\sigma] \not\subseteq \mathbf{O}[\phi_1] \cup \dots \cup \mathbf{O}[\phi_n]$. Then R_f relates points from $\mathbf{K}[\sigma]$ to all of Y which implies $\mathbf{O}[\tau] = Y$. Because of $o(\perp \Rightarrow \tau) = o(\sigma \Rightarrow \top) = [X \Rightarrow Y]$ this implies $\Vdash_{[L \Rightarrow M]} (\sigma \Rightarrow \tau)$.

If $\mathbf{K}[\sigma]$ is covered by the $\mathbf{O}[\phi_i]$ then we proceed as follows. For each $x \in \mathbf{K}[\sigma]$ we determine the $\mathbf{O}[\phi_i]$ that x is an element of. If we get

$$x \in \mathbf{O}[\phi_{i_1}] \cap \dots \cap \mathbf{O}[\phi_{i_k}]$$

where x is not contained in any other $\mathbf{O}[\phi_j]$, then we have

$$f(x) = \mathbf{K}[\psi_{i_1}] \cap \dots \cap \mathbf{K}[\psi_{i_k}]$$

and thus

$$\mathbf{K}[\psi_{i_1}] \cap \dots \cap \mathbf{K}[\psi_{i_k}] \subseteq \mathbf{O}[\tau].$$

Putting it all together we get a covering of σ by a union of intersections of sets $\mathbf{O}[\phi_i]$ such that the corresponding set made up from the $\mathbf{K}[\psi_i]$ instead of the $\mathbf{O}[\phi_i]$ is contained in $\mathbf{O}[\tau]$.

This leads to the single rule:

$$\frac{\sigma \Vdash_L t(\phi_1, \dots, \phi_n) \quad t(\psi_1, \dots, \psi_n) \Vdash_M \tau}{(\phi_1 \Rightarrow \psi_1), \dots, (\phi_n \Rightarrow \psi_n) \Vdash_{[L \Rightarrow M]} (\sigma \Rightarrow \tau)}$$

where t is a term in n variables using the connectives \perp, \top, \vee and \wedge .

Since we allow the usage of the constants \perp and \top we have the ‘special’ case as an instance if we pick $t := \top$. The rule is clearly sound and our discussion shows that this rule is sufficient to derive all $\Vdash_{[L \Rightarrow M]}$ -sequents with a singleton on the right.

We can reformulate the rule in a way that does not resort to the logical connectives: Beginning from the two sequents

$$\sigma \Vdash_L t(\phi_1, \dots, \phi_n) \quad \text{and} \quad t(\psi_1, \dots, \psi_n) \Vdash_M \tau$$

we apply backwards rules until we end up with sequents

$$\begin{array}{ccc} \sigma \Vdash_L \Phi_1 & & \Psi_1 \Vdash_M \tau \\ \vdots & & \vdots \\ \sigma \Vdash_L \Phi_k & & \Psi_l \Vdash_M \tau \end{array}$$

such that the Φ_i are subsets of $\{\phi_1, \dots, \phi_n\}$ and the Ψ_j of $\{\psi_1, \dots, \psi_n\}$, i.e. the sequents no longer contain any logical connectives. We now claim that these sequents satisfy a slightly modified form of the side condition of (Cut*) (see page 61): Let $i: \{\phi_1, \dots, \phi_n\} \rightarrow \{\psi_1, \dots, \psi_n\}$ be the bijection that maps ϕ_1 to ψ_1 , ϕ_2 to ψ_2 , and so on. Using i we can express the condition as:

$$\left(\forall f \in \prod_i \Phi_i \right) (\exists j) \Psi_j \subseteq i[\{f_1, \dots, f_k\}] \quad (3.3)$$

To understand why the sequents we constructed earlier satisfy this condition we take a closer look at the Φ_i and Ψ_j . The sets Ψ_j give rise to the following disjunctive normal form of $t(\psi_1, \dots, \psi_n)$:

$$\bigwedge \Psi_1 \vee \dots \vee \bigwedge \Psi_l$$

The Φ_i , on the other hand, are more naturally understood as giving rise to a conjunctive normal form of $t(\phi_1, \dots, \phi_n)$. To get a disjunctive normal form we have to distribute all the conjunctions over the disjunctions. This leads precisely to choice functions:

$$\bigvee_{f \in \prod_i \Phi_i} (f_1 \wedge \dots \wedge f_k)$$

Both these disjunctive normal forms come essentially from the same term t . In general they need not be identical, but considered as elements of the free distributive lattice on n generators they must be equivalent. Thus, we see that the condition (*) must be satisfied.

Hence, we can give the following alternative rule for the relation space:

$$\frac{\begin{array}{ccc} \sigma \Vdash_L \Phi_1 & & \Psi_1 \Vdash_M \tau \\ \vdots & & \vdots \\ \sigma \Vdash_L \Phi_k & & \Psi_l \Vdash_M \tau \end{array}}{(\phi_1 \Rightarrow \psi_1), \dots, (\phi_n \Rightarrow \psi_n) \Vdash_{[L \Rightarrow M]} (\sigma \Rightarrow \tau)}$$

where the sets Φ_i and the Ψ_j are contained in $\{\phi_1, \dots, \phi_n\}$ and $\{\psi_1, \dots, \psi_n\}$, respectively, and they satisfy the above side condition (3.3).

Of course we could also formulate the side condition in a more symmetric fashion by adding the isomorphism i to the condition given in Lemma 2.1.2. Moreover, note that the possibility to have empty left or right hand sides in the sequents allows us to dispense with the constants \perp and \top .

As before our derivation of the rule shows that it is complete. To prove that it is sound we only have to translate what the side condition means for the disjunctive normal forms: It says that for every conjunction term of ϕ_i appearing there is a conjunction of ψ_j that contains at most the formulae corresponding to the ϕ_i . This implies that if the normal form on the left is given by $s(\phi_1, \dots, \phi_n)$, and the one on the right by $t(\psi_1, \dots, \psi_n)$ then we have $s \leq t$ in the free distributive lattice. From this we infer $\mathbf{K}[\sigma] \subseteq s(\mathbf{O}[\phi_1], \dots, \mathbf{O}[\phi_n])$ and

$$s(\mathbf{K}[\psi_1], \dots, \mathbf{K}[\psi_n]) \subseteq t(\mathbf{K}[\psi_1], \dots, \mathbf{K}[\psi_n]) \subseteq \mathbf{O}[\tau]$$

which implies soundness.

3.3 Relations and Functions


As our final topic illuminating the relationship between the function space and the relation space we come back to the last problem of Section 3.1.4, namely characterising the ‘functions’ among the morphisms of StCp^* . In this section we do it in categorical terms.

The idea is simple: Take a category with finite products. If we choose binary products and a terminal, then the products give rise to a symmetric monoidal structure on the category. Moreover, the unique maps to the terminal and the diagonals $\Delta_A: A \rightarrow A \times A$ are natural transformations. Now consider the relationship between Set and Rel . Cartesian products in Set can be extended to a bifunctor on Rel although they are no longer categorical products. Moreover, a relation is a function if and only if the diagonals and the unique functions into the chosen singleton are natural with respect to it. This suggests to define ‘functions’ in these terms.

The plan of this section is as follows: We begin by introducing categories with diagonals. In such a category we can define an intrinsic notion of a morphism being a function. We then explore how this relates to the standard construction of a category of relations from any regular category. All this material is essentially category theory folklore. Finally we apply this machinery to our category StCp .


3.3.1 Diagonals

Suppose \mathcal{C} is a symmetric monoidal category. Recall that a commutative comonoid in \mathcal{C} is an object $A \in |\mathcal{C}|$ together with morphisms $t: A \rightarrow I$ and $\Delta: A \rightarrow A \otimes A$ that makes the three diagrams



commute. Here I is the unit for the tensor, the unlabelled arrows are the isomorphisms that are part of the definition of the monoidal structure and $\kappa_{A,A}$ is the isomorphism that makes \mathcal{C} symmetric monoidal.

Definition 3.3.1. A *diagonal structure* on a symmetric monoidal category \mathcal{C} is a family of commutative comonoids $\langle t_A, \Delta_A \rangle_{A \in |\mathcal{C}|}$, indexed by the objects of \mathcal{C} , that respects the symmetric monoidal structure of \mathcal{C} in the sense that the following diagrams commute:



We also call a symmetric monoidal category with a chosen diagonal structure a *category with diagonals*.

The guiding example is Rel , the category of sets with relations as morphisms. This category is self dual since for every relation R from X to Y the opposite relation $\{ \langle x, y \rangle \mid y R z \}$ is a morphism from Y to X and taking the opposite relation is clearly functorial. As the self duality fixes objects, products and coproducts agree, and in fact they are given by disjoint union. Cartesian products and any terminal, say $I := \{*\}$, give rise to a symmetric monoidal structure on Rel : For objects we set $X \otimes Y := X \times Y$ and for morphisms we take the relations component-wise:

$$\langle x, y \rangle R \otimes S \langle x', y' \rangle \quad :\iff \quad x R x' \text{ and } y S y'$$

In the following we consider the following diagonal structure on Rel :

$$\begin{aligned} t_A &:= A \times \{*\} = \{ \langle a, * \rangle \mid a \in A \} \\ a \Delta_A \langle b, c \rangle &:\iff a = b = c \end{aligned}$$

The t_A and Δ_A are readily seen to define commutative comonoids and to satisfy the conditions of the previous definition. We will see shortly that this is a consequence of the fact that in Set the tensor is just the categorical product and the singleton is a terminal.

Let us return to the general situation. Without additional assumptions neither $(t_A)_A$ nor $(\Delta_A)_A$ is a natural transformation. They are if and only if the symmetric monoidal structure on \mathcal{C} is given by finite products. One direction of the proof is easy: If \mathcal{C} has (chosen) finite products, then we can take the unique arrows into the terminal as t_A and let $\Delta_A := \langle A, A \rangle$. These morphisms are always natural since they are determined by a universal property and they define a diagonal structure with respect to binary products and the terminal. The converse is the content of the following proposition.

Proposition 3.3.2. 1. If $(t_A)_A$ is a natural transformation from the identity on \mathcal{C} to the constant I functor, then I is a terminal.

2. If in addition $(\Delta_A)_A$ is natural from the identity functor \mathcal{C} to $(- \otimes -)$, then \otimes is the categorical product.

Proof. The object I is weakly terminal, and if $f: A \rightarrow I$ is any arrow to I then naturality means

$$\begin{array}{ccc}
 A & \xrightarrow{t_A} & I \\
 f \downarrow & & \parallel \\
 I & \xrightarrow{t_I} & I
 \end{array}$$

which implies $f = t_A$.


Suppose $\Delta: \mathcal{C} \rightarrow (- \otimes -)$ is natural. To show $\otimes = \times$ we first have to define the projections:

$$\begin{array}{ccccc}
 A \otimes I & \xleftarrow{A \otimes t_B} & A \otimes B & \xrightarrow{t_A \otimes B} & I \otimes B \\
 \cong \downarrow & \swarrow \pi_0 & & \searrow \pi_1 & \downarrow \cong \\
 A & & & & B
 \end{array}$$

If we are given two morphisms $f: C \rightarrow A$ and $g: C \rightarrow B$ we define $\langle f, g \rangle$ to be the composition


$$\langle f, g \rangle: C \xrightarrow{\Delta_C} C \otimes C \xrightarrow{f \otimes g} A \otimes B.$$

Now, we have to consider the composition with the projections:



The top square commutes by naturality, the bottom one since we already know that I is a terminal, and the triangle on the right is one of the axioms defining comonoids. Hence, we have shown $\pi_0 \circ \langle f, g \rangle = f$, and the proof for $\pi_1 \circ \langle f, g \rangle = g$ is analogous.

It remains to show that $\langle f, g \rangle$ is unique. For another mediating arrow $h: C \rightarrow A \otimes B$ we consider:



For most sub-diagrams commutativity is obvious. The rightmost rectangle is the tensor of the equations $(A \otimes t) \circ \Delta = A$ and $(t \otimes B) \circ \Delta = B$, both of which are instances of one of the axioms for comonoids. To understand


$$(A \otimes t_A) \otimes (t_B \otimes B) \circ i = (A \otimes t_B) \otimes (t_A \otimes B)$$

we have to take a closer look at the isomorphism i . It is made up from the isomorphisms saying that \otimes is associative and commutative. These isomorphisms are natural transformations by definition. Note that $I \otimes I \cong I$ and


so $I \otimes I$ is also terminal which implies that the commutativity morphism $\kappa_{I,I}: I \otimes I \rightarrow I \otimes I$ is the identity. This concludes the proof that the bottom middle triangle commutes. If we trace the arrows at the lower edge of the diagram we see that we have shown $h = \langle f, g \rangle$. \square

We now use the lack of naturality of $(t_A)_A$ and $(\Delta_A)_A$ to define an abstract notion of function.

Definition 3.3.3. We say that a morphism $f: A \rightarrow B$ is *total* if t is natural with respect to it, or in other words if



commutes. It is *deterministic* if Δ is natural for it, i.e. if the diagram



commutes.


If f is total and deterministic, then we call it *functional* or a *function*, and we denote the subcategory of functions by $\text{Fun}(C)$.

The following lemma ensures that $\text{Fun}(C)$ is indeed a category.

Lemma 3.3.4. *The classes of total and of deterministic morphisms are closed under composition and contain all identity morphisms.*

Proof. All claims follow immediately from the definitions. \square

To get a better understanding of the properties that we have just defined we investigate what they mean for our principal example Rel . A relation $R: X \dashrightarrow Y$ satisfies



if and only if for every $x \in X$ there is a $y \in Y$ such that $x R y$. This is equivalent to R being a total relation in the usual sense.

For determinism consider:

$$\begin{array}{ccc}
 X & \xrightarrow{\Delta_X} & X \otimes X \\
 R \downarrow & & \downarrow R \otimes R \\
 Y & \xrightarrow{\Delta_Y} & Y \otimes Y
 \end{array}$$

For $x R y \Delta_Y \langle y, y \rangle$ we clearly also have $x \Delta_X \langle x, x \rangle (R \otimes R) \langle y, y \rangle$. But if conversely

$$x \Delta_X \langle x, x \rangle (R \otimes R) \langle y, y' \rangle$$

holds, then we can find a y'' satisfying

$$x R y'' \Delta_Y \langle y, y' \rangle$$

only if $y = y' = y''$. This means that R is deterministic as a morphism if and only if it is a single-valued relation. Putting the two things together we see that in Rel the functional relations are precisely those that are functions in the classical, concrete sense.

All structure maps are functional, but since we do not need this fact later on, we simply state the result without proof:

Proposition 3.3.5. *Let C be a symmetric monoidal category with diagonals. Then the structural natural isomorphisms that make C symmetric monoidal, as well as all arrows t_A and Δ_A of the diagonal structure are functional.*

With the new terminology we can reformulate Proposition 3.3.2 as follows:

Corollary 3.3.6. *If C is a symmetric monoidal category with diagonals, then the category of functions $\text{Fun}(C)$ is the largest subcategory of C such that the tensor— together with the diagonal structure—gives rise to a product.*

3.3.2 Regular Categories

There is a standard way of talking about relations in an arbitrary category. A relation from A to B is simply a subobject $R \hookrightarrow A \times B$. For two arrows $u: C \rightarrow A$ and $v: C \rightarrow B$ the mediating morphism $\langle u, v \rangle: C \rightarrow A \times B$ is monic if and only if u and v are jointly monic. Hence, we can equivalently think of relations as jointly monic pairs $r_0: R \rightarrow A$ and $r_1: R \rightarrow B$. To be


more precise we should say that a relation is an equivalence class of such pairs; the equivalence is given by isomorphisms that make the obvious diagram commute.

To define a proper composition of relations we need more structure on the category \mathcal{C} ; it has to be a regular category. We quickly review the basic definitions and results. Our exposition follows [McL92, Chapter 25], in particular we will make use of generalised elements which will be introduced shortly.

An arrow $q: A \rightarrow B$ is *surjective* or a *surjection* if the only subobject of B it factors through is B . In the following we will use arrows of the shape \longrightarrow to denote surjections. If \mathcal{C} has equalisers then all surjections are epi. A *surjective image* of an arrow $f: A \rightarrow B$ is a factorisation $A \xrightarrow{q} Q \xrightarrow{e} B$ such that q is surjective and e is monic. A *regular category* is a category with finite limits such that all morphisms have surjective images and surjections are stable under pullback.

The meaning of these definitions will become clearer when we use them to compose relations. But to explain this composition it is useful to introduce the language of generalised elements first.

A *generalised element* is simply an arrow $x: A \rightarrow B$. We write $x \in_A B$ and call A the *stage of definition* of the element x . Given another morphism $f: B \rightarrow C$, we can apply it to x to get the generalised element $f \circ x \in_A C$ which we also write as $f(x) \in_A C$. If $x \in_A B$ is a generalised element and $i: C \rightarrow B$ monic, then we write $x \in i$ if x factors through i :



In this case we also say that the element x is a *member* of i . Provided that i is implicitly understood we may also write $x \in C$ and say x is a member of C . One use of this notation lies in the following:

Proposition 3.3.7. *Let $i: A \rightarrow B$ and $i': A' \rightarrow B$ be two monics. Then we have $i \subseteq i'$, i.e. i factors through i' and hence is also a subobject of i' , if and only if all generalised elements that are members of i are members of i' .*

Proof. The “only if” part follows immediately from the definitions. For the other implication we consider the generalised element $i \in_A B$ which is clearly a member of A . If this implies that i is also a member of i' , then this says precisely that i factors through i' . \square


As an immediate consequence we see that generalised elements characterise subobjects:

Corollary 3.3.8. *Two monics $i: A \rightarrow B$ and $i': A' \rightarrow B$ represent the same subobject of B if and only if they have the same generalised elements.*

Before we come back to the composition of relations we have to investigate how membership of a generalised element in a subobject is affected by changing the stage of definition. Suppose $x \in_A B$ is a generalised element which is a member of a subobject $C \rightarrow B$. For a morphism $y: A' \rightarrow A$ we get a generalised element $x \circ y \in_{A'} B$, and this element clearly satisfies $x \circ y \in C$. The converse is false in general but holds if the earlier stage A' covers A , i.e. if y is surjective:

Proposition 3.3.9. *Let C be a subobject of B in a regular category. If $x \in_A B$ is a generalised element and $y: A' \rightarrow A$ surjective, then $x \circ y \in C$ implies $x \in C$.*

Proof. We take the pullback of the monic $C \rightarrow B$ and get the situation:



Since $x \circ y$ factors through C we get a unique arrow d as in the diagram. Pullbacks of monics are monic and as y is surjective we infer that the monic at the top of the square is an isomorphism. This implies that x factors through C . □

Putting the proposition and the discussion preceding it together we can say that membership of generalised elements in subobjects is invariant under going to covering stages of definition.

Remark. *We can also read this proposition as the non-trivial part of the proof that in a regular category the surjections and the monics form a factorisation system. Both these classes of arrows contain all isomorphisms and are closed under composition with isomorphisms. Every morphism factors as a surjection followed by a monic by definition, and the diagonal fill-in property is exactly the property given in the proposition.*


The relevance of the proposition is the following. Suppose R is a relation from A to B and S another one from B to C . Trying to mimic the situation in Rel we might try to define the composition of R and S as the subobject of $A \times C$ that contains a pair $\langle a, c \rangle \in_X A \times C$ if and only if there

is an element $b \in_X B$ such that $\langle a, b \rangle \in R$ and $\langle b, c \rangle \in S$. Unfortunately, it is possible that we cannot find such a b at stage X but that we can at a covering stage $X' \twoheadrightarrow X$, contradicting the previous proposition. This means that the best we can try is to define the elements of the composition at stage X as those $\langle a, c \rangle \in_X A \times C$ that satisfy:

$$(\exists X' \twoheadrightarrow X, b \in_{X'} B) \langle a \circ u, b \rangle \in R \text{ and } \langle b, c \circ u \rangle \in S \quad (3.4)$$

If such a subobject of $A \times C$ exists, then Corollary 3.3.8 tells us that this condition determines it uniquely. In the following we show that there is such an object in \mathbf{C} using the constructions provided by the definition of a regular category.

We first take the pullback of r_1 and s_0 :




This gives us a map from P to $A \times C$ which we factorise as

$$P \twoheadrightarrow R \circ S \twoheadrightarrow A \times C.$$

The monic part of this factorisation gives us a subobject of $A \times C$. By composing with the projections we can also consider it as a jointly monic pair of morphisms from $R \circ S$ to A and C as indicated in the diagram.


Lemma 3.3.10. *The composition $R \circ S$ contains precisely the generalised elements given in (3.4).*

Proof. Suppose we have a surjection $u: X' \twoheadrightarrow X$ and a generalised element $b \in_{X'} B$ such that $\langle a \circ u, b \rangle \in R$ and $\langle b, c \circ u \rangle \in S$. By the definition of the pullback this yields a unique arrow $d: X' \rightarrow P$ such that




commutes. We infer that $\langle a \circ u, c \circ u \rangle$ factors through $R \circ S$ which implies $\langle a, c \rangle \in R \circ S$ by Proposition 3.3.9.

For the converse we take the pullback as indicated in the diagram:




This gives us the required generalised element $b \in_{X'} B$. □

As an immediate consequence we see that composition of relations is associative; the proof is essentially the one for Rel expressed in terms of generalised elements. It is also clear that the relations



act as identities. From now on we let $\text{Rel}(\mathcal{C})$ denote the category with the same objects as \mathcal{C} and relations as morphisms. We can embed the original category \mathcal{C} in $\text{Rel}(\mathcal{C})$ by taking a morphism $f: A \rightarrow B$ to its graph Γ_f :



At every stage of definition the graph contains exactly the generalised elements $\langle a, f(a) \rangle$. This immediately implies that the embedding is functorial and faithful because we can test it on the generic element $A \in_A A$.

Conversely, if for a relation $\langle r_0, r_1 \rangle: R \rightarrow A \times B$ the arrow r_0 is an isomorphism, then the relation is equivalent to one of the form Γ_f and hence corresponds to a ‘function’ in the original category.

The category $\text{Rel}(\mathcal{C})$ has a lot of additional structure; it is in particular, an allegory. For such categories there is a standard way of defining a notion of ‘function’. As it turns out, in $\text{Rel}(\mathcal{C})$ they are precisely the morphisms embedded from the original category \mathcal{C} . For more details see [FS90, 2.132].

In the previous section we discussed an alternative version of defining ‘functions’ in a given category, namely via a diagonal structure. The reason

for this is that some categories that look like categories of relations are not allegories. One such example is $\text{MLS} \cong \text{StCp}^*$. It satisfies most axioms of an allegory, but its self-duality does not fix objects. In the following we will show that for $\text{Rel}(C)$, instead of the allegory structure, we can also use an induced diagonal structure to recover C .

Having assumed that C has finite limits we can extend the products to a symmetric tensor of $\text{Rel}(C)$: We choose any terminal object 1 in C as I , and for the tensor on objects we simply define $A \otimes B := A \times B$. Since we can embed the symmetric monoidal structure given by products in C into the new category $\text{Rel}(C)$, it is clear that this makes the latter category symmetric monoidal if we can define the tensor on arbitrary relations. Suppose R is a relation from A to A' , and S one from B to B' : We construct

$$\begin{array}{ccc} & R \otimes S = R \times S & \\ \begin{array}{c} \nearrow \\ \tau_0 + s_0 \end{array} & & \begin{array}{c} \searrow \\ \tau_1 + s_1 \end{array} \\ A \times B & & A' \times B' \end{array}$$

which is a jointly monic pair of arrows and hence represents a relation from $A \otimes B$ to $A' \otimes B'$. In terms of generalised elements this relation is characterised by

$$\langle \langle x, y \rangle, \langle x', y' \rangle \rangle \in R \otimes S \iff \langle x, x' \rangle \in R \quad \text{and} \quad \langle y, y' \rangle \in S$$

at all stages of definition. From this description it is clear that $(- \otimes -)$ is a bifunctor and that it extends products, i.e. $\Gamma_f \otimes \Gamma_g = \Gamma_{f \times g}$.

The category $\text{Rel}(C)$ also inherits the diagonal structure from the original category. We set


$$\begin{aligned} t &:= \Gamma_! \quad \text{and} \\ \Delta_A &:= \Gamma_{\langle A, A \rangle} \end{aligned}$$

where $!$ is the unique arrow from an object into the chosen terminal. It is clear that this satisfies the axioms of a diagonal structure given in Definition 3.3.1 since all maps are embedded from C where the corresponding diagrams commute. Using the same argument we see that all morphisms Γ_f are functional with respect to this diagonal structure. The converse also holds:

Theorem 3.3.11. *The functional morphisms in $\text{Rel}(C)$ are exactly those of the form Γ_f , or in other words $\text{Fun}(\text{Rel}(C)) \cong C$.*

Proof. We only have to show that all functional relations in $\text{Rel}(C)$ come from a morphisms in C . Suppose R , given by the two maps $r_0: R \rightarrow A$

and $r_1: R \rightarrow B$, is such a relation. From totality we infer that we have a commuting diagram



which implies that r_0 is surjective. We now show that it is also a monic and hence an isomorphism. Suppose we have two parallel arrows $x, y: X \rightarrow R$ such that $r_0 \circ x = r_0 \circ y =: a$, then we consider the generalised elements $\langle a, r_1 \circ x \rangle, \langle a, r_1 \circ y \rangle \in_X A \times B$. By construction we have

$$\begin{aligned} & \langle a, r_1 \circ x \rangle, \langle a, r_1 \circ y \rangle \in R \\ \implies & \langle \langle a, a \rangle, \langle r_1 \circ x, r_1 \circ y \rangle \rangle \in R \otimes R \\ \implies & \langle a, \langle r_1 \circ x, r_1 \circ y \rangle \rangle \in \Delta_A \circ (R \otimes R) = R \circ \Delta_B \\ \implies & (\exists u: X' \longrightarrow X, b \in_{X'} B) \langle a \circ u, b \rangle \in R \\ & \text{and } \langle b, \langle r_1 \circ x \circ u, r_1 \circ y \circ u \rangle \rangle \in \Delta_B. \end{aligned}$$

From this we infer $b = r_1 \circ x \circ u = r_1 \circ y \circ u$, and as u is epi this means $r_1 \circ x = r_1 \circ y$. The morphisms r_0 and r_1 are jointly monic which allows us to conclude that x and y are equal, thus showing that r_0 is monic. Since we have seen before that it is also a cover it must be an isomorphism and hence the relation R is a graph Γ_f . \square

3.3.3 Closed Relations

Now, we apply this machinery to our category StCp^* . Although, it is not of the form $\text{Rel}(\text{StCp})$ we still have an embedding from the cartesian category StCp to StCp^* , and the tensor extends the product as discussed in Section 3.2.4. This implies that we can simply embed the diagonal structure as well. Explicitly, the diagonals Δ_X are given by

$$x \Delta_X \langle x', x'' \rangle \quad :\iff \quad x \sqsubseteq x', x''$$

and the nullary equivalent t_X relate all elements in X to the one element of I .

As in the previous section it is also clear that the embedding of a continuous function from StCp is functional in StCp^* . The converse is a refinement of the proof for Rel .

Theorem 3.3.12. *The functional morphisms in StCp^* are precisely the hypergraphs of continuous functions between stably compact spaces.*

Proof. We only have to prove that a functional closed relations come from continuous functions. So, let us take such a relation $R: X \dashrightarrow Y$. If it is total, then for all $x \in X$ there is a $y \in Y$ such that $x R y$. If we consider the corresponding function $f_R: X \rightarrow \mathcal{K}(Y)$ this means $f_R(x) \neq \emptyset$, for all $x \in X$. Let us fix an x and suppose y, y' are in the compact saturated set $f_R(x)$. We get

$$x \Delta_X \langle x, x \rangle R \otimes R \langle y, y' \rangle$$

and if R is deterministic there must be a $y'' \in Y$ such that

$$x R y'' \Delta_Y \langle y, y' \rangle$$

which implies $y'' \sqsubseteq y, y'$. This shows that for a functional R the set $f_R(x)$ is filtered, or directed with respect to the co-compact topology.

The space Y_κ is sober by Corollary 1.3.12 and thus Proposition 1.2.9 implies that the directed supremum of $f_R(x)$ exists and that it does not lie in the open set $Y \setminus f_R(x)$. In terms of the original space Y this says that $f_R(x)$ has a least element, i.e. it is a principal filter. We have thus shown that R corresponds to a continuous function from X to Y . \square

This means that the categorical characterisation of functions agrees with the concrete one given at the end of Section 3.1.4.

Conclusion

Let us come back to the aims laid out in the introduction to see what has been achieved and which questions are still open.

The results on cut elimination of Section 2.2 and the representation Theorem 3.2.2 allow us to perform a number of domain constructions in logical form and they are defined for all stably compact spaces. This class contains the bifinite domains of Abramsky's theory as well as the FS domains which include the continuous domains most commonly studied in semantics. Of the power domain constructions we have only considered the Smyth power domain, and so the Hoare and the Plotkin power domains are canonical candidates to look at. A more interesting object of study, however, is the probabilistic power domain since there is no equivalent for algebraic domains. The probabilistic power domain for a stably compact *domain* is again stably compact [JT98] but, apart from this, very little is known about its structure. It is not even clear whether we can extend the construction to arbitrary stably compact spaces or not. This means that performing this construction in logical form will be difficult. We also have not looked at bilimits, but as it is straightforward in Abramsky's case this should be quite easy. To tackle them in our setup means that we have to characterise the embeddings in MLS.

The situation for function spaces is a very unsatisfactory. It is obvious that we cannot have functions spaces of arbitrary stably compact spaces, but we can for subclasses like FS domains. At the moment we are not even able to perform the syntactic construction for the much more restrictive class of continuous Scott domains. The problem is Abramsky's primality predicate that, in the algebraic case, allows him to talk about points which is necessary for his construction. It is not clear how to translate that to our setup or what other modification to the system could replace it.

This has a direct connection to the question whether it is possible to do domain theory in *purely* logical form, i.e. without referring back to the semantics. The construction of products and coproducts in Section 2.2.3 show how this could be done. Because of our categorical characterisation of relation spaces in Theorem 3.2.15 and the syntactic characterisation of functions given in Theorem 3.1.44 something similar should be possible for the relation space construction. The problem, again, is the function space,

and the first step is to translate Abramsky's construction for algebraic domains into a form that can be expressed in RMLS. As it is not known to what extent this can be done, the question remains open.

As we have seen in Section 3.1.3, an MLS formula has an open and a compact interpretation. It has been argued in [JS96] that compact and open sets play dual rules, and the discussion of consistency in Section 3.1.1 shows that the compact interpretation can be understood as the *negative* information contained in the formula. This poses the question what this signifies on the syntactic side, and in particular what it means for the semantics of a programming language using this logic.

Compared to the function spaces, the relation space construction is very smooth. This may be an indication that it could be beneficial to focus more on relational rather than on functional semantics, as is traditionally done in denotational semantics. It remains to be seen what exactly can be expressed in such a semantics and what the possible benefits are.

Index

Categories

For the hierarchy of full subcategories of Top , the category of topological spaces, see Figure 1.2 on page 48. It contains all categories listed in (1) and (2) below.

1. Categories of dcpo's with Scott-continuous functions

AlgScott	algebraic Scott domains, 21
Alg	algebraic domains, 19
DCPO	dcpo's, 19
Dom	continuous domains, 19
FS	FS domains, 22
ΛCp	Lawson-compact domains, 44
RSFP	retracts of bifinite domains, 22
Scott	continuous Scott domains, 21
SFP	bifinite (or SFP) domains, 21
$\frac{2}{3}\text{SFP}$	$2/3$ -bifinite domains, 50

2. Categories of topological spaces with continuous functions

CpOpen	spaces with a basis of compact opens, 34
LocCp	locally compact spaces, 33
Sob	sober spaces, 27
Spec	spectral spaces, 49
StCp	stably compact spaces, 38
Stone	Stone spaces, 49
Top	topological spaces

3. Other categories

ASL	arithmetic semilattices with Scott-continuous semilattice morphisms, 107
CLat	complete lattices with frame homomorphisms, 25
Frm	frames with frame morphisms, 24
Loc	locales, the opposite category of Frm, 24

MLS	continuous sequent calculi with compatible consequence relations, 67
RMLS	reflexive sequent calculi (i.e. with identity axiom) with compatible consequence relations, 69
Rel	sets with relations
SPL	strong proximity lattices with approximable relations, 100
SPL _w	strong proximity lattices with weak approximable relations, 100
Set	sets with functions
StCp*	stably compact spaces with closed relations, 123
StCp _K	Kleisli category for the Smyth power monad on stably compact spaces, 120

Notation

1. Order

\mathfrak{P} , 31
 \sqsubseteq , 16
 \uparrow , 17
 \downarrow , 16
 \sqcup^\uparrow , 19
 \ll , 19
 \uparrow , 20
 \downarrow , 19
 $K(X)$, 19
 \lll , 47
 \lesssim , 51
 \approx , 51
 \prec , 98
 \equiv , 101

2. Topology

$L(\mathcal{U})$, 36
 $N(x)$, 16
 $N^\circ(x)$, 25
 $\Omega(X)$, 17
 $\mathcal{K}(X)$, 17
 $\square O$, 34
 $\square K$, 136
 $K\Omega(X)$, 34
 $\Sigma(X)$, 19
 X_κ , 17
 X_π , 17
 X^\uparrow , 41
 O_x , 26
 \mathcal{O}_I , 111
 \mathcal{O}_ϕ , 112
 \mathcal{K}_F , 111
 \mathcal{K}_ϕ , 112
 $\mathbf{O}[\![\phi]\!]$, 132
 $\mathbf{K}[\![\phi]\!]$, 132
 $\hat{\mathcal{O}}_\phi$, 117
 $\llbracket \cdot \rrbracket$, 52
 $\mathbf{2}$, 24
 X_\perp , 42

$X \otimes Y$, 146
 $\mathcal{P}_S(X)$, 136
 Y^X , 137
 $N(A, B)$, 141
 $(O \searrow y)$, 138
 $[X \Rightarrow Y]$, 145

3. MLS

\vdash , 60
 \Vdash , 65
 \circ , 61
 $(\cdot)^{\text{op}}$, 89
 $X[\vdash]$, 92
 $\phi[\vdash]$, 92
 $[\vdash]X$, 92
 $[\vdash]\phi$, 92
 $[\phi]$, 74, 101
 R^+ , 70
 B^+ , 70
 $\vdash|_B^C$, 70
 $r(\Gamma)$, 70
 $rc(\Gamma)$, 70
 $g(\Gamma)$, 70

4. Named

filt , 92
 $\text{Fun}(\mathbf{C})$, 158
 idl , 92
 pt , 25
 $\text{Rel}(\mathbf{C})$, 163
 spec , 49, 111

Bibliography

- [Abr87] S. Abramsky. *Domain Theory and the Logic of Observable Properties*. PhD thesis, University of London, 1987.
- [Abr90] S. Abramsky. The lazy lambda calculus. In D. Turner, editor, *Research Topics in Functional Programming*, pages 65–117. Addison Wesley, 1990.
- [Abr91a] S. Abramsky. A domain equation for bisimulation. *Information and Computation*, 92:161–218, 1991.
- [Abr91b] S. Abramsky. Domain theory in logical form. *Annals of Pure and Applied Logic*, 51:1–77, 1991.
- [AJ94] S. Abramsky and A. Jung. Domain theory. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 3, pages 1–168. Clarendon Press, 1994.
- [Bir67] G. Birkhoff. *Lattice Theory*, volume 25 of *AMS Colloq. Publ.* American Mathematical Society, Providence, third edition, 1967.
- [Bon96] M. Bonsangue. *Topological Dualities in Semantics*. PhD thesis, Vrije Universiteit Amsterdam, November 1996. 258pp.
- [Bou88] N. Bourbaki. *General Topology—Chapters 1–4*. Elements of Mathematics. Springer Verlag, 1988.
- [Dij76] E. W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, Englewood Cliffs, New Jersey, 1976.
- [DP90] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, 1990.
- [Eda94] A. Edalat. Domain theory and integration. In *Ninth Annual IEEE Symposium on Logic in Computer Science*, pages 115–124. IEEE Computer Society Press, 1994.

- [Eda95] A. Edalat. Dynamical systems, measures and fractals via domain theory. *Information and Computation*, 120(1):32–48, 1995.
- [EE96] A. Edalat and M. H. Escardó. Integration in real PCF. To appear in LICS'96, 1996.
- [EH95] A. Edalat and R. Heckmann. Computational models for metric spaces, October 1995. Manuscript.
- [Esc96] M. H. Escardó. PCF extended with real numbers. *Theoretical Computer Science*, 162:79–115, 1996.
- [Esc98] M. H. Escardó. Properly injective spaces and function spaces. *Topology and Its Applications*, 89(1–2):75–120, November 1998.
- [FS90] P. J. Freyd and A. Scedrov. *Categories, Allegories*. North-Holland, 1990.
- [Gen34] G. Gentzen. Untersuchungen über das logische Schließen. *Mathematische Zeitschrift*, 39:176–210, 405–431, 1934. (Reprint 1969, Wissenschaftliche Buchgesellschaft Darmstadt).
- [GHK⁺80] G. Gierz, K. H. Hofmann, K. Keimel, J. D. Lawson, M. Mislove, and D. S. Scott. *A Compendium of Continuous Lattices*. Springer Verlag, 1980.
- [GLS83] G. Gierz, J. D. Lawson, and A. Stralka. Quasicontinuous posets. *Houston Journal of Mathematics*, 9:191–208, 1983.
- [HM81] K. H. Hofmann and M. Mislove. Local compactness and continuous lattices. In B. Banaschewski and R.-E. Hoffmann, editors, *Continuous Lattices, Proceedings Bremen 1979*, volume 871 of *Lecture Notes in Mathematics*, pages 209–248. Springer Verlag, 1981.
- [Hoa69] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12:576–580, 1969.
- [Hof81] R.-E. Hoffmann. Continuous posets, prime spectra of completely distributive complete lattices, and Hausdorff compactification. In B. Banaschewski and R.-E. Hoffmann, editors, *Continuous Lattices, Proceedings Bremen 1979*, volume 871 of *Lecture Notes in Mathematics*, pages 159–208. Springer Verlag, 1981.
- [JKM97] A. Jung, M. Kegelmann, and M. A. Moshier. Multi lingual sequent calculus and coherent spaces. In S. Brookes and M. Mislove, editors, *13th Conference on Mathematical Foundations of Programming Semantics*, volume 6 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science Publishers B.V., 1997. 18 pages.

- [JKM99] A. Jung, M. Kegelmann, and M. A. Moshier. Multi lingual sequent calculus and coherent spaces. *Fundamenta Informaticae*, 1999. (Accepted for publication.).
- [Joh81] P. T. Johnstone. Scott is not always sober. In B. Banaschewski and R.-E. Hoffmann, editors, *Continuous Lattices, Proceedings Bremen 1979*, volume 871 of *Lecture Notes in Mathematics*, pages 282–283. Springer Verlag, 1981.
- [Joh82] P. T. Johnstone. *Stone Spaces*, volume 3 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1982.
- [Joh83] P. T. Johnstone. The point of pointless topology. *Bulletin of the American Mathematical Society*, 8:41–53, 1983.
- [Jon90] C. Jones. *Probabilistic Non-Determinism*. PhD thesis, University of Edinburgh, Edinburgh, 1990. Also published as Technical Report No. CST-63-90.
- [JP89] C. Jones and G. Plotkin. A probabilistic powerdomain of evaluations. In *Logic in Computer Science*, pages 186–195. IEEE Computer Society Press, 1989.
- [JS96] A. Jung and Ph. Sünderhauf. On the duality of compact vs. open. In S. Andima, R. C. Flagg, G. Itzkowitz, P. Misra, Y. Kong, and R. Kopperman, editors, *Papers on General Topology and Applications: Eleventh Summer Conference at the University of Southern Maine*, volume 806 of *Annals of the New York Academy of Sciences*, pages 214–230, 1996.
- [JS98] A. Jung and Ph. Sünderhauf. Uniform approximation of topological spaces. *Topology and its Applications*, 83(1):23–38, 1998.
- [JT98] A. Jung and R. Tix. The troublesome probabilistic powerdomain. In A. Edalat, A. Jung, K. Keimel, and M. Kwiatkowska, editors, *Proceedings of the Third Workshop on Computation and Approximation*, volume 13 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science Publishers B.V., 1998. 23 pages.
- [Jun89] A. Jung. *Cartesian Closed Categories of Domains*, volume 66 of *CWI Tracts*. Centrum voor Wiskunde en Informatica, Amsterdam, 1989.
- [Jun90] A. Jung. The classification of continuous domains. In *Proceedings, Fifth Annual IEEE Symposium on Logic in Computer Science*, pages 35–40. IEEE Computer Society Press, 1990.

- [Kel75] J. L. Kelley. *General Topology*, volume 27 of *Graduate Texts in Mathematics*. Springer Verlag, 1975.
- [Kop94] R. Kopperman. The skew fields of topology. Preprint, 1994.
- [KP94] K. Keimel and J. Paseka. A direct proof of the Hofmann-Mislove theorem. *Proceedings of the AMS*, 120:301–303, 1994.
- [KT96] K. Keimel and R. Tix. Valuations and measures. Unpublished manuscript, 1996.
- [Law79] J. D. Lawson. The duality of continuous posets. *Houston Journal of Mathematics*, 5:357–394, 1979.
- [Law88] J. D. Lawson. The versatile continuous order. In M. Main, A. Melton, M. Mislove, and D. Schmidt, editors, *Mathematical Foundations of Programming Language Semantics*, volume 298 of *Lecture Notes in Computer Science*, pages 134–160. Springer Verlag, 1988.
- [Law91] J. D. Lawson. Order and strongly sober compactification. In G. M. Reed, A. W. Roscoe, and R. F. Wachter, editors, *Topology and Category Theory in Computer Science*, pages 179–205. Clarendon Press, 1991.
- [Law97] J. D. Lawson. Spaces of maximal points. *Mathematical Structures in Computer Science*, 7(5):543–555, October 1997.
- [LS86] J. Lambek and P. J. Scott. *Introduction to Higher Order Categorical Logic*, volume 7 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1986.
- [Mac71] S. Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer Verlag, 1971.
- [McL92] C. McLarty. *Elementary Categories, Elementary Toposes*, volume 21 of *Oxford Logic Guides*. Oxford University Press, 1992.
- [MT92] K. Meinke and J. V. Tucker. Universal algebra. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 1. Clarendon Press, 1992.
- [Nac48] L. Nachbin. Sur les espaces uniformes ordonnés. *C.R.Acad.Sci*, MR 9, 455:774–775, 1948. English translation in [Nac65].
- [Nac65] L. Nachbin. *Topology and Order*. Von Nostrand, Princeton, N.J., 1965. Reprinted by Robert E. Kreiger Publishing Co., Huntington, NY, 1967.

- [Plo76] G. D. Plotkin. A powerdomain construction. *SIAM Journal on Computing*, 5:452–487, 1976.
- [Plo81] G. D. Plotkin. Post-graduate lecture notes in advanced domain theory (incorporating the “Pisa Notes”). Dept. of Computer Science, Univ. of Edinburgh, 1981.
- [Ran53] G. N. Raney. A subdirect-union representation for completely distributive complete lattices. *Proceedings of the AMS*, 4:518–522, 1953.
- [RB96] I. Rewitzky and C. Brink. Unification of four versions of program semantics. Research Reports RR 10/96, The University of Cape Town, October 1996.
- [Sch93] A. Schalk. *Algebras for Generalized Power Constructions*. Doctoral thesis, Technische Hochschule Darmstadt, 1993. 174 pp.
- [Sco70] D. S. Scott. Outline of a mathematical theory of computation. In *4th Annual Princeton Conference on Information Sciences and Systems*, pages 169–176, 1970.
- [Sco76] D. S. Scott. Data types as lattices. *SIAM J. Computing*, 5:522–587, 1976.
- [Sco80] D. S. Scott. Relating theories of lambda calculus. In J. R. Hindley and J. P. Seldin, editors, *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus and Formalism*, pages 403–450. Academic Press, 1980.
- [Smy83a] M. B. Smyth. The largest cartesian closed category of domains. *Theoretical Computer Science*, 27:109–119, 1983.
- [Smy83b] M. B. Smyth. Powerdomains and predicate transformers: a topological view. In J. Diaz, editor, *Automata, Languages and Programming*, volume 154 of *Lecture Notes in Computer Science*, pages 662–675. Springer Verlag, 1983.
- [Smy92a] M. B. Smyth. Stable compactification I. *Journal of the London Mathematical Society*, 45:321–340, 1992.
- [Smy92b] M. B. Smyth. Topology. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, vol. 1, pages 641–761. Clarendon Press, 1992.
- [Sto36] M. H. Stone. The theory of representations for Boolean algebras. *Trans. American Math. Soc.*, 40:37–111, 1936.

- [Sün94] Ph. Sünderhauf. *Discrete Approximation of Spaces*. PhD thesis, Technische Hochschule Darmstadt, 1994. 91 pp.
- [SW63] N. Smythe and C. A. Wilkins. Minimal Hausdorff and maximal compact spaces. *Journal of the Australian Mathematical Society*, 3:167–171, 1963.
- [Vic89] S. J. Vickers. *Topology Via Logic*, volume 5 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1989.
- [Zha91] G.-Q. Zhang. *Logic of Domains*. Progress in Theoretical Computer Science. Birkhäuser, 1991.