

EXTENSIONAL AND INTENSIONAL MODELS FOR BOUNDED AND UNBOUNDED NONDETERMINISM

J. LAIRD

Department of Computer Science, University of Bath, UK

ABSTRACT. We give extensional and intensional characterizations of functional programs with nondeterminism: as structure preserving functions between *biorders*, and as nondeterministic *sequential algorithms* on *ordered concrete data structures* which compute them. A fundamental result establishes that these extensional and intensional representations are equivalent, by showing how to construct the unique sequential algorithm which computes a given monotone and stable function, and describing the conditions on sequential algorithms which correspond to continuity with respect to each order.

We illustrate by defining may-testing and must-testing denotational semantics for sequential functional languages with bounded and unbounded choice operators. We prove that these are computationally adequate, despite the non-continuity of the must-testing semantics of unbounded nondeterminism.

In the bounded case, we prove that our continuous models are fully abstract with respect to may-testing and must-testing by identifying a simple universal type, which may also form the basis for models of the untyped λ -calculus. In the unbounded case we observe that our model contains computable functions which are not denoted by terms, by identifying a further “weak continuity” property of the definable elements, and use this to establish that it is not fully abstract.

1. INTRODUCTION

This paper describes denotational models of sequential higher-order functional computation with nondeterminism, which may be given as an explicit choice operation or arise via an under-specified operational semantics or abstract interpretation of a richer language. Interpreting nondeterminism in a sequential setting presents some novel challenges for denotational semantics — for example, accurately representing the branching points within a program at which choices are made. Moreover *unbounded* nondeterminism allows programs to be written which will always return a value but may take an unbounded number of steps to do so (corresponding to the notion of *fairness* [Dji76]). Any sound denotational interpretation of such programs will not be a continuous function, meaning that many standard semantic techniques are not available.

We take an approach which relates extensional and intensional representations of nondeterministic functional programs, by developing the theory of *biorders* (sets with two partial orders) and applying it to nondeterminism, and defining a new notion of *ordered*

Key words and phrases: Nondeterminism, Denotational Semantics, Biorders, Sequential algorithms.

concrete data structure — a game-like structure representing the interaction between a program and its environment at a given type. Our key results are based on an isomorphism between the *monotone stable* functions between the biorders of states of ordered concrete data structures, and the sequential algorithms which compute them. We show that these equivalent representations may be used to give may-and-must testing models of an “observably sequential” functional programming language with nondeterminism, $\Lambda^+(N)$. These models are fully abstract when restricted to bounded choice, as is shown by identifying a simple universal type, of which every other computational type is a retract. However, the model of must testing for unbounded non-determinism contains a class of elements which are not definable as terms, leading to a failure of full abstraction and suggesting that it could be further refined by restricting to “weakly continuous” functions.

1.1. Related Work: Extensional Semantics for Non-determinism. There is a body of research describing the appropriate order-theoretic structures for representing nondeterminism denotationally, and establishing their relationship to notions of testing and equivalence — e.g. [HA80, Plo82]. The principal difficulty which arises when interpreting countable nondeterminism — the non-continuous nature of such a semantics — may be resolved by weakening the continuity properties demanded (e.g. to ω_1 -continuity [AP86]). However, this admits many undefinable functions and leaves fewer principles with which to reason about program behaviour (not least, proving computational adequacy).

In this work we focus on the relationship between sequentiality of (higher-order) functions and non-determinism, axiomatizing the structure required to interpret choice in this setting. More specifically, our approach is based on two orders — an extensional order corresponding to the observable input/output behaviour of functions, and a *stable* order, capturing, in effect, the minimal computation required to produce a given result. This is foreshadowed by Roscoe’s work on the semantics of CSP (in which the stable order is called the “strong order”) [Ros93a, Ros93b] and Berry’s introduction of *biorders* [Ber78], which combine both orders in constructing models of functional (but deterministic) languages. In previous work, the author has observed that stable and continuous functions on biorders with a (extensionally) greatest element are sequential in the sense of Milner and Vuillemin, and used them to construct fully abstract models of sequential languages such as the lazy λ -calculus [Lai03]. However, although these models technically carry information about the way higher-order programs evaluate their arguments, this must be recovered from the graphs or traces of the functions they denote, suggesting that an alternative representation would be valuable for program analysis.

One approach is suggested by the work of Curien, Plotkin and Winskel [CWP97], which describes a fully faithful functor into the category of bidomains from a Cartesian closed category constructed from a model of linear logic based on *bistructures*. This is a more general representation of bidomains than that given by ordered concrete data structures (i.e. the biorders which are finite-branching states of the latter are strictly included in those which correspond to cliques of the former). Thus it includes stable and extensionally continuous functions which are not sequential, such as Gustav’s function [Ber78]. As a continuous semantics, it does not capture programs with unbounded nondeterminism.

Intensional Semantics for Non-determinism. Given the success of intensional techniques — in particular, game semantics — in characterizing deterministic sequential computation through fully abstract models of programming languages (with and without side-effects),

extending these models with non-determinism by relaxing the determinacy constraint on strategies is a natural step. This was taken by Harmer and McCusker [HM98], who gave a game semantics of Idealized Algol with bounded non-determinism, fully abstract with respect to may-and-must testing equivalence. This bears clear parallels with our intensional model, e.g. in the explicit recording of divergent traces. Representing the branching behaviour of functional programs without state is more challenging; witness the difficulty of combining innocence with nondeterminism [Har99], resolved in work by Tsukada and Ong [TO15] based on a sheaved model, as well as in concurrent games models [RW11]. This succeeds in characterizing the definable elements in a model of the λ -calculus with choice, although it does not capture their observational equivalence in a direct or effective manner. Intensional models of unbounded nondeterminism encounter challenges related to non-continuity — as noted in [HM98], representations of strategies as collections of finite sequences of moves are insufficient to capture the distinction between infinite interactions and finite, but unbounded ones, and Levy [Lev08] describes a semantics of recursively defined nondeterministic programs in terms of their infinite traces. Our model similarly contains representations of infinitary interaction, although these are positions which may be reached by (potentially, many different) ordinal chains of events.

The intensional semantics for nondeterminism developed here is based on *concrete data structures*, which were introduced by Kahn and Plotkin [KP93], as part of a definition of sequentiality for (deterministic) higher-order functionals (see further discussion in Section 6.2), although they offer an appealing model of computation in their own right, via the notion of *sequential algorithm* [BC82]. On the one hand, concrete data structures correspond to a *positional* form of games, and sequential algorithms to positional strategies (see e.g. [HS02]). On the other, sequential algorithms may be related to purely extensional models: in the deterministic case, Cartwright, Curien and Felleisen [CCF94] have established that they compute, and are equivalent to “observably sequential” functions; the author has given a more abstract representation of the latter as *bistable* functions on bistable biorders [Lai07a, Lai05].

Berry and Curien’s deterministic sequential algorithms are *strongly sequential* — in each position the data required to compute further is *unique* (each *cell* is filled with a unique *value*). By contrast, we describe a *weakly sequential* representation (cells may be filled with multiple values) — which may be required either because evaluation is explicitly non-deterministic, or because evaluation order is not observable — by abandoning this consistency condition. However, this also requires an ordering on cells and values (corresponding to game positions), to reflect the fact that (for example) any program which may diverge in response to a given argument may still diverge in response to an argument with a wider range of behaviours. This notion of an ordered concrete data structure was first introduced in [Lai09], in which stable and continuous functions were shown to correspond to *finite-branching* sequential algorithms. Extension to unbounded nondeterminism [Lai15] requires a new notion of transfinite (ordinal-indexed) interaction, to distinguish computations which are infinite from those which are finite but unbounded.

1.2. Outline of the Paper. Section 2 recalls and brings together some earlier work by using it to give a biorder semantics of a minimal functional language with non-determinism, Λ^+ . This illustrates some features of sequential nondeterministic higher-order functions, and motivates the rest of the paper which expands and elaborates on this model. Section 3 extends it with infinitary data (natural numbers) and fixed points, allowing the expression of

$$\begin{array}{c}
\overline{\Downarrow \Downarrow \Downarrow} \quad \overline{\lambda x. t \Downarrow \lambda x. s} \\
\frac{s \Downarrow^{must} \Downarrow \quad t \Downarrow^{must} \Downarrow}{s + t \Downarrow^{must} \Downarrow} \quad \frac{s \Downarrow^{may} \Downarrow \quad t \Downarrow^{may} \Downarrow}{s + t \Downarrow^{may} \Downarrow}
\end{array}
\quad
\frac{s \Downarrow \lambda x. s' \quad s'[t/x] \Downarrow C}{s \Downarrow C}$$

Table 1: Big-step rules for May and Must Testing

countable nondeterminism, and shows that these may be interpreted soundly by introducing notions of completeness to biorders. Section 4 introduces *ordered concrete data structures* (ocds) and their biorders of states, and Section 5 introduces non-deterministic sequential algorithms and shows that they compute monotone stable functions between biorders. Section 6 extends these results to show that each such function is computed by a unique sequential algorithm, and thus establish Cartesian closure for our categories of ocds. Section 7 returns to the semantics of nondeterminism, proving full abstraction and universality results for the models of bounded nondeterminism, and describing a further “weak continuity” property of definable elements in the must testing semantics of unbounded nondeterminism, leading to a proof that it is not fully abstract.

2. NON-DETERMINISM AND WEAK SEQUENTIALITY

Let Λ^+ be the simply-typed λ -calculus over a single base type o with two ground type constants — $\Omega : o$ (representing divergence) and $\Downarrow : o$ (representing immediate termination, due to error, deadlock etc.) — and a binary choice operation $+$: $o \rightarrow o \rightarrow o$, which we write infix. Programs (closed terms of type o) may be evaluated by (leftmost outermost) reduction) of $\lambda x. r s$ to $r[s/x]$ and $r + s$ to either r or s . If every such reduction of a program $t : o$ terminates with \Downarrow then we say that t *must* converge ($t \Downarrow^{must}$) and if some reduction of t terminates with \Downarrow then we say that t *may* converge ($t \Downarrow^{may}$).

Alternatively, these may and must convergence predicates may be defined via the “big step” reduction rules in Table 1 (in which the rules in the top row apply to both may and must convergence predicates). They are used as the basis for the following notions of observational approximation on terms of Λ^+ :

May Testing: $s \lesssim^{may} t$ if for all compatible program contexts $C[\cdot] : o$, $C[s] \Downarrow^{may} \Downarrow$ implies $C[t] \Downarrow^{may} \Downarrow$.

Must Testing: $s \lesssim^{must} t$ if for all compatible program contexts $C[\cdot] : o$, $C[s] \Downarrow^{must} \Downarrow$ implies $C[t] \Downarrow^{must} \Downarrow$.

May-and-Must Testing: $s \lesssim_{must}^{may} t$ if $s \lesssim^{may} t$ and $s \lesssim^{must} t$.

We obtain an equivalence by the intersection of each approximation relation with its converse. For example, let \mathbf{bool} be the type $o \rightarrow o \rightarrow o$, which contains as its equivalence classes (in each case) the values $\mathbf{tt} \triangleq \lambda x. \lambda y. x$ and $\mathbf{ff} \triangleq \lambda x. \lambda y. y$, the error and divergence terms $\Downarrow_{\mathbf{bool}} \triangleq \lambda x. \lambda y. \Downarrow$ and $\Omega_{\mathbf{bool}} \triangleq \lambda x. \lambda y. \Omega$ and the non-deterministic choice $\mathbf{tt} + \mathbf{ff} \triangleq \lambda x. \lambda y. (x + y)$. We have $\Omega \lesssim^{may} \mathbf{tt}, \mathbf{ff} \lesssim^{may} \mathbf{tt} + \mathbf{ff} \lesssim^{may} \Downarrow$ and $\Omega \lesssim^{must} \mathbf{tt} + \mathbf{ff} \lesssim^{must} \mathbf{tt}, \mathbf{ff} \lesssim^{must} \Downarrow$.

We make the following observations:

- The notions of approximation and equivalence are extensional, in the sense that closed programs of function type will be observationally equivalent if and only if they are

equivalent when applied to any closed term. So, in particular, there are only finitely many equivalence classes at each type.

- In this finitary setting, may testing and must testing (and thus the corresponding notions of approximation and equivalence) are dual, in the sense that $t \Downarrow^{may}$ if and only if $t[\Omega/\mathcal{U}, \mathcal{U}/\Omega] \Downarrow^{must}$. So one model will suffice for both.
- The operational semantics for must testing is the same as that given for *Unary PCF* in e.g. [Loa98]. Unary PCF is the simply-typed λ -calculus over a ground type containing a single value, and a single binary operator on this type which tests both of its arguments for convergence to that value. The latter is usually regarded as a sequential composition, but since the order of evaluation is not observable, the arguments may be evaluated in either order, or in parallel, or non-deterministically.

By relating non-determinism to Unary PCF we may use some important facts about the latter — including that its observational equivalence is decidable [Loa98], and that it has a universal model in the category of biorders and monotone stable functions [Lai03]. Conversely, our sequential algorithms model may be seen as a solution to the problem of giving a corresponding intensional semantics of Unary PCF.

Unary PCF (and thus Λ^+) is a sequential language in the sense of Milner-Vuillemin: for any term $t : T_1 \rightarrow \dots T_k \rightarrow o$ which is strict (i.e. there is some sequence of arguments on which it diverges) there is a *sequentiality index* $i \in \{1, \dots, k\}$ such that $t s_1 \dots s_k \Downarrow^{must}$ implies $s_i \not\Downarrow^{must} \lambda \vec{x}. \Omega$. In fact, it is *weakly* sequential — there are strict but non-divergent terms with more than one sequentiality index — e.g. $\lambda xy. x + y$. As in this instance, weak sequentiality may arise because evaluation is non-deterministic, or because the order of evaluation is underspecified or unobservable. Other examples of weakly sequential languages include PCF itself (for which the fully abstract model is known to not be effectively presentable [Loa00]), and the call-by-value and lazy λ -calculi [Abr90] — e.g. in the absence of side-effects such as mutable state, a call-by-value application $t s$ may be evaluated by first evaluating t to a value, then s , or vice-versa (or, indeed, non-deterministically or in parallel) — which do have effectively presentable fully abstract models based on biorders [Lai03].

Intensional (game) semantics are typically based on explicitly sequential representations of interaction such as sequences of moves: one of the aims of the current work is to investigate a model of higher-order computation which is intrinsically weakly sequential. The following example illustrates why this is needed to give fully abstract semantics of non-deterministic functional programs.

Example 2.1. Define the following terms $x_0 : \text{bool}, x_1 : \text{bool} \vdash t_{ij} : \text{bool}$

- $t_{0,1} \triangleq \text{If } x_0 \text{ then } x_1 \text{ else } x_1$, (i.e. $\lambda ab. (x_0 ((x_1 a) b)) ((x_1 a) b)$)
- $t_{1,0} \triangleq \text{If } x_1 \text{ then } x_0 \text{ else } x_0$,
- $t_{0,0} \triangleq \text{If } x_0 \text{ then (If } x_1 \text{ then tt else tt) else (If } x_1 \text{ then ff else ff)}$,
- $t_{1,1} \triangleq \text{If } x_1 \text{ then (If } x_0 \text{ then tt else tt) else (If } x_0 \text{ then ff else ff)}$.

Each term evaluates both arguments and returns one of them — t_{ij} evaluates argument i first and returns the value of argument j . Thus they are pairwise (may or must) observationally distinguishable — $t_{0,1}$ and $t_{0,0}$ may be distinguished from $t_{1,0}$ and $t_{1,1}$ by application to Ω and \mathcal{U} , and $t_{1,0}$ and $t_{0,0}$ may be distinguished from $t_{0,1}$ and $t_{1,1}$ by application to tt and ff . However the terms $t_{0,1} + t_{1,0}$ and $t_{0,0} + t_{1,1}$ are may-and-must equivalent — both can test their arguments in either order, and return either one of them. Note that they are

observationally distinguishable using mutable state (e.g. in erratic Idealized Algol), by using an imperative variable to record the order in which the arguments *actually were* tested.

A second (well-known) example illustrates another difficulty of giving an intensional semantics of non-deterministic functions — in addition to capturing the order in which they explore their arguments, it is necessary to represent the *branching points* where different evaluation paths may be chosen.

Example 2.2. The following approximations are strict:

$$\lambda f.(f \text{ tt}) + (f \text{ ff}) \lesssim^{\text{may}} \lambda f.f(\text{tt} + \text{ff}) \text{ and } \lambda f.f(\text{tt} + \text{ff}) \lesssim^{\text{must}} \lambda f.(f \text{ tt}) + (f \text{ ff})$$

Indeed, these terms are separated by application to $\lambda x.\text{If } x \text{ then } \Omega \text{ else } (\text{If } x \text{ then } \mathcal{U} \text{ else } \Omega)$ (may-testing) and $\lambda x.\text{If } x \text{ then } \mathcal{U} \text{ else } (\text{If } x \text{ then } \Omega \text{ else } \mathcal{U})$ (must-testing).

To capture the difference between these functions at the level of the strategies which compute them, it is necessary to record whether different calls by f to its argument may be supplied with different values. For strategies which are innocent (i.e. not stateful) this requires a representation of interaction which captures branching points, such as the sheaved model of [TO15]

2.1. Biorders and Stable Functions. Turning to denotational semantics, we recall, redefine and develop the notion of *biorder* [Ber78, CWP97]. At its most basic a biorder is a set with two partial orders, related as follows:

Definition 2.3. A biorder $(|D|, \sqsubseteq_E, \leq_S)$ is a set with two partial orders (the *extensional* and *stable* orders) such that any non-empty set $X \subseteq |D|$ which is bounded above in the stable order (for which we may write $\uparrow X$) has a greatest lower bound $\bigwedge X$ in that order, which is also a least upper bound $\bigsqcup X$ for X in the extensional order.

Remark 2.4. Note that it is a consequence of this definition that the stable order is contained in the *reverse* of the extensional order, unlike the original definition [Ber78], where it appears as a refinement of that order. This is a presentational choice — it is possible to reframe the following definitions in the original style — which has some advantages in the context of interpreting non-deterministic computation:

- May-and-must duality — as we have noted, the may and must orders (which correspond to the extensional order and its dual) are dual to each other in the finite setting of Λ^+ — one is a meet semilattice, the other is a join semilattice. Our definition allows a single notion of stability, rather than (in the may testing case) using its dual, co-stability.
- Continuity — introducing fixed points breaks the duality between may and must (and opens the space between finite and unbounded nondeterminism in the must case). Our definition of biorder accommodates this by allowing the definition of stably continuous and extensionally monotone functions (must testing) versus extensionally continuous and stable functions (may testing) in the same setting.
- In our intensional model, which essentially extends sequential algorithms with explicit divergences, it is natural to interpret the extensional order as a *containment* between states, but the stable order as a *refinement* which removes divergences.
- We identify a new property of programs with unbounded nondeterminism — that they preserve suprema of extensional chains, but only up to stable approximation — which is most naturally presented as a weak continuity property.

We shall say that a biorder is *extensionally pointed* if it has an extensionally least element, \perp_E , *stably pointed* if it has a stably least element \perp_S (note that the latter element is greatest in the extensional order), and just *pointed* if it has both.

Key constructions on biorders include:

- For any family of biorders $\{D_i \mid i \in I\}$, the (set-theoretic) product $\prod_{i \in I} D_i$ and disjoint sum $\coprod_{i \in I} D_i$ with both orders defined pointwise.
- If D is a biorder, its stable lifting D_\uparrow is given by adding a new element to D which is least with respect to the stable order (and thus greatest with respect to the extensional order).

From these constructions we obtain the empty biorder, the one-element biorder 1 and its stable lifting — the (pointed) biorder Σ , with two elements $\{\perp_S, \perp_E\}$ such that $\perp_E \sqsubseteq_E \perp_S$ and $\perp_S \leq_S \perp_E$. For any set X , we write \tilde{X} for the disjoint sum $\coprod_{x \in X} 1$ (in which the extensional and stable orders are both the discrete order).

To define a category of biorders, we recall Berry's notion of *stable function*.

Definition 2.5. A function $f : D \rightarrow E$ between partial orders is *stable* if it is monotone and for any $x \in D$ and $y \leq f(x)$ there exists $m \leq x$ such that $y \leq f(m)$ and for all $z \leq x$, $y \leq f(z)$ implies $m \leq z$.

The relationship between stability and bounded infima may be summed up as follows:

Lemma 2.6. *If partial orders D and E have bounded infima then $f : D \rightarrow E$ is stable if and only if it preserves all bounded infima.*

Proof. From left-to-right, suppose Z is bounded above by x , then taking $y = \bigwedge_{z \in Z} f(z)$ there exists m such that $y \leq f(m)$ and for all $z \in Z$, $m \leq z$ and thus $\bigwedge_{z \in Z} f(z) = f(\bigwedge Z)$ as required.

From right-to left, given $y \leq x$, take $m = \bigwedge \{a \leq x \mid y \leq f(a)\}$. □

We define a category \mathcal{B} in which objects are biorders, and morphisms from D to E are functions from $|D|$ to $|E|$ which are *monotone stable* — i.e. monotone with respect to the extensional order and stable with respect to the stable order. This concrete category is Cartesian closed — the product of biorders is a cartesian product and the internal hom $[D, E]$ is the biorder consisting of the set of monotone stable functions from D to E with the extensional (Scott) and stable (Berry) orders on functions — i.e.

$$\begin{aligned} f \sqsubseteq_E g &\text{ if for all } x \in D, f(x) \sqsubseteq_E g(x) \\ f \leq_S g &\text{ if for all } x, y \in D, x \leq_S y \implies f(x) \leq_S g(y) \text{ and } f(x) = f(y) \wedge g(x). \end{aligned}$$

The infimum of a stably bounded set of monotone stable functions is given pointwise — i.e. $(\bigwedge F)(x) = \bigwedge_{f \in F} f(x)$, which is clearly stable and monotone, and a \sqsubseteq_E -supremum for F .

It is well known (and an interesting exercise to prove) that stability of $f : D \times E \rightarrow F$ implies stability of its currying $\Lambda(f) : D \rightarrow [E, F]$ and vice-versa. Thus $\mathcal{B}(D \times E, F)$ and $\mathcal{B}(D, [E, F])$ are naturally isomorphic and so \mathcal{B} is Cartesian closed, as is its subcategory of pointed objects. So, in particular, it contains may and must testing denotational semantics of Λ^+ in which the ground type o is interpreted as the biorder Σ , the choice operator as the extensional join operator $\sqcup : \Sigma \times \Sigma \rightarrow \Sigma$ (which is monotone stable), and the constants \mathcal{U} and \mathcal{O} as \perp_S and \perp_E (respectively) in the may testing model and \perp_E and \perp_S in the must testing model.

Proposition 2.7 (Computational Adequacy). *$t \Downarrow^{may} \mathcal{U}$ if and only if $\llbracket t \rrbracket_{may} = \llbracket \mathcal{U} \rrbracket_{may}$ and $t \Downarrow^{must} \mathcal{U}$ if and only if $\llbracket t \rrbracket_{must} = \llbracket \mathcal{U} \rrbracket_{must}$*

Proof is simple based on strong normalization of the λ -calculus and soundness of the reduction rules. We will give a proof of adequacy for a conservative extension of this model (Proposition 3.23).

2.2. Sequentiality and Universality. Sequentiality for higher-order functions may be defined in a variety of ways, which will be discussed further in later sections: via correspondence with an explicitly sequential representation such as games or sequential algorithms (see Section 5) via Kahn-Plotkin sequentiality indices [KP93] (see Section 6), or via universality and full abstraction results for an interpretation of a sequential language (here, for Λ^+ and in Section 7 for a more expressive extension). First, we recall from [Lai03] a simple result establishing that conditionally multiplicative functions on biorders are sequential in the “first-order” sense of Milner and Vuillemin.

Definition 2.8. Let $\{A_i \mid i \in I\}$ be pointed biorders. A function $f : \prod_{i \in I} A_i \rightarrow B_\uparrow$ is *strict* if $f(\perp_S) = \perp_S$ and *i -strict* if $\pi_i(x) = \perp_S$ implies $f(x) = \perp_S$.

Proposition 2.9 (Milner-Vuillemin Sequentiality). *Any strict morphism $f : \prod_{i \in I} A_i \rightarrow B_\uparrow$ is i -strict for some $i \in I$.*

Proof. Define (for $x \in \prod_{i \in I} A_i$ and $y \in A_i$), the insertion $x[y]_i \in \prod_{i \in I} A_i$ by $\pi_j(x[y]_i) = y$ if $i = j$ and $\pi_j(x[y]_i) = \pi_j(x)$, otherwise.

Then $\{\perp_E[\perp_S]_i \mid i \in I\}$ is stably bounded above by \perp_E , and has stable greatest lower bound \perp_S so $f(\bigwedge_{i \in I} \perp_E[\perp_S]_i) = f(\perp_S) = \perp_S$ and so by stability of f , $f(\perp_E[\perp_S]_i) = \perp_S$ for some $i \in I$. Hence $f(x[\perp_S]_i) = \perp_S$ for all $x \in \prod_{i \in I} A_i$, as $\perp_E[\perp_S]_i \sqsubseteq_E x[\perp_S]_i$. \square

Note that a function may be i -strict for several values of i (for example, the function from $\Sigma \times \Sigma$ to Σ sending x, y to $x \sqcup y$). In other words our model contains weakly sequential functions, as expected.

It was shown in [Lai03] that any (Milner-Vuillemin) sequential and order-extensional model of Unary PCF in a CCC which interprets the ground type as the two-point order, is universal (i.e. every element is denoted by a term) and thus fully abstract. We sketch a simplified version of the proof here, which is the basis for a proof of universality for an infinitary extension of Λ^+ in Section 7.

Proposition 2.10. *The biorder semantics of Λ^+ are universal.*

Proof. By duality, we may give the proof for the must testing model only. For any types S and T , and $n \in \mathbb{N}$, write $S \trianglelefteq T^n$ if there is a *definable retraction* from S to T^n – that is, there are terms $\mathbf{inj}_1 : S \rightarrow T, \dots, \mathbf{inj}_n : S \rightarrow T$ and $\mathbf{proj} : T \rightarrow \dots \rightarrow T \rightarrow S$ such that $\llbracket \lambda x. \mathbf{proj} \mathbf{inj}_1(x) \dots \mathbf{inj}_n(x) \rrbracket = \llbracket \lambda x. x \rrbracket$.

The following facts are evident:

- If $S \trianglelefteq T^n$ then universality at type T implies universality at type S .
- Universality holds at type $o \rightarrow o$, since the only elements of $[\Sigma, \Sigma]$ are the functions denoted by $\lambda x. \Omega$, $\lambda x. x$ and $\lambda x. \mathcal{U}$.

So it suffices to prove for all types T , that $T \trianglelefteq (o \rightarrow o)^n$ for some n , by induction on T , for which the key step is to show that $(o \rightarrow o)_1 \rightarrow \dots \rightarrow (o \rightarrow o)_n \rightarrow (o \rightarrow o) \trianglelefteq (o \rightarrow o)^{2^{2n+1}}$, which follows from these two facts:

- $(o \rightarrow o) \rightarrow o \trianglelefteq (o \rightarrow o)^2$ — by the retraction $\mathbf{inj}_1 \triangleq \lambda f. \lambda x. f \lambda y. x$, $\mathbf{inj}_2 \triangleq \lambda f. \lambda x. f \lambda y. y$ and $\mathbf{proj} \triangleq \lambda x. \lambda y. \lambda z. x (z (y \Omega))$. For each element $e \in \llbracket (o \rightarrow o) \rightarrow o \rrbracket$ (i.e. $e \in \{\lambda y. \Omega, \lambda y. \mathcal{U}, \lambda y. y \Omega, \lambda y. y \mathcal{U}\}$) we have $(\mathbf{proj} \mathbf{inj}_1(e)) \mathbf{inj}_2(e) = e$.

- $o \rightarrow o \rightarrow o \leq (o \rightarrow o)^2$ by the retraction $\text{inj}_1 \triangleq \lambda f. \lambda x. (f x) \mathcal{U}$, $\text{inj}_2 \triangleq \lambda f. \lambda x. (f \mathcal{U}) x$ and $\text{proj} \triangleq \lambda f. \lambda g. \lambda x. \lambda y. f(x) + g(y)$. For any element $e \in \llbracket o \rightarrow o \rightarrow o \rrbracket$, $(\text{proj } \text{inj}_1(e)) \text{inj}_2(e) = \lambda x. \lambda y. ((e x) \top_S) \wedge ((e \top_S) y) = \lambda x. \lambda y. (e x) y = e$ by stability. \square

Proposition 2.11 (Full abstraction). $t \lesssim^{\text{may}} s$ iff $\llbracket t \rrbracket_{\text{may}} \sqsubseteq_E \llbracket s \rrbracket_{\text{may}}$ and $t \lesssim^{\text{must}} s$ iff $\llbracket s \rrbracket_{\text{must}} \sqsubseteq_E \llbracket t \rrbracket_{\text{must}}$.

Proof. This follows from adequacy, universality and extensionality of the model following standard arguments as in [Plo77]. \square

Thus, we have given a fully abstract semantics of a simple, sequential non-deterministic functional language using biorders. In the remainder of the paper, we will extend it to a fully expressive model of computation, and give an intensional characterization of its types and programs.

3. $\Lambda^+(N)$ AND ITS SEMANTICS

Λ^+ is too limited to describe generalized higher-order non-deterministic computation because it lacks infinitary datatypes. In particular, this means that there is no real distinction between may and must testing (witness the formal duality between them) and there is no possibility of describing unbounded nondeterminism and comparing it with the binary choice operation. So we will define a new prototypical language by extending Λ^+ with a type N of natural number *values* — i.e. terms of type N do not have computational effects (nondeterminism or non-termination). Maintaining this distinction between computation and value types (as in *call-by-push-value* [Lev04]) simplifies programming with non-determinism and allows construction of fully-abstract semantics of richer languages by CPS interpretation.

Types of $\Lambda^+(N)$ are either *pointed* (computation) types P or possibly unpointed (value) types V , given by the grammars:

$$\begin{array}{l} V ::= N \mid P \\ P ::= o \mid V \rightarrow P \end{array}$$

Terms are formed by extending Λ^+ with the following constants and operations:

- Expressions of type N , given by the grammar $e ::= 0 \mid x : N \mid \text{succ}(e)$.
- Equality testing: given expressions $s, t : N$, $\text{Eq}(s, t) : \text{bool}$.
- Fixed point combinators: $\mathbf{Y}_P : (P \rightarrow P) \rightarrow P$ for each pointed type P .
- Unbounded Choice: $?_N : (N \rightarrow o) \rightarrow o$, which nondeterministically chooses a numeral to which its argument is applied.

$\Lambda^+(N)$ may be regarded as a target language for *continuation-passing-style* translation of languages with bounded and unbounded nondeterminism, with o as the return or answer type. Most simply, by defining the type nat of natural number computations to be the type $(N \rightarrow o) \rightarrow o$, we may express the language SPCF (observably sequential PCF [CF92]) extended with bounded (binary) or unbounded (natural number) choice. In particular, note that we may define $\text{pred} : N \rightarrow \text{nat} \triangleq \lambda u. \lambda f. (\mathbf{Y} \lambda g. \lambda v. (\text{Eq}(u, \text{succ}(v)) (f v)) (g \text{succ}(v))) 0$

SPCF is itself an extension of PCF with a simple non-local control operator (e.g. we may define $\text{callcc} : ((\text{nat} \rightarrow \text{nat}) \rightarrow \text{nat}) \rightarrow \text{nat} \triangleq \lambda f. \lambda k. (f \lambda xy. x k) k$) and some error terms which immediately abort computation (in this case, a single such term, \mathcal{U}). SPCF provides a complete syntactic representation of the (strongly) observably sequential functions (and the corresponding deterministic sequential algorithms) [CCF94, Lai07a] as demonstrated

$$\begin{array}{c}
\overline{\text{Eq}(\mathbf{n}, \mathbf{n}) \Downarrow \lambda x. \lambda y. x} \quad \overline{\text{Eq}(\mathbf{m}, \mathbf{n}) \Downarrow \lambda x. \lambda y. y} m \neq n \\
\overline{\mathbf{Y} \Downarrow \lambda f. f} (\mathbf{Y} f) \quad \frac{\exists n. t \mathbf{n} \Downarrow^{\text{may}} \mathcal{U}}{?_N t \Downarrow^{\text{may}} \mathcal{U}} \quad \frac{\forall n. t \mathbf{n} \Downarrow^{\text{must}} \mathcal{U}}{?_N t \Downarrow^{\text{must}} \mathcal{U}}
\end{array}$$

Table 2: Operational Semantics for $\Lambda^+(N)$

by full abstraction and universality results — this paper may be seen as an extension of these results to weakly observably sequential functions and non-deterministic sequential algorithms.

3.1. Operational Semantics. We give may and must testing operational semantics for programs (closed terms of type o) by extending the rules for Λ^+ (Table 1) with the additional rules in Table 2. The definitions of may, must and may-and-must approximation and equivalence extend directly from Λ^+ . Also as in Λ^+ , the presence of the error element makes the language extensional — any terms of function type are observationally equivalent if (and only if) applying to the same argument returns equivalent results. (This is essentially Milner’s Context Lemma [Mil77] and may be proved using the same techniques.) In particular, we will use the result for must equivalence in Section 7.

Lemma 3.1 (Function Extensionality). *For any closed terms $s, t : S \rightarrow T$, if $s r \simeq_{\text{must}} t r$ for all $r : S$ then $s \simeq_{\text{must}} t$.*

3.2. Unbounded Nondeterminism. $\Lambda^+(N)$ gives us a setting in which we may express programs with *unbounded* non-determinism — they may evaluate to (countably) infinitely many different values without diverging. Evidently, we may express bounded choice up to may-and-must equivalence using unbounded choice — e.g. $s + t \simeq_{\text{must}}^{\text{may}} ?_N (\lambda k. (\text{Eq}(k, 0) s) t)$. Conversely:

Proposition 3.2. *$?_N$ is macro-expressible up to may equivalence in bounded $\Lambda^+(N)$.*

Proof. The terms $?_N$ and $\lambda k. (\mathbf{Y}_{\text{nat}} \lambda f. \lambda u. (k u) + (f (\text{suc } x))) 0$ are may-equivalent. (They denote the same elements in our fully abstract model of may testing.) \square

Note that this equivalence fails with respect to must equivalence (evaluation may always take the right hand branch of the binary choice and so diverge), as will any such attempt to define countable choice using bounded choice (a term which reduces to infinitely many different values will have an infinite reduction path by König’s lemma and may therefore diverge). In other words:

Proposition 3.3. *$?_N$ is not macro-expressible up to must equivalence using binary choice.*

Formal proof of this claim is straightforward using our fully abstract model (Section 7).

Remark 3.4 (Fairness). One reason for studying unbounded nondeterminism is its close connection with the notion of *fairness*, according to which an event will eventually occur, even if it is not possible to bound the number of steps taken before this happens. For example, suppose we wish to implement a type `coin` with constants `new_coin` : $(\text{coin} \rightarrow o) \rightarrow o$ (which supplies a coin to its argument) and `toss` : $\text{coin} \rightarrow (\text{coin} \rightarrow o) \rightarrow (\text{coin} \rightarrow o) \rightarrow o$

which tosses its first argument and chooses its second or third argument depending on whether the result is heads or tails, and passes the tossed coin to it.

A necessary condition for an implementation of `coin` to be *fair* is that tossing a new coin will eventually come up heads, but may require an unbounded number of tosses to do so. In other words, a program which repeatedly tosses a coin and increments a counter until it gets a head, and then returns the contents of the counter may return any number value (but not diverge) — i.e.

$$\lambda k.((\mathbf{Y}_{N \rightarrow \text{coin} \rightarrow o} \lambda f. \lambda u : N. \lambda x : \text{coin}.\text{toss } x) \lambda y.(k u) (f \text{ suc}(u))) 0) \text{ new_coin}$$

is must equivalent to $?_N$. So a fair coin may be used to express unbounded choice and is therefore not definable in bounded $\Lambda^+(N)$ by Proposition 3.3. Conversely, in unbounded $\Lambda^+(N)$ we may implement a coin which passes the fairness test, e.g. by defining the macros $\text{coin} \triangleq N$, $\text{new_coin} \triangleq ?_N$ and $\text{toss} \triangleq \lambda u. \lambda f. \lambda g. (\text{Eq}(u, 0) (?_N f)) (\text{pred}(u) g)$.

Remark 3.5 (Abstract Interpretation). Another setting in which unbounded nondeterminism may arise is in the “abstract interpretation” of a domain with infinitely many states as a finite domain — the “concretization” of such an abstract interpretation may then represent a single abstract state as an unbounded choice of concrete states. Observe that in $\Lambda^+(N)$, countable nondeterminism may be used to define an embedding-projection pair from the type $N \rightarrow o$ to o — the terms $e \triangleq \lambda x. \lambda u. x : o \rightarrow N \rightarrow o$ and $p \triangleq ?_N : (N \rightarrow o) \rightarrow o$ satisfy $\lambda x. p(e x) \simeq_{\text{must}} \lambda x. x$ and $\lambda y. e(p y) \lesssim_{\text{must}} \lambda y. y$. In other words, there is a Galois insertion between $N \rightarrow o$ and o , giving an *abstract interpretation* of the former, infinitary type inside the latter, finitary one, which we may use to define, inductively, a Galois connection between each pointed type of $\Lambda^+(N)$ and the corresponding type of Λ^+ obtained by erasing all instances of N .

3.3. Continuity. A challenge encountered in defining and reasoning about an adequate denotational semantics of must testing for programs with unbounded non-determinism is that it cannot be continuous with respect to the must ordering (on must equivalence classes of terms).

For instance, define $t_n : N \rightarrow o \triangleq \mathbf{Y}_{N \rightarrow o} \lambda f. \lambda u. (\text{Eq}(n, u) \mathcal{U}) (f \text{ suc}(u))$, so that $t_n \mathbf{i} \Downarrow_{\text{must}}$ iff $i \leq n$. It is easy to see (and prove using our fully abstract model) that $t_n \lesssim_{\text{must}} t_{n+1}$ for each $n \in \mathbb{N}$, and that the \lesssim_{must} -least upper bound for this chain of terms (up to must equivalence) is the term $\lambda f. \mathcal{U}$. However, whereas $?_N t_n$ may diverge for all n , $?_N (\lambda f. \mathcal{U})$ must converge — i.e. $?_N$ cannot denote a function which is continuous with respect to the must ordering.

Another example shows that function application is not continuous in the function component either — or in other words, that least upper bounds of chains of functions are not given pointwise, in general. For instance, let $s_k : (N \rightarrow o) \rightarrow o \triangleq \lambda f. ?_N \lambda x. f \text{ suc}^k(x)$, which nondeterministically supplies to its argument a numeral greater than or equal to k . It is evident that $s_k \lesssim_{\text{must}} s_{k+1}$ for each $k \in \mathbb{N}$, and that the \lesssim_{must} -least upper bound for this chain of terms (up to must equivalence) is the term $\lambda f. \mathcal{U}$. However, $s_k \lambda u. \Omega$ diverges for all k , but $(\lambda f. \mathcal{U}) \lambda u. \Omega$ must converge.

This suggests that we cannot define *well-behaved* fixed points for functions by taking least fixed points of chains of approximants with respect to the must approximation ordering. This problem will be resolved semantically by defining least fixed points with respect to the stable order rather than the extensional (observational) order. Application is stably-continuous

with respect to chains of functions, although not their arguments. (The t_i defined above denote a stable chain in our model, but the s_k do not.)

Failure of continuity also entails that there are must inequivalent terms which cannot be distinguished by finitary tests (i.e. by applying them to arguments which are compact in the must ordering). We give an example at the end of this section, after developing the necessary interpretation of fixed points.

3.4. Complete Biorders. To extend our model of Λ^+ to $\Lambda^+(N)$ we require biorders which are directed complete with respect to at least one of their orders, in order to define least fixed points with respect to that order. We now define these notions of complete biorder.

Definition 3.6. An I -indexed family x of elements of a biorder D is \sqsubseteq_E -directed if for all $i, j \in I$ there exists $k \in I$ such that $x_i, x_j \sqsubseteq_E x_k$. D is *extensionally complete* if the extensional order is directed complete — i.e. every such family has a least upper bound, and for any \sqsubseteq_E -directed I -indexed families x and y , $x_i \sqsubseteq_E y_i$ for all $i \in I$ implies $\bigsqcup_{i \in I} x_i \leq_S \bigsqcup_{i \in I} y_i$.

It is not hard to show that the product, disjoint sum, lifting and function-space preserve extensional completeness.

The requirements for a biorder to be *stably complete* are slightly more involved. They include a “bounded distributivity condition” requiring certain bounded meets to distribute over directed joins, which will be used to define CCCs of stably complete biorders.

Definition 3.7. A *stable upper bound* for a set X of stably directed, I -indexed families over a biorder D is an I -indexed, stably directed family y such that for all $x \in X$ and $i \in I$, $x_i \leq_S y_i$, and for all $i, j \in I$, $x_i \wedge y_j = x_j \wedge y_i$.

D is *stably complete* if the stable order is directed complete and for any stably bounded set X of stably directed I -indexed families of elements, $\bigwedge_{x \in X} \bigvee_{i \in I} x_i = \bigvee_{i \in I} \bigwedge_{x \in X} x_i$.

A biorder is *bicomplete* if it is both extensionally and stably complete.

It is easy to see that the coproduct and product of families of stably complete domains are stably complete. To show that lifting preserves stable completeness (so in particular Σ is stably complete) is a little harder. Note that the above distributivity property does not hold in general for unbounded sets of stably directed families over Σ (i.e. it is not a completely distributive lattice) — take, for example the set of families $\{\{n_i \mid i \in \omega\} \mid n \in \mathbb{N}\}$ with $n_i = \top_S$ if $i > n$ and $n_i = \perp_S$ otherwise, which satisfies $\bigvee_{i \in \omega} \bigwedge_{n \in \mathbb{N}} n_i = \perp_S$ and $\bigwedge_{n \in \mathbb{N}} \bigvee_{i \in \omega} n_i = \top_S$.

Lemma 3.8. *If D is stably complete then D_\uparrow is stably complete.*

Proof. We check that the bounded distributivity property holds. Suppose the set X of I -indexed directed families is bounded above by $\{y_i \mid i \in I\}$. By definition, $\bigvee_{i \in I} \bigwedge_{x \in X} x_i \leq_S \bigwedge_{x \in X} \bigvee_{i \in I} x_i$. If $\bigvee_{i \in I} \bigwedge_{x \in X} x_i \neq \perp_S$ then defining $J = \{i \in I \mid \bigwedge_{x \in X} x_i \neq \perp_S\}$, we have a set of J -indexed directed sets of elements of D , bounded above by $\{y_j \mid j \in J\}$ and so $\bigvee_{i \in I} \bigwedge_{x \in X} x_i = \bigvee_{j \in J} \bigwedge_{x \in X} x_j = \bigwedge_{x \in X} \bigvee_{j \in J} x_j = \bigwedge_{x \in X} \bigvee_{i \in I} x_i$ by stable completeness of D .

So we need to consider the case in which (a) $\bigvee_{i \in I} \bigwedge_{x \in X} x_i = \perp_S$ and (b) $\bigwedge_{x \in X} \bigvee_{i \in I} x_i \neq \perp_S$. We claim that this implies $x_i = \perp_S$ for all $i \in I$ and $x \in X$ and so in fact $\bigwedge_{x \in X} \bigvee_{i \in I} x_i = \perp_S$ as required. To prove the claim, choose any $i \in I$. There exists $z \in X$ such that $z_i = \perp_S$ by (a) and $j \in I$ such that $z_j \neq \perp_S$ by (b). Because $z_i \wedge y_j = y_i \wedge z_j$, we have $y_i = \perp_S$, and hence $x_i = \perp_S$ for all $x \in X$ as required. \square

Lemma 3.9. *If E is stably complete then $[D, E]$ is stably complete.*

Proof. Joins of stably directed sets are defined pointwise (i.e. function application is stably continuous with respect to functions). To show that these are stable, we need to show that they commute appropriately with bounded meets using the bounded distributivity property. Suppose F is a stably directed set of monotone stable functions from D to E , and X is a set of elements stably bounded above by y . Then the set $\{\{f(x) \mid f \in F\} \mid x \in X\}$ of stably directed F -indexed sets is stably bounded above by $\{f(y) \mid f \in F\}$ — for any x , we have $f(x) \leq_S f(y)$ for all f and for any $g \in F$ there exists $h \in F$ with $f, g \leq_S h$ and so $f(x) \wedge g(y) = f(y) \wedge h(x) \wedge g(y) = f(y) \wedge g(x)$. Thus, by the bounded distributivity condition: $(\bigvee F)(\bigwedge X) = \bigvee_{f \in F} \bigwedge_{x \in X} f(x) = \bigwedge_{x \in X} \bigvee_{f \in F} f(x) = \bigwedge_{x \in X} (\bigvee F)(x)$. \square

So, in particular, the full subcategory of \mathcal{B} consisting of bicomplete biorders is Cartesian closed. We define the following subcategories of \mathcal{B} in which morphisms are continuous:

- $\mathcal{B}_{\mathcal{E}}$ — objects are extensionally complete biorders and morphisms from D to E are functions from D to E which are continuously stable — i.e. they are monotone stable and preserve suprema of extensionally directed sets.
- $\mathcal{B}_{\mathcal{S}}$ — objects are stably complete biorders and morphisms from D to E are functions from D to E which are stably continuous — i.e. they are monotone stable and preserve suprema of stably directed sets.

Each of these categories is Cartesian closed, with products given pointwise, and internal homs given by the relevant biorders of functions — i.e. $[D, E]_{\mathcal{B}_{\mathcal{E}}}$ is the restriction of $[D, E]$ to continuously stable functions and $[D, E]_{\mathcal{B}_{\mathcal{S}}}$ is the restriction of $[D, E]$ to stably continuous functions. We need to check that in each case this defines an object of our category — i.e. is (respectively) extensionally complete or stably complete. We give the second (more interesting) case.

Lemma 3.10. *If D and E are stably complete then $[D, E]_{\mathcal{B}_{\mathcal{S}}}$ is stably complete.*

Proof. The proof follows that of Lemma 3.9. We also need to show that if F is a set of continuously stable functions, stably bounded above by a (continuous) function g , then $\bigwedge F$ is continuous: if X is a stably directed set then $\{\{f(x) \mid x \in X\} \mid f \in F\}$ is stably bounded above by the stably directed set $\{g(x) \mid x \in X\}$ since for all $x, x' \in X$ there exists $x'' \in X$ with $x, x' \leq_S x''$ and thus $f(x) \wedge g(x') = f(x'') \wedge g(x) \wedge g(x') = f(x'') \wedge g(x') \wedge g(x) = f(x') \wedge g(x)$. Hence $(\bigwedge F)(\bigvee X) = \bigwedge_{f \in F} f(\bigvee X) = \bigwedge_{f \in F} \bigvee_{x \in X} f(x) = \bigvee_{x \in X} \bigwedge_{f \in F} f(x) = \bigvee_{x \in X} (\bigwedge F)(x)$. \square

Finally, we need to check that the isomorphism of biorders $\Lambda : [D \times E, F] \rightarrow [D, [E, F]]$ restricts to each of our categories — i.e. $\Lambda(f)$ is respectively continuously stable or stably continuous if and only if f is continuously stable/stably continuous. This is straightforward, based on these properties for the Scott and Berry orders. Thus we have shown:

Proposition 3.11. *The categories $\mathcal{B}_{\mathcal{E}}$ and $\mathcal{B}_{\mathcal{S}}$ are Cartesian closed.*

As an example we show that for any set X , the function space $[\Sigma^X, \Sigma] \cong [[\tilde{X}, \Sigma], \Sigma]$ corresponds to a powerdomain biorder on X , justifying the identification of the $\Lambda^+(N)$ type $(N \rightarrow o) \rightarrow o$ with the type `nat` of natural number computations.

Definition 3.12. Given a set X , let $\mathcal{P}(X)$ be the biorder consisting of subsets of X , with the extensional order being inclusion and the stable order discrete.

There is an evident strict morphism $\phi : \mathcal{P}(X)_\uparrow \rightarrow [\Sigma^X, \Sigma]$ sending $Y \subseteq X$ to $\bigsqcup_{y \in Y} \pi_y$. To show that this is an isomorphism by giving an inverse, we define:

Definition 3.13. Given monotone stable $f : \prod_{i \in I} A_i \rightarrow \Sigma$ let $\text{si}(f)$ be the set of $i \in I$ such that f is i -strict (the sequentiality indices of f). By Proposition 2.9 this is non-empty if and only if f is strict.

Lemma 3.14. *If $\pi_i(x) = \perp_E$ for all $i \in \text{si}(f)$ then $f(x) = \perp_E$.*

Proof. If f is not strict then $f(x) \sqsubseteq_E f(\perp_S) = \top_E$. Otherwise, noting that $i \notin \text{si}(f)$ implies $f(\perp_E[\perp_S]_i) = \perp_E$, we have $f(x) \sqsubseteq_E f(\bigwedge_{i \notin \text{si}(f)} \perp_E[\perp_S]_i) = \bigsqcup_{i \notin \text{si}(f)} f(\perp_E[\perp_S]_i) = \perp_E$ by stability of f . \square

Proposition 3.15. *The biorders $\mathcal{P}(X)_\uparrow$ and $[\Sigma^X, \Sigma]$ are isomorphic in \mathcal{B} .*

Proof. By Lemma 3.14, $f(x) = \perp_S$ if and only if $\pi_i(x) \neq \perp_E$ for some $i \in \text{si}(f)$, and thus we may define an inverse to $\phi : \mathcal{P}(X)_\uparrow \rightarrow [\Sigma^X, \Sigma]$ by $\phi^{-1}(f) = \perp_S$, if $f = \perp_S$ and $\phi^{-1}(f) = \text{si}(f)$, otherwise. \square

ϕ is continuously stable (so $\mathcal{P}(X)_\uparrow \cong [\Sigma^X, \Sigma]_{\mathcal{B}_\varepsilon}$) but not stably continuous (cf the non-continuity of unbounded choice with respect to must testing).

Lemma 3.16. *If f is stably continuous then $\text{si}(f)$ is finite.*

Proof. Given $J \subseteq I$ define $a_J \in \prod_{i \in I} A_i$ by $\pi_i(a_J) = \perp_E$ if $i \in J$ and $\pi_i(a_J) = \perp_S$ otherwise. Then $\bigvee_{J \subseteq_{\text{fin}} I} a_J = \perp_E$ and so if f is stably continuous, $f(a_J) = \perp_E$ for some finite $J \subseteq I$. Then for all $i \in \text{si}(f)$, i -strictness of f entails that $\pi_i(a_J) = \perp_E$. Thus $\text{si}(f) \subseteq J$ and thus $\text{si}(f)$ is finite. \square

So defining $\mathcal{P}_{\text{fin}}(X)$ to be the restriction of $\mathcal{P}(X)$ to finite subsets of X , we have: $\mathcal{P}_{\text{fin}}(X)_\uparrow \cong [\Sigma^X, \Sigma]_{\mathcal{B}_S}$.

3.5. Denotational Semantics of $\Lambda^+(N)$. To interpret recursive higher-order programs, we require fixed points of endomorphisms $f : D \rightarrow D$, where D is extensionally or stably complete and pointed. These may be obtained in standard fashion, as (respectively, extensional or stable) suprema of the chain of approximants $\perp, f(\perp), \dots, f^i(\perp), \dots$. In the categories \mathcal{B}_ε and \mathcal{B}_S , which are cpo-enriched, the fixed point is reached at f^ω , whereas in the category of biorders and monotone stable functions it may be necessary to continue the chain beyond ω , but it will still reach a stationary point [AP86].

Proposition 3.17. *If D is stably (resp. extensionally) complete and pointed then every monotone stable function $f : D \rightarrow D$ has \leq_S -least (resp. \sqsubseteq_E -least) fixed points.*

Proof. Define the chain of stable approximants $f^\lambda \in D$ for each ordinal λ by:

- $f^0 = \perp_S$,
- $f^{\kappa+1} = f(f^\kappa)$,
- $f^\lambda = \bigvee_{\kappa < \lambda} f^\kappa$ if $\lambda = \bigcup_{\kappa < \lambda} \kappa$.

By Hartog's Lemma this has a stationary point, which is a \leq_S -least fixed point for f — evidently if f is λ -continuous (preserves suprema of λ -chains) then this stationary point is f^λ . \square

We now have the basis for may testing and must testing semantics of $\Lambda^+(N)$ in our categories of biorders — more precisely:

- May testing and must testing semantics of unbounded $\Lambda^+(N)$ in the category of bicomplete biorders and monotone stable functions. (As we have noted, unbounded choice is expressible up to may-testing in bounded $\Lambda^+(N)$, so the bounded/unbounded distinction is moot in this case and we focus on must testing for unbounded nondeterminism).
- A must testing semantics of bounded $\Lambda^+(N)$ in the category of stably complete and stably continuous functions.
- A may testing semantics of bounded/unbounded $\Lambda^+(N)$ in the category of extensionally complete and continuously stable functions.

(Pointed) types are interpreted as (pointed) biorders : o denotes the biorder Σ , the type N denotes the discrete biorder of natural numbers $\tilde{\mathbb{N}}$ and $S \rightarrow T$ the relevant function-space. Terms $x_1 : S_1, \dots, x_n : S_n \vdash t : T$ denote morphisms from $\llbracket S_1 \rrbracket \times \dots \times \llbracket S_n \rrbracket$ to $\llbracket T \rrbracket$, derived from the Cartesian closed structure of each category, together with:

- The extensional join operation: $s + t$ denotes the bounded (binary) join $\llbracket s \rrbracket \sqcup \llbracket t \rrbracket$ and $?_N$ denotes the infinite join $\bigsqcup_{n \in \mathbb{N}} \lambda f. f(n)$.
- Error: In the must testing models, \mathcal{U} denotes the extensionally least element \perp_E . In the may testing model, it denotes the stably least element (greatest extensional element) \perp_S .
- Fixed points: $\mathbf{Y}_P : (P \rightarrow P) \rightarrow P$ denotes a least fixed point of the function $Y \triangleq F \vdash \lambda f. f(Ff)$. In the must testing models it denotes the stably least fixed point and in the may testing models it denotes the extensionally least fixed point.

3.6. Computational Adequacy. Soundness of the reduction rules with respect to the denotational semantics follows from the structure identified in our categories, yielding:

Proposition 3.18. *$t \Downarrow^{\text{may}} C$ implies $\llbracket t \rrbracket_{\text{may}} = \llbracket C \rrbracket_{\text{may}}$ and $t \Downarrow^{\text{must}} C$ implies $\llbracket t \rrbracket_{\text{must}} = \llbracket C \rrbracket_{\text{must}}$.*

We establish the converse for terms of type o (computational adequacy) by defining “approximation relations” in the style of Plotkin [Plo85]. The models in which this proof is not completely standard are those which are not ω -cpo-enriched — in particular, the case of unbounded nondeterminism with must testing, which we give here. Plotkin’s proof still adapts to this case because fixed point combinators denote least upper bounds of stable chains, which are defined pointwise at function type.

For each type T we define a relation \triangleleft_T between elements of $\llbracket T \rrbracket$ and closed terms of type T :

- $e \triangleleft_N t$ if $e = \llbracket t \rrbracket$.
- $e \triangleleft_o t$ if $e = \llbracket \mathcal{U} \rrbracket$ implies $t \Downarrow^{\text{must}} \mathcal{U}$.
- $f \triangleleft_{S \rightarrow P} t$ if $e \triangleleft_S s$ implies $f(e) \triangleleft_P t s$.

We want to show that every program is approximated by its denotation, for which we use the following two lemmas, proved by straightforward induction on type structure.

Lemma 3.19. *If $t \Downarrow^{\text{must}} C$ implies $s \Downarrow^{\text{must}} C$ then $e \triangleleft_T t$ implies $e \triangleleft_T s$.*

Lemma 3.20. *If $\langle e_\alpha \mid \alpha < \lambda \rangle$ is a stable chain of functions such that $e_\alpha \triangleleft_T t$ for all $\alpha < \lambda$ then $\bigvee_{\alpha < \lambda} e_\alpha \triangleleft_T t$.*

The key lemma for our proof establishes that the fixed point combinator is approximated by its denotation.

Lemma 3.21. *For any pointed type P , $\llbracket \mathbf{Y}_P \rrbracket \triangleleft_{(P \rightarrow P) \rightarrow P} \mathbf{Y}_P$.*

Proof. We show by induction on λ that $Y^\lambda \triangleleft_{(P \rightarrow P) \rightarrow P} \mathbf{Y}$ for each approximant Y^λ to $\llbracket \mathbf{Y} \rrbracket$ — i.e. if, $g \triangleleft_{P \rightarrow P} t$ then $(Y^\lambda g) \triangleleft_P \mathbf{Y} t$.

- For $\lambda = 0$, we have $(Y^\lambda g) = \perp_S \triangleleft_P \mathbf{Y} t$.
- If $\lambda = \kappa + 1$ then by hypothesis $Y^\kappa \triangleleft \mathbf{Y}$, and so $Y^\lambda g = g(Y^\kappa g) \triangleleft_P t(\mathbf{Y} t)$. Since $\mathbf{Y} t \Downarrow^{must} C$ implies $t(\mathbf{Y} t) \Downarrow^{must} C$, $Y^\lambda g \triangleleft_P \mathbf{Y} t$ by Lemma 3.19.
- For $\lambda = \bigcup_{\kappa < \lambda} \kappa$, we have $Y^\lambda = \bigvee_{\kappa < \lambda} Y^\kappa \triangleleft_P \llbracket \mathbf{Y} \rrbracket$ by Lemma 3.20.

So $\llbracket \mathbf{Y} \rrbracket \triangleleft_{(P \rightarrow P) \rightarrow P} \mathbf{Y}$ as required. \square

The rest of the proof is now standard — the approximation relation is extended to open terms by defining $f : \llbracket \Gamma \rrbracket \rightarrow \llbracket T \rrbracket \triangleleft_{\Gamma, T} \Gamma \vdash t : T$ if $\Gamma = x_1 : S_1, \dots, x_n : S_n$ and $e_1 \triangleleft_{S_1} s_1, \dots, e_n \triangleleft_{S_n} s_n$ implies $f(e_1, \dots, e_n) \triangleleft_T t[s_1/x_1, \dots, s_n/x_n]$.

Proposition 3.22. *If $\Gamma \vdash t : T$ then $\llbracket \Gamma \vdash t : T \rrbracket \triangleleft_{\Gamma, T} t$*

Proof. By structural induction on t .

- If t is a value or application then $\llbracket t \rrbracket \triangleleft_{\Gamma, N} t$ by definition of the approximation relation
- If $t = \lambda x : S. t'$ and $e \triangleleft_S s$, then $\llbracket t \rrbracket \triangleleft t'[s/x]$ by hypothesis, and $t'[s/x] \Downarrow^{must} C$ implies $\lambda x. t \Downarrow^{must} C$ and so $\llbracket t \rrbracket(e) \triangleleft t s$ by Lemma 3.19 above.
- If $t = ?_N$, then for any $e \triangleleft_{N \rightarrow o} s$, either $\llbracket s \rrbracket(e) = \perp_S$ or $e(i) = \top_S$ for all $i \in \omega$ and so by hypothesis $s i \Downarrow^{must} \mathcal{U}$ for all i and hence $t s \Downarrow^{must}$ as required.
- If $t = \text{Eq}(s, s')$ then $\llbracket t \rrbracket \triangleleft_{o \rightarrow o \rightarrow o} t$ by definition.
- The case $t = \mathbf{Y}_T$ is dealt with in Lemma 3.21 above. \square

Thus, by definition of the approximation relation.

Proposition 3.23. *For any program $t : o$, $\llbracket t \rrbracket_{must} = \llbracket \mathcal{U} \rrbracket_{must}$ implies $t \Downarrow_{must} \mathcal{U}$.*

By a standard argument, computational adequacy for our models of bounded and unbounded $\Lambda^+(N)$ implies inequational soundness.

Corollary 3.24. *If $\llbracket s \rrbracket_{must} \sqsubseteq_E \llbracket t \rrbracket_{must}$ then $t \lesssim_{must} s$, and if $\llbracket t \rrbracket_{may} \sqsubseteq_E \llbracket s \rrbracket_{may}$ then $t \lesssim_{may} s$.*

As promised, we may now give an example of must inequivalent terms which cannot be separated by any finite approximant to the argument which does separate them. Let $U = N \rightarrow o \rightarrow o$, so that an element $e \in \llbracket U \rrbracket$ is compact with respect to must approximation if and only if there exists j such that $\llbracket e \rrbracket(k) = \perp_S$ for all $k > j$.

Example 3.25. Consider the terms:

- $s : U \rightarrow o \triangleq \lambda f. (\mathbf{Y}_{N \rightarrow o} \lambda F. \lambda u. (f u) (F(\text{succ}(u)))) 0$
- $t : U \rightarrow o \triangleq \lambda f. ?_N (\lambda v. (\mathbf{Y}_{N \rightarrow o} \lambda F. \lambda u. (\text{Eq}(u, v) \mathcal{U}) ((f u) (F(\text{succ}(u)))))) 0$.

Then $\llbracket s \rrbracket$ is the least upper bound of the stable chain of approximants $\langle \llbracket s_i \rrbracket \mid i \in \omega \rangle$, where $s_i = (f0) ((f1) (\dots ((fi) \Omega) \dots))$ and $\llbracket t \rrbracket$ is the greatest lower bound of the descending extensional chain $\langle \llbracket t_i \rrbracket \mid i \in \omega \rangle$, where $t_i = (f0) ((f1) (\dots ((fi) \mathcal{U}) \dots))$. Then:

- $\llbracket s_i \rrbracket_{must} \sqsubseteq_E \llbracket t_i \rrbracket_{must}$ for all $i \in \omega$, so $\llbracket s \rrbracket_{must} \sqsubseteq_E \llbracket t \rrbracket$ and thus $s \lesssim_{must} t$.
- s and t are distinguished by application to $\lambda u. \lambda x. x$, since we have $s \lambda u. \lambda x. x \Downarrow^{must} \mathcal{U}$ but $t \lambda u. \lambda x. x \Downarrow^{must} \mathcal{U}$.
- If $e : U$ is a finite \lesssim^{must} -approximant to $\lambda u. \lambda x. x$ (i.e. \lesssim^{must} -dominates finitely many elements, which is equivalent to compactness) then $s e \Downarrow^{must}$ if and only if $t e \Downarrow^{must}$, since there exists j such that $\llbracket e \rrbracket(k) = \perp$ for all $k > j$ and thus $s_i e \Downarrow^{must}$ if and only if $t_i e \Downarrow^{must}$ for all $i > j$.

We will study the completeness (full abstraction) properties of our models in Section 7. First, we develop a setting in which we can give intensional characterizations of them.

4. ORDERED CONCRETE DATA STRUCTURES

We will now define the notion of ordered concrete data structure and use it to give an intensional representation of monotone stable functions. It is based on the original notion of concrete data structure, introduced by Kahn and Plotkin [KP93] and further developed by Berry and Curien [BC82].

A concrete data structure consists of sets of *cells*, *values* and *events* (which are pairs of cells and values), and an *enabling relation* between events and cells. The idea is that each step of a sequential computation is represented as an event (the filling of a cell with a value), which may be dependent on some combination of previous events having occurred (as specified by the enabling relation). Programs correspond to *states* which specify values for filling enabled cells — sets of events which satisfy two conditions: *consistency* — every cell must be filled with a unique value — and *safety* — for every filled cell there is a finite chain of enablings of filled cells within the state, back to an “initial cell” which does not depend on any prior events. In order to model nondeterministic computation we adapt this setting in the following ways:

- Removing the consistency condition, so that a cell may be filled with multiple different values.
- Placing an ordering on cells and values, (and thus events) and requiring states to be upwards closed under this ordering. This reflects the fact that (for example) the responses of a program to an input which is a nondeterministic choice of a and b must include all of its responses to both a and b .
- Including a distinct element \bullet — representing failure by divergence or error, depending on may or must testing interpretation — with which any cell may be filled (cf the representation of divergence in the game semantics of must testing in [HM98]).
- Extending the safety condition to allow *transfinite* ordinal chains of enabling events. (Capturing the distinction between infinite, and finite but unbounded interaction.)

Remark 4.1. We may think of a concrete data structure as a kind of *two-player game* in which one player chooses a cell for the second player to fill with a value of their choice, enabling a choice of cells for the first player, and so on. In the deterministic setting, this analogy can be made into a precise correspondence between the category of concrete data structures and sequential algorithms, and a category of *graph games* and strategies [HS02]. However, this correspondence depends on determinacy to decompose the CCC of sequential algorithms into a model of linear logic from which the graph games model is constructed: we leave the description of a more explicitly games-like characterization of ordered concrete data structures as future work.

4.1. Ordered Concrete Data Structures.

Definition 4.2. A (filiform) *ordered concrete data structure* (ocds) A is given by a tuple $(C(A), V(A), \vdash_A, E(A))$ where $C(A), V(A)$ (the *cells* and *values* of A) are partial orders not containing the distinguished element \bullet , $E(A) \subseteq C(A) \times V(A)$ is a set of *events* and $\vdash_A \subseteq (E(A) \cup \{\bullet\}) \times C(A)$ is a relation (*enabling*) such that $(c, v) \vdash c'$ implies $c < c'$.

We write $V(A)_\bullet$ for the partial order on $V(A) \cup \{\bullet\}$ with $a \leq a'$ if $a = \bullet$ or $a \leq a'$ in $V(A)$.

We write $E(A)_\bullet$ for the partial order on $E(A) \cup (C(A) \times \{\bullet\})$, with $(c, a) \leq (c', a')$ if $c \leq c'$ in $C(A)$ and $a \leq a'$ in $V(A)_\bullet$.

Definition 4.3. A *proof* of an event $e \in E(A)_\bullet$ is an *ordinal sequence* of events $(e_\alpha)_{\alpha \leq \kappa}$ in $E(A)_\bullet$ such that $e_\kappa \leq e$ and for $\alpha \leq \kappa$, where $e_\alpha = (c_\alpha, a_\alpha)$:

- If $\alpha = 0$ then $* \vdash c_\alpha$ (c_0 is initial),
- If $\alpha = \beta + 1$ then $e_\beta \vdash_A c_\alpha$,
- If $\alpha = \bigvee_{\beta < \alpha} \beta$, then $c_\alpha = \bigvee_{\beta < \alpha} c_\beta$.

For $x \subseteq E(A)_\bullet$, we write $x \vdash^* e$ if there is a proof of e in x . If this has length κ , we write $x \vdash^\kappa e$. We write $x \vdash^* c$ if $x \cup \{(c, \bullet)\} \vdash^* (c, \bullet)$ (resp. $x \vdash^\kappa c$ if $x \cup \{(c, \bullet)\} \vdash^\kappa (c, \bullet)$).

Note that:

- Every event in a proof in A except (perhaps) the last one is in $E(A)$.
- For any proof $(c_\alpha, a_\alpha)_{\alpha \leq \kappa}$, if $\beta < \alpha \leq \kappa$ then $c_\beta < c_\alpha$.

Definition 4.4. A *state* of an ocds A is a set of events $x \subseteq E(A)_\bullet$ which satisfies:

Upwards Closure: If $e \in x$ and $e \leq e'$ then $e' \in x$.

Safety: If $e \in x$ then there exists $e' \in x$ such that $e' \leq e$ and $x \vdash^* e'$.

We write $D(A)$ for the set of states of A . Some further conditions on states will be useful.

Definition 4.5. A state $x \in D(A)$ satisfies:

- *totality* if $x \subseteq E(A)$ (i.e. no cell is filled with \bullet in x).
- *finite safety* if $e \in x$ implies $x \vdash^i e$ for some $i < \omega$ — i.e. every event in x has a finite proof in x .
- *finite-branching* if $(c, \bullet) \notin x$ implies $\{v \in V(A) \mid (c, v) \in x\}$ is finite, and x satisfies the following *infinite livelock* condition:
if $x \vdash^\lambda c$ and $\lambda \geq \omega$ implies $(c, \bullet) \in x$ — i.e. every cell which has an infinite proof in x is filled with \bullet in x .

Note that any total state which satisfies infinite livelock also satisfies finite safety.

Example 4.6. For any set X , the ocds $\widehat{X} = (\{c\}, X, \{(*, c)\}, \{c\} \times X)$, has a single initial cell c , which may be filled with any value in X (ordered discretely). Clearly, every state of \widehat{X} trivially satisfies finite safety and infinite livelock

- The total states of \widehat{X} are in one-to-one correspondence with the subsets of X ,
- The total and finite-branching states are in one-to-one correspondence with the *finite* subsets of X .
- There is a single non-total (and finite-branching) state containing (c, \bullet) and thus every event of \widehat{X} .

A second key example is the following “universal” ocds, which will play an important role in our full abstraction result.

Definition 4.7. Let \mathcal{U} be the ocds with a countable set of cells $\{c.i \mid i \in \omega\}$, all of which are initial and may be filled by a single value \star .

Notwithstanding the simplicity of \mathcal{U} , for any countable ocds A , each state $x \in A$ may be uniquely represented as a state of \mathcal{U} : let $\{e_i \mid i \in \omega\}$ be an enumeration of $E(A)_\bullet$ and define $\text{inj}(x) \in D(\mathcal{U})$:

$(c.i, \star) \in \text{inj}(x)$ iff $e_i \in x$,

$(c.i, \bullet) \in \text{inj}(x)$ iff $e_i \in x$, where $e_i = (c, \bullet)$ for some c .

The following construction illustrates the role of the partial order on the cells of an ocds. There is a subclass of ocds for which this order is given by the transitive closure of the relation: $c \ll c'$ if there exists v such that $(c, v) \vdash c'$. In fact, this forms a sub-CCC of our category of ocds. However, it is not closed under the following *lifting* operation, which adds an unfillable cell which can therefore not enable any other cells.

Definition 4.8 (Stable Lifting). For any ocds A , let $A_\uparrow = (C(A)_\perp, V(A), \vdash_A \cup \{(*, c_\perp)\}, E(A))$, where $C(A)_\uparrow$ is the lifting of A by adding a new least element c_\perp .

Any state of A_\uparrow which contains (c_\perp, \bullet) must contain all events of $E(A)$ — i.e. the states of A_\uparrow are the states of A , together with one new state containing all possible events.

For a state $x \in D(A)$, we identify the following sets of cells (subsets of $C(A)$):

- $F(x) = \{c \in C(A) \mid \exists a \in V(A)_\bullet. (c, a) \in x\}$ — the set of *filled* cells of x .
- $En(x) = \{c \in C(A) \mid x \vdash^* c\}$ — the set of *enabled* cells of x .
- $A(x) = En(x) - F(x)$ — the set of *accessible* cells of x .

If $c \in En(x)$, and $(c, a) \in E(A)_\bullet$, we write $x + (c, a)$ for the state $x \cup \{e \in E(A)_\bullet \mid (c, a) \leq e\}$.

We will also write $x + (c, V)$ for $\bigcup_{v \in V} (x + (c, v))$.

Say that a state x is *complete* when every cell with a finite proof in x is filled.

Example 4.9. Consider the ocds of *lazy* natural numbers: the set of cells is $\{c_i \mid i \leq \omega\}$, ordered by i . For each $i < \omega$, c_i can be filled with either of the values 0 or suc . c_0 is initial and c_{i+1} is enabled by (c_i, suc) .

Proofs are ascending chains of events $(c_0, \text{suc}), (c_1, \text{suc}), \dots$, and thus:

- The total, complete and finitely safe states are in one-to-one correspondence with the non-empty subsets of \mathbb{N} .
- The total and complete finitely branching states are in one-to-one correspondence with the non-empty finite subsets of \mathbb{N} (since any state in which infinitely many cells are filled must contain a proof of (c_ω, \bullet) by the infinite livelock condition).

The states which are either non-total or incomplete correspond to a subset of \mathbb{N} bounded above by $i \leq \omega$, together with a proof of c_i , which is either filled with \bullet or (for $i < \omega$) unfilled. (Filling c_ω with \bullet corresponds to presence of the recursively defined element $n := \text{suc}(n)$.)

4.2. Ordered Concrete Data Structures as Biorders. For each ocds, we define a biorder in which the set of elements is the set $D(A)$ of states of A . The extensional order is set-inclusion: since any union of states satisfies the safety and upwards-closure conditions, $(D(A), \subseteq)$ is a complete (join) lattice, with least element the empty state and greatest element $\bigcup D(A)$.

Definition 4.10. The stable order on states is defined as follows:

$$x \leq_S y \text{ if } y \subseteq x \text{ and for all } (c, v) \in x, \text{ either } (c, v) \in y \text{ or } (c, \bullet) \in x.$$

$(D(A), \subseteq, \leq_S)$ is a biorder — if X is a non-empty set of states bounded above by y , then $\bigcup X$ is their greatest lower bound in the stable order: if $(c, v) \in \bigcup X$ then $(c, v) \in x$ for

some $x \in X$ and either $(c, \bullet) \in x \subseteq \bigcup X$ or else $(c, v) \in y$ and so for all $x' \in X$, $(c, v) \in x'$. If $z \leq_S x$ for all $x \in X$ then $\bigcup X \subseteq z$ and if $(c, v) \in z$ either $\forall x. (c, v) \in x \subseteq \bigcup X$ or $(c, \bullet) \in z$. $(D(A), \subseteq, \leq_S)$ is pointed — its stably least element is $\bigcup D(A)$.

It is easy to see that $D(A)$ is extensionally complete, but what of stable completeness? In fact, there are pathological examples of ocds in which this property does not hold, but we will show that the ocds for which $D(A)$ is stably complete form a Cartesian closed category. Note that in any case, if Δ is a non-empty set of states with a stable upper bound then its stable l.u.b. is $\bigcap \Delta$.

Evidently, the sub-biorder $D^*(A)$ consisting of states which are finitely safe is also an (extensionally complete) biorder, as this property is preserved by arbitrary unions of states. The sub-biorder $D_*(A)$ consisting of finite branching states is a biorder, although not, in general extensionally complete, since the infinite union of finite-branching states may not be finite-branching. However, it is stably complete.

Lemma 4.11. *If $X \subseteq D_*(A)$ is stably bounded above by $y \in D_*(A)$, then $\bigcup X \in D_*(A)$.*

Proof. Suppose there is a cell c and infinite set of values V such that for all $v \in V$, $(c, v) \in \bigcup X$. If every such event occurs in y , then $(c, \bullet) \in y \subseteq \bigcup X$. Otherwise there exists $x \in X$ with $(c, v) \in x$ and $(c, v) \notin y$ and so $(c, \bullet) \in x \subseteq \bigcup X$. Suppose there is an event (c, \bullet) which has an infinite proof in $\bigcup X \cup \{(c, \bullet)\}$. If this proof is in y then $(c, \bullet) \in y \subseteq \bigcup X$. Otherwise there is an event (c', v) in the proof such that $(c, v) \notin y$, and hence $x \in X$ such that $(c', \bullet) \in x \subseteq \bigcup X$ such that $(c', \bullet) \in x$ and so $(c, \bullet) \in \bigcup X$ by upwards closure. \square

We shall use the following alternative characterizations of the stable order and stable boundedness.

Definition 4.12. Given a state x and $C \subseteq \text{En}(x)$, let $x_C = \bigcup \{x + (c, \bullet) \mid c \in C\}$.

Lemma 4.13. *$y \leq_S x$ iff $y = x_C$ for some $C \subseteq \text{En}(x)$.*

Proof. Evidently, $x_C \leq_S x$, so it suffices to show that every element $y \leq_S x$ has this form. Let $C = \{c \in \text{En}(x) \mid (c, \bullet) \in y\}$, so that $x_C \subseteq y$. We claim that $y \subseteq x_C$: suppose $(c, a) \in y$, but $(c, a) \notin x$. Let $\langle (c_\alpha, v_\alpha) \mid \alpha < \kappa \rangle$ be a proof of (c, a) in y . Then either $(c, a) \in x$ or there exists a least value α such that $(c_\alpha, v_\alpha) \notin x$. So $c_\alpha \in E(x)$, and $(c, \bullet) \in y$ by stability and so $c_\alpha \in C$. Since $c_\alpha \sqsubseteq_E c$, $(c, a) \in x_C$ as required. \square

Lemma 4.14. *A non-empty set X of states is stably bounded above if (and only if):*

(\dagger) *for all $(c, v) \in x \in X$, either $(c, \bullet) \in x$ or $(c, v) \in \bigcap X$.*

Proof. From left-to-right, this follows from the fact that X is bounded above by $\bigcap X$. From right-to-left, we need to show that if X satisfies (\dagger) then $\bigcap X$ satisfies the safety condition, since it is then a stable upper bound. Suppose $(c, v) \in \bigcap X$ — choose $x \in X$ and let P be a shortest proof of (c, v) in $x \in X$ (i.e. there is no shorter proof). Then $P \subseteq \bigcap X$, since for every preceding element (c', v') in P we have $(c', \bullet) \notin x$ (or else a shorter proof would exist) and so $(c', v') \in \bigcap X$. \square

Returning to our examples, we may observe that:

- For any set X , $D(\widehat{X}) \cong \mathcal{P}(X)_\uparrow$ and $D_*(\widehat{X}) \cong \mathcal{P}_{fin}(X)_\uparrow$. In particular, $D(\widehat{\emptyset})$ is the two-element space Σ .
- $D(U) \cong \llbracket U \rrbracket$, where $U = N \rightarrow o \rightarrow o$ — it is the countable product of copies of $\llbracket o \rightarrow o \rrbracket \cong \widehat{1}$. Moreover, for any ocds A , the map sending $x \in D^*(A)$ to $\text{inj}(x)$ is a continuously stable

function, with a left inverse — $\mathbf{proj}(y) = \bigcap \{x \in D^*(A) \mid \mathbf{inj}(x) \leq_S y\}$ — i.e. \mathcal{U} is a *universal object* among the countable ocds and continuously stable functions.

- For any ocds A , $D(A_\uparrow) \cong D(A)_\uparrow$ (and $D^*(A_\uparrow) \cong D^*(A)_\uparrow$, $D_*(A_\uparrow) \cong D_*(A)_\uparrow$) — lifting adds a single element which is stably least/extensionally greatest.

5. SEQUENTIAL ALGORITHMS ON ORDERED CONCRETE DATA STRUCTURES

By the results in the previous section, we may define the following categories:

- \mathcal{OC} — objects are ordered concrete data structures, morphisms from A to B are monotone stable functions from $D(A)$ to $D(B)$.
- $\mathcal{OC}_\mathcal{E}$ — objects are ocds and morphisms from A to B are continuously stable functions from $D^*(A)$ to $D^*(B)$.
- $\mathcal{OC}_\mathcal{S}$ — objects are ocds and morphisms from A to B are stably continuous functions from $D_*(A)$ to $D_*(B)$.

Each category has products, given by the disjoint union of ocds: $\prod_{i \in I} (C_i, V_i, E_i, \vdash_i) = (\prod_{i \in I} C_i, \prod_{i \in I} V_i, \bigcup_{i \in I} \{(c.i, v.i) \mid (c, v) \in E_i\}, \bigcup_{i \in I} \{((c.i, v.i), d.i) \mid ((c, v), d) \in E_i\})$. The full, identity-on-morphisms functor $D : \mathcal{OC} \rightarrow \mathcal{B}$ sending each concrete data structure to its set of states (which restricts to functors $D^* : \mathcal{OC}_\mathcal{E} \rightarrow \mathcal{B}_\mathcal{E}$ and $D_* : \mathcal{OC}_\mathcal{S} \rightarrow \mathcal{B}_\mathcal{S}$) preserves products, and so to establish Cartesian closure in each case, it remains to define an internal hom ocds for each pair of objects and show that its biorder of states is isomorphic to the internal hom (function space) in the corresponding category of biorders. This is a key result, since it establishes that every stable and monotone function between the biorder of states of an ocds is computed by a unique state of this exponential ocds or *sequential algorithm*.

Focussing first on the general case of monotone stable functions, for each pair of ocds A, B we define an ordered concrete data structure $A \Rightarrow B$ (cf. the analogous definition of unordered concrete data structure [BC82]) and show that $D(A \Rightarrow B) \cong [D(A), D(B)]_\mathcal{B}$ in \mathcal{B} .

Cells: A cell of $A \Rightarrow B$ is given by a pair of a total state of A and a cell of B :

$$C(A \Rightarrow B) = (D(A) \cap \mathcal{P}(E(A))) \times C(B) \text{ — with } (x, c) \leq (x', c') \text{ if } x \subseteq x' \text{ and } c \leq c'.$$

Values: A *value* of $A \Rightarrow B$ is either a cell from A or a value from B — the order being determined pointwise from that of $V(B)$ and the reverse of $C(A)$:

$$V(A \Rightarrow B) = C(A)^c \uplus V(B)$$

Events: A cell (x, c) of $A \Rightarrow B$ may be *filled* with either a cell accessible from x in A or a value in B which can fill c :

$$E(A \Rightarrow B) = \{((x, c), c'.1) \mid (x, c) \in C(A \Rightarrow B) \wedge c' \in A(x)\} \cup \{((x, c), v.2) \mid (x, c) \in C(A \Rightarrow B) \wedge (c, v) \in E(B)\}$$

Enabling: An event $((x, c), c'.1)$ enables a cell (x', c) if x' may be obtained from x by filling c' , and an event $((x, c), v.2)$ enables a cell (x, c') if (c, v) enables c' in B : $\vdash_{A \Rightarrow B} = \{(((x, c), c'.1), (x', c)) \in E(A \Rightarrow B) \times C(A \Rightarrow B) \mid \exists V \subseteq V(A). x' = x + (c', V)\} \cup \{(((x, c), v.2), (x, c')) \in E(A \Rightarrow B) \times C(A \Rightarrow B) \mid (c, v) \vdash_B c'\}$

A *sequential algorithm* from A to B is a state of $A \Rightarrow B$. In the following, we will assume that the sets of cells of A and values of B are disjoint, and so omit explicit tagging.

5.1. Stable Functions from Sequential Algorithms. Every sequential algorithm $\sigma \in D(A \Rightarrow B)$ computes a monotone stable function $\mathbf{fun}(\sigma)$ from $D(A)$ to $D(B)$. Given a state $x \in D(A)$, define: $\mathbf{fun}(\sigma)(x) =$

$$\{(c, a) \in E(B)_\bullet \mid \exists x' \subseteq x. ((x', c), a) \in \sigma \vee \exists c'. (c', \bullet) \in x \wedge ((x', c), c') \in \sigma\}$$

Lemma 5.1. *For any $x \in D(A)$, $\text{fun}(\sigma)(x)$ is a state in $D(B)$.*

Proof. $\text{fun}(\sigma)(x)$ satisfies:

Upwards closure: Suppose $(c', a') \geq (c, a) \in \text{fun}(\sigma)(x)$. If there exists $x' \subseteq x$ with $((x', c), a) \in \sigma$ then $((x', c'), a') \geq ((x', c), a)$ and so $((x', c'), a') \in \sigma$ and $(c', a') \in \text{fun}(\sigma)(x)$.

If there exists $(c'', \bullet) \in x$ such that $((x', c), c'') \in \sigma$ then $((x', c'), c'') \geq ((x', c), c'')$ so $((x', c'), c'') \in \sigma$ and hence $(c', a') \in \text{fun}(\sigma)(x)$.

Safety: Suppose $(c, a) \in \text{fun}(\sigma)(x)$. Then there exists $((y, c), b) \in \sigma$ with $y \subseteq x$, and a proof of (y, c) in σ which therefore restricts to a proof of c in $f(y)$. \square

Say that a function $f : D(A) \rightarrow D(B)$ is computed by the sequential algorithm σ if $f(x) = \text{fun}(\sigma)(x)$ for all $x \in D(A)$. We now show that any such function is monotone stable.

Lemma 5.2. *If $\uparrow X$ and $(c, a) \in \text{fun}(\sigma)(\bigcup X)$ then either $(c, a) \in \text{fun}(\sigma)(x)$ for all $x \in X$ or $(c, \bullet) \in \text{fun}(\sigma)(x')$ for some $x' \in X$.*

Proof. If $(c, a) \in \text{fun}(\sigma)(\bigcup X)$ then there exists an event $e \in \sigma$ such that either $e = ((w, c), a)$, where $w \subseteq \bigcup X$ or $e = ((w, c), c')$, where $w + (c', \bullet) \subseteq \bigcup X$. If $w \subseteq x$ for every $x \in X$ then in the first case $(c, a) \in \bigcup \text{fun}(\sigma)(x)$ for every $x \in X$, and in the second case there exists $x' \in X$ with $w + (c', \bullet) \subseteq x'$ and so $(c, \bullet) \in \text{fun}(\sigma)(x')$.

So suppose $w \not\subseteq x$ for some $x \in X$. Fixing a proof of e in σ , let $e' = ((w', c'), a')$ be the first element in this proof such that $w' \not\subseteq x$ for some $x \in X$. Then there is an immediately preceding event $((w'', c''), c'')$ such that $w' = w'' + (c'', V)$ for some set of values V , including a value u such that $(c'', u) \not\subseteq x$. Because $w' \subseteq \bigcup X$, there exists $x' \in X$ with $(c'', u) \in x'$. Since $x \uparrow x'$, therefore $(c'', \bullet) \in x'$, and hence $(c', \bullet) \in \text{fun}(\sigma)(y)$. Since $c' \leq c$, we have $(c, \bullet) \in \text{fun}(\sigma)(x')$ as required. \square

Proposition 5.3. *For any sequential algorithm σ , $\text{fun}(\sigma)$ is a monotone stable function.*

Proof. Evidently, if $x \subseteq y$ then $\text{fun}(\sigma)(x) \subseteq \text{fun}(\sigma)(y)$. If $\uparrow X$, then $\text{fun}(\sigma)(\bigcup X) = \bigcup_{x \in X} \text{fun}(\sigma)(x)$: if $(c, a) \in \text{fun}(\sigma)(\bigcup X)$ then by Lemma 5.2, either $(c, a) \in \text{fun}(\sigma)(x)$ for all $x \in X$ or $(c, \bullet) \in \bigcup_{x \in X} \text{fun}(\sigma)(x)$. \square

5.2. Examples. We now describe the sequential algorithms which compute some of the monotone stable functions described in the examples of Sections 2 and 3.

Example 5.4. For any set X , consider the ocds $\widehat{\mathcal{O}}^X \Rightarrow \widehat{\mathcal{O}}$ — i.e. $\widehat{\mathcal{O}}^X$ is the X -indexed product of copies of $\widehat{\mathcal{O}}$, which has (initial) cells $\{c_x \mid x \in X\}$ and no values, so $\widehat{\mathcal{O}}^X \Rightarrow \widehat{\mathcal{O}}$ has a single initial cell $(\{-\}, c)$, which may be filled by any of the cells c_x for $x \in X$. In other words, it is the same ocds as \widehat{X} , up to renaming of cells and values, and so $D(\widehat{\mathcal{O}}^X \Rightarrow \widehat{\mathcal{O}}) \cong D(\widehat{X}) \cong \mathcal{P}(X)_\uparrow \cong [\widehat{\mathcal{O}}^X, \widehat{\mathcal{O}}]_{\mathcal{B}}$ by Proposition 3.15, and $D_*(\widehat{\mathcal{O}}^X \Rightarrow \widehat{\mathcal{O}}) \cong D_*(\widehat{X}) \cong \mathcal{P}_{\text{fin}}(X)_\uparrow \cong [\widehat{\mathcal{O}}^X, \widehat{\mathcal{O}}]_{\mathcal{S}\mathcal{B}}$.

By considering the algorithms computing the functions of Example 2.1, we may see that unlike a non-deterministic strategy in e.g. the games model of [HM98], a non-deterministic sequential algorithm may have different decompositions into a non-deterministic choice between two (complete and total) algorithms.

Example 5.5. Consider the OCDS $\widehat{\mathbb{B}} \times \widehat{\mathbb{B}} \Rightarrow \widehat{\mathbb{B}}$.

- Cells are pairs (x, c) , where x is a total state of $\widehat{\mathbb{B}} \times \widehat{\mathbb{B}}$ (corresponding to a pair of subsets of \mathcal{B}).
- Values are either a Boolean (filling (x, c) for any x) or a cell $(c_0$ or $c_1)$ of $\widehat{\mathbb{B}} \times \widehat{\mathbb{B}}$, filling (x, c) provided it is not filled in x .

So a sequential algorithm in $\widehat{\mathbb{B}} \times \widehat{\mathbb{B}} \Rightarrow \widehat{\mathbb{B}}$ may investigate (both, either or neither) of its arguments in any order, and return a Boolean value. The functions denoted by the terms $t_{i,j}$ of Example 2.1 are computed by sequential algorithms $\sigma_{i,j}$ as follows:

$$\begin{array}{ccc}
 \widehat{\mathbb{B}} \times \widehat{\mathbb{B}} & \Rightarrow & \widehat{\mathbb{B}} \\
 & & (\{-\}, c) \\
 c_i & & (\{(c_i, v_i)\}, c) \\
 c_{i^*} & & (\{(c_0, v_0), (c_1, v_1)\}, c) \\
 & & v_j
 \end{array}$$

Observe that the cell $(\{(c_0, v_0), (c_1, v_1)\}, c)$ does not depend on i or j , and thus $\sigma_{0,0} \cup \sigma_{1,1} = \sigma_{0,1} \cup \sigma_{1,0}$.

In fact, we can present sequential algorithms from A into $\widehat{\mathcal{O}}$ (which might be considered as continuations on A) more simply by ignoring the unfillable cell of $\widehat{\mathcal{O}}$ — i.e. cells correspond to states of A ordered by inclusion, values to cells of A , and an event to a pair (x, c) of a state of A and a cell which is accessible from it, which enables the states (new cells) of the form $x + (c, V)$. The sequential algorithms computing the functions of Example 2.2 may be presented in this way:

Example 5.6. Consider the concrete data structure $(\widehat{\mathbb{B}} \Rightarrow \widehat{\mathcal{O}}) \Rightarrow \widehat{\mathcal{O}}$. Ignoring the unfillable cell of $\widehat{\mathcal{O}}$, its initial cell is the empty state, which can be filled (only) by the initial cell of $\widehat{\mathbb{B}} \Rightarrow \widehat{\mathcal{O}}$ (also the empty state). (\emptyset, \emptyset) enables the cell (state of $\widehat{\mathbb{B}} \Rightarrow \widehat{\mathcal{O}}$) $\{(\emptyset, c)\}$, where c is the (unique) initial cell of $\widehat{\mathbb{B}}$, which may be filled with the values (non-empty states of $\widehat{\mathbb{B}}$) $\{(c, \mathbf{tt})\}$, $\{(c, \mathbf{ff})\}$ and $\{(c, \mathbf{tt}), (c, \mathbf{ff})\}$. There is a non-trivial order on these values: $\{(c, \mathbf{tt})\}, \{(c, \mathbf{ff})\} \geq \{(c, \mathbf{tt}), (c, \mathbf{ff})\}$, so by upwards closure any cell which is filled with the last must also be filled with the first two. This distinguishes the sequential algorithms computing $\lambda f.f(\mathbf{tt} + \mathbf{ff})$ and $\lambda f.(f \mathbf{tt}) + (f \mathbf{ff})$: in the first case $\{(\emptyset, c)\}$ is filled with $\{(c, \mathbf{tt}), (c, \mathbf{ff})\}$, in the second case it is not. (See Figure 1 for the extensional hierarchy of sequential algorithms at this type.)

The above examples relate to finitary (or, at least, finitely safe) sequential algorithms. Recasting Example 3.25 furnishes an instance of an event with an infinite proof which is required to separate the sequential algorithms computing the terms s and t defined there.

Example 5.7. Consider the ocds $\mathcal{U} \Rightarrow \widehat{\mathcal{O}}$: again, ignoring the unfillable cell on the right this has as cells total states of \mathcal{U} (which correspond to subsets of \mathbb{N}), and as values, cells of \mathcal{U} (which correspond to elements of \mathbb{N}) — an event is thus a pair (X, x) such that $x \notin X$,

$$\begin{array}{c}
\emptyset (\perp_E) \\
\{(\emptyset, \emptyset)\} (\lambda f.f \perp_E) \\
\{(\emptyset, \emptyset), ((\emptyset, c), \{(c, \mathbf{tt})\})\} (\lambda f.f \mathbf{tt}) \quad \{(\emptyset, \emptyset), ((\emptyset, c), \{(c, \mathbf{ff})\})\} (\lambda f.f \mathbf{ff}) \\
\{(\emptyset, \emptyset), ((\emptyset, c), \{(c, \mathbf{tt})\}), ((\emptyset, c), \{(c, \mathbf{ff})\})\} (\lambda f.(f \mathbf{tt}) + (f \mathbf{ff})) \\
[\{(\emptyset, \emptyset), ((\emptyset, c), \{(c, \mathbf{tt}), (c, \mathbf{ff})\})\}] (\lambda f.f (\mathbf{tt} + \mathbf{ff})) \\
[\{(\emptyset, \emptyset), ((\emptyset, c), \bullet)\}] (\lambda f.f \top_E) \\
[\{(\emptyset, \bullet)\}] (\top_E)
\end{array}$$

Figure 1: Sequential Algorithms on $(\widehat{\mathbb{B}} \Rightarrow \widehat{\emptyset}) \Rightarrow \widehat{\emptyset}$ (ordered extensionally)

with $(X, x) \vdash Y$ iff $Y = X \cup \{x\}$. A sequential algorithm on $\mathcal{U} \Rightarrow \widehat{\emptyset}$ thus corresponds to a set of finite non-repeating sequences over \mathbb{N} , together with a set of divergences — finite and infinite sequences over \mathbb{N} ending with a cell filled with \bullet — such that any sequence which extends a divergence appears in both sets. Compare with Harner and McCusker’s game semantics of non-determinism [HM98].

The sequential algorithms on $\mathcal{U} \Rightarrow \widehat{\emptyset}$ which compute the functions denoted by s and t defined in Example 3.25 contain the same finite sequences of events (corresponding to finite prefixes of 0123...). However, the algorithm computing s contains a divergence corresponding to the infinite sequence 0123..., whereas the algorithm computing t does not. To be more precise, for each $j \leq \omega$, let P_j be the state $\{(c.i, \star) \mid i < j\}$ in $D(\mathcal{U})$, and for each $k \leq \omega$ let Q_k be the (upwards closure of) $\{((P_j, c), c_j) \mid j < k\}$ in $D(\mathcal{U} \Rightarrow \widehat{\emptyset})$. Then $\llbracket s \rrbracket$ is computed by $Q_\omega \cup \{((P_\omega, c), \bullet)\}$ and $\llbracket t \rrbracket$ is computed by Q_ω — i.e. they are differentiated only by the event $((P_\omega, c), \bullet)$, which is the conclusion of the infinite proof $((P_j, c), c_j)_{j < \omega}$ in Q_ω .

One might wonder whether the only infinite proofs which are needed to distinguish terms of $\Lambda(N)^+$ are the ones that lead to a divergence in this way. However, at the next order of types (third order) we may give examples of terms denoting *total* sequential algorithms which are distinguished only by events with infinite proofs. Define $p : (U \rightarrow o) \rightarrow o \triangleq \lambda g.g \lambda u.\lambda x.x$, and $r : (U \rightarrow o) \rightarrow o \triangleq \lambda u.\lambda x.?_N \lambda v.(1\mathbf{eq}(u, v) x) \mathcal{U}$, so that $\llbracket p \rrbracket$ is computed by applying its argument to the infinite state $P_\omega \in D(\mathcal{U})$ and $\llbracket r \rrbracket$ by non-deterministically applying its argument to some finite state P_j in \mathcal{U} . Writing $\llbracket p \rrbracket$ and $\llbracket r \rrbracket$ for these sequential algorithms:

- every event in $\llbracket p \rrbracket$ with a *finite proof* (i.e. the finite states P_j for $j < \omega$) also occurs in $\llbracket r \rrbracket$.
- p and r are distinguished by application to t — $p t \not\Downarrow_{\text{must}}$ and $r t \Downarrow_{\text{must}}$, since $\llbracket p \rrbracket$ contains the event $((Q_\omega, c), P_\omega)$ — which has the proof $((Q_k, c), P_k)_{k < \omega}$ in $\llbracket p \rrbracket$ — but $\llbracket r \rrbracket$ does not.

We may extend these examples to give terms which are distinguished only by events with proofs of length $\omega + l$, for any $l < \omega$. Defining $U^0 \triangleq U$ and $U^{l+1} \triangleq (U \rightarrow o) \rightarrow o$, we give terms $p^l : U^l$ and $t^l, s^l : U^l \rightarrow o$ by $p^0 \triangleq \lambda u.\lambda x.x$ and $p^{l+1} \triangleq \lambda f.f p^l$, $t^0 \triangleq t$ and $t^{l+1} \triangleq \lambda f.f t^l$ and $s^0 \triangleq s$ and $s^{k+1} \triangleq \lambda f.f s^k$. Then $t^l p^l \not\Downarrow_{\text{must}}$ for each l , but $s^l p^l \Downarrow_{\text{must}}$.

Denotationally, for each $l \in \omega$, $\llbracket t^l \rrbracket = \llbracket s^l \rrbracket \cup \{((P_{\omega+l}, c), \bullet)\}$, where $P_j^l \in D(U^l \Rightarrow \widehat{\emptyset})$ and $Q_j^l \in D(\llbracket U^l \Rightarrow \widehat{\emptyset} \rrbracket \Rightarrow \widehat{\emptyset})$ are defined inductively for $1 \leq j \leq \omega + l$ by $P_j^0 = P_j$ and $Q_j^0 = Q_j$, and $P_j^{l+1} = \{((Q_i^l, c), P_i^l) \mid i < j\}$ and $Q_j^{l+1} = \{((P_i^{l+1}, c), Q_i^l) \mid i < j\}$.

Observe that the proof of $(P_{\omega+l}, c)$ in $\llbracket t^l \rrbracket$ is the $\omega + l$ sequence of events $((P_i^{l+1}, c), Q_i^l)_{j < \omega + l}$. In the simply-typed language $\Lambda(N)^+$, it appears that this is as far as it is necessary to go —

we conjecture that any inequivalent terms at a type of order at most $2l$ may be distinguished by an event with a proof no longer than $\omega + l$.

6. STABLE FUNCTIONS AND SEQUENTIALITY

We now return to the task of showing that the functions between the biorders of states of ocdds which are computed by sequential algorithms are precisely the monotone and stable ones. Recall that concrete data structures were originally introduced by Kahn and Plotkin in order to give a description of sequentiality for higher-order deterministic functionals [KP93]. Essentially, in this setting a function between the states of concrete data structures A and B is Kahn-Plotkin sequential if any argument (state of A) x , and cell c of B which is filled in $f(y)$ for some y which extends x , can be associated with a cell, accessible from x , which must be filled in any state z (which extends x) such that c is filled in $f(z)$. However, in this original setting, divergence is represented *implicitly*, by not filling an enabled cell (rather than as an explicit divergence by filling a cell with \bullet), and inclusion of states corresponds to the *stable* order. Thus we translate this original definition of Kahn-Plotkin sequentiality to the current setting by (essentially) replacing the role of “accessible cell” with that of “cell filled with \bullet ”, and “filled cell” with “enabled cell not filled with \bullet ”. We define a partial order (\preceq) on total states (which plays the role of the stable order in the original definition of Kahn-Plotkin sequentiality):

$$x \preceq y \text{ if } x \subseteq y \text{ and if } c \in F(x) \text{ then } (c, v) \in y \text{ implies } (c, v) \in x.$$

Definition 6.1. A function $f : D(A) \rightarrow D(B)$ is *explicitly sequential* if whenever x, y are total states such that $x \preceq y$ then for any event $(c, v) \in f(y)$, either $(c, v) \in f(x)$, or there exists $c' \in A(x) \cap F(y)$ such that if $\uparrow \{x, z\}$ and $(c', \bullet) \in z$ then $(c, \bullet) \in f(z)$.

We will now show that all monotone stable functions are explicitly sequential.

Lemma 6.2. *If $x \preceq y$ then $y \subseteq x_{A(x) \cap F(y)}$.*

Proof. Suppose $(c, v) \in y$ but $(c, v) \notin x$. Let P be a proof of c in y , and (c', v') the least element of P which is not in x . Then c' is accessible (enabled but not filled) in x , as $x \preceq y$ — i.e. $c' \in A(x) \cap F(y)$, so $(c, v) \geq (c', \bullet) \in x_{A(x) \cap F(y)}$. \square

Proposition 6.3. *Every monotone stable function $f : D(A) \rightarrow D(B)$ is explicitly sequential.*

Proof. Suppose $x \preceq y$ and $(c, v) \in f(y)$ but $(c, v) \notin f(x)$. Since $y \subseteq x_{A(x) \cap F(y)}$, $(c, v) \in f(x_{A(x) \cap F(y)})$ and as $f(x_{A(x) \cap F(y)}) \leq_S f(x)$, so $(c, \bullet) \in f(x_{A(x) \cap F(y)})$.

By conditional multiplicativity, $f(x_{A(x) \cap F(y)}) = \bigcup_{c' \in A(x) \cap F(y)} f(x + (c', \bullet))$ and so there is a cell $c' \in A(x) \cap F(y)$ such that $(c, \bullet) \in f(x + (c', \bullet))$ as required. \square

We establish the converse — that every explicitly sequential function is monotone stable — by showing below that every explicitly sequential function from $D(A)$ to $D(B)$ is computed by a sequential algorithm.

6.1. Sequential Algorithms from Monotone Stable Functions. We will now use the explicit sequentiality property to establish that every monotone stable function $f : D(A) \rightarrow D(B)$ is computed by a sequential algorithm $\text{strat}(f) \in D(A \Rightarrow B)$. Define this to be the set of events:

$$\begin{aligned} & \{((x, c), a) \in C(A \Rightarrow B) \times V(B)_\bullet \mid (c, a) \in f(x)\} \\ \cup & \{((x, c), c') \in C(A \Rightarrow B) \times C(A) \mid (c, \bullet) \in f(x + (c', \bullet))\} \end{aligned}$$

Lemma 6.4. $\text{strat}(f)$ is an upper set.

Proof. Suppose $((x', c'), a') \geq ((x, c), a) \in \text{strat}(f)$. If $a, a' \in V(B)_\bullet$ — i.e. $(c, a) \in f(x)$ — then $(c', a') \geq (c, a) \in f(x') \supseteq f(x)$, and so $((x', c'), a') \in \text{strat}(f)$. If $a, a' \in C(A)$, so that $a' \leq a$ and hence $x + (a, \bullet) \subseteq x' + (a', \bullet)$, then $(c, \bullet) \in f(x' + (a', \bullet))$ and so $((x', c'), a') \in \text{strat}(f)$ as required. \square

Lemma 6.5. $\text{strat}(f)$ satisfies the safety property.

Proof. We construct a proof of each event $((x, c), a)$ in $\text{strat}(f)$ using the explicit sequentiality property for f . Suppose $a \in V(B)_\bullet$: let $\langle (c_\beta, v_\beta) \mid \beta \leq \lambda \rangle$ be a proof of (c, a) in $f(x)$.

For each ordinal α , we define:

- A state $x_\alpha \preceq x$.
- An ordinal $\kappa(\alpha) \leq \alpha$
- An event $e_\alpha \in E(A \Rightarrow B)$ such that $e_\alpha = ((x_\alpha, c_{\kappa(\alpha)}), v_{\kappa(\alpha)})$ or $e_\alpha = ((x_\alpha, c_{\kappa(\alpha)}), c')$ for some $c' \in F(x)$.

such that if $\kappa(\alpha) < \lambda$, then $\langle e_\gamma \mid \gamma \leq \alpha \rangle$ is a proof of e_α .

- Let $x_0 = \{\}$ and $\kappa(0) = 0$,
- If $\alpha = \bigcup \{\beta < \alpha\}$ then $x_\alpha = \bigcup_{\beta < \alpha} x_\beta$ and $\kappa(\alpha) = \bigcup \{\kappa(\beta) \mid \beta < \alpha\}$,
- For all α , if $(c_{\kappa(\alpha)}, v_{\kappa(\alpha)}) \in f(x_\alpha)$ then let $e_\alpha = ((x_\alpha, c_{\kappa(\alpha)}), v_{\kappa(\alpha)})$ and $x_{\alpha+1} = x_\alpha$ and $\kappa(\alpha + 1) = \min\{\lambda, \kappa(\alpha) + 1\}$.

Otherwise, by the explicit sequentiality of f (Proposition 6.3) there is a cell $c' \in A(x_j) \cap F(x)$ such that $(c_{\kappa(\alpha)}, \bullet) \in f(x_\alpha + (c', \bullet))$ and so we may set $e_\alpha = ((x_\alpha, c_{\kappa(\alpha)}), c')$, and $x_{\alpha+1} = \bigcup \{x_\alpha + (c', v) \mid (c', v) \in x\}$ and $\kappa(\alpha + 1) = \kappa(\alpha)$.

Since $\kappa(\alpha) \neq \lambda$ implies $e_\alpha \neq e_\beta$ for all $\beta < \alpha$, there must be some (least) α (no greater than the cardinality of $E(A \Rightarrow B)$) such that $\kappa(\alpha) = \lambda$ and so $\langle e_j \mid \beta \leq \alpha \rangle$ is a proof of $((x, c), a)$ in $\text{strat}(f)$. The case in which a is a cell in $C(A)$ is similar. \square

Hence we have shown that:

Proposition 6.6. For any monotone stable function $f : D(A) \rightarrow D(B)$, $\text{strat}(f)$ is a well-defined state of $A \Rightarrow B$.

We now show that fun and strat are inverse — i.e. a monotone stable function f is computed by $\sigma \in D(A \Rightarrow B)$ if and only if $\sigma = \text{strat}(f)$. Let $x^\top = \{(c, v) \in x \mid (c, \bullet) \notin x\}$ (the maximal total state contained in x).

Lemma 6.7. x^\top is a well-defined state.

Proof. x^\top is evidently an upper set. For safety, suppose $(c, v) \in x$: for any event (c', v') in any proof of c , $c' \leq c$ and so if $(c', v) \notin x^\top$ then $(c, v) \notin x^\top$ — i.e. if $(c, v) \in x$ then any proof of c in x is a proof of c in x^\top . \square

By definition, if $(c, v) \in x$ then either $(c, \bullet) \in x$ or $(c, v) \in x^\top$ — i.e. $x \leq_S x^\top$.

Lemma 6.8. *For any monotone stable function $f : D(A) \rightarrow D(B)$ and state $x \in D(A)$: $\text{fun}(\text{strat}(f))(x) = f(x)$.*

Proof. Suppose $(c, a) \in \text{fun}(\text{strat}(f))(x)$. Then either:

- there exists a total $y \subseteq x$ with $((y, c), a) \in \text{strat}(f)$, and so $(c, a) \in f(y) \subseteq f(x)$.
- or there exists $y \subseteq x$ with $((y, c), c') \in \text{strat}(f)$ and $(c', \bullet) \in x$, and so $(c, \bullet) \in f(y + (c', \bullet)) \subseteq f(x)$.

For the converse, suppose $(c, a) \in f(x)$. If $(c, a) \in f(x^\top)$ then $((x^\top, c), a) \in \text{strat}(f)$, and so $(c, a) \in \text{fun}(\text{strat}(f))(x)$ as required.

Otherwise $(c, \bullet) \in f(x)$, and by the explicit sequentiality of f (Proposition 6.3), there exists $c' \in A(x^\top) \cap F(y)$ such that $(c', \bullet) \in x$ and $(c, \bullet) \in f(x^\top + (c', \bullet))$ and so $((x^\top, c), c') \in \text{strat}(f)$ and $(c, a) \in \text{fun}(\text{strat}(f))(x)$ as required. \square

Lemma 6.9. *For all sequential algorithms $\sigma \in D(A \Rightarrow B)$, $\text{strat}(\text{fun}(\sigma)) = \sigma$.*

Proof. Suppose $((x, c), a) \in E(A \Rightarrow B)_\bullet$. Then:

- If $a \in V(B)_\bullet$ then $((x, c), a) \in \sigma$ if and only if $(c, a) \in \text{fun}(\sigma)(x)$ if and only if $((x, c), a) \in \text{strat}(\text{fun}(\sigma))$.
- If $a \in C(A)$ then $((x, c), a) \in \sigma$ if and only if $(c, \bullet) \in \text{fun}(\sigma)(x + (c', \bullet))$ if and only if $((x, c), a) \in \text{strat}(\text{fun}(\sigma))$. \square

To complete the proof that fun is a *(bi)order-isomorphism* from $D(A \Rightarrow B)$ to $[D(A), D(B)]$ we need to show that it is a biorder-embedding — i.e. that fun and strat are monotone in both orders (from which stability follows). Monotonicity with respect to the extensional order is straightforward, leaving the stable order.

Lemma 6.10. *If $\sigma \leq_S \tau$ then $\text{fun}(\sigma) \leq_S \text{fun}(\tau)$.*

Proof. First, we show that for all x , $\text{fun}(\sigma)(x) \leq_S \text{fun}(\tau)(x)$. Suppose $(c, v) \in \text{fun}(\sigma)(x)$ but $(c, \bullet) \notin \text{fun}(\sigma)(x)$. Then there exists $x' \subseteq x$ with $((x', c), v) \in \sigma$ and $((x', c), \bullet) \notin \sigma$ and so $((x', c), v) \in \tau$ and $(c, v) \in \text{fun}(\tau)(x)$ as required.

Now we need to show that for all $x \leq_S y$, $\text{fun}(\sigma)(x) = \text{fun}(\tau)(x) \cup \text{fun}(\sigma)(y)$. Suppose $(c, a) \in \text{fun}(\sigma)(x)$, we need to show that $(c, a) \in \text{fun}(\tau)(x)$ or $(c, a) \in \text{fun}(\sigma)(y)$. By Proposition 5.3, if $(c, a) \notin \text{fun}(\sigma)(y)$, then $(c, \bullet) \in \text{fun}(\sigma)(x)$, and so we may assume that $a = \bullet$. So either there exists an event $((z, c), \bullet) \in \sigma$ with $z \subseteq x$ or else there exists $z + (c', \bullet) \subseteq x$ such that $((z, c), c') \in \sigma$. But this latter case reduces to the first one, since either $((z, c), c') \in \tau$ (and so $(c, \bullet) \in \text{fun}(\tau)(x)$ and we are done), or else $((z, c), \bullet) \in \sigma$.

Assuming $(c, \bullet) \notin \text{fun}(\sigma)(y)$, let P be a proof of $((z, c), \bullet)$ in σ , and let $((z', c'), a)$ be the least element of P such that $z' \not\subseteq y$. Then there must be an immediately preceding event in P of the form $((z'', c'), c'')$, where $z' = z'' + (c'', V)$ for some V , and hence $z'' + (c'', \bullet) \subseteq x$, as $\uparrow \{x, y\}$. If $((z'', c'), c'') \in \tau$ then $(c, \bullet) \in \text{fun}(\tau)(x)$. Otherwise $((z'', c'), \bullet) \in \sigma$ and so $(c', \bullet) \in \text{fun}(\sigma)(y)$, and so $(c, \bullet) \in \text{fun}(\sigma)(y)$ as required. \square

Lemma 6.11. *For $f, g : D(A) \rightarrow D(B)$, if $f \leq_S g$ then $\text{strat}(f) \leq_S \text{strat}(g)$.*

Proof. Suppose $((x, c), a) \in \text{strat}(f)$. If a is a value in B then $(c, a) \in f(x) \leq_S g(x)$ and so either $((x, c), a) \in \text{strat}(g)$ or $((x, c), \bullet) \in \text{strat}(f)$. If a is a cell in A , so that $(c, \bullet) \in f(x + (a, \bullet)) = f(x) \cup g(x + (a, \bullet))$ (since $f \leq_S g$), then either $((x, c), \bullet) \in f$ or $((x, c), a) \in g$ as required. \square

Hence we have shown that:

Theorem 6.12. $[D(A), D(B)] \cong D(A \Rightarrow B)$ in the category of biorders and monotone stable functions.

Corollary 6.13. The category of ordered concrete data structures and monotone stable functions is Cartesian closed.

6.2. Continuous Functions and Sequential Algorithms. We now consider stably continuous and continuously stable functions, showing that these correspond to sequential algorithms which satisfy (respectively) the finite-branching and finite safety conditions. The more interesting case is the former one. We define an internal hom $A \Rightarrow_* B$ for \mathcal{OC}_S by restricting $A \Rightarrow B$ to cells (x, c) such that x is a total state in $D_*(A)$ and defining $\text{fun}_* : D_*(A \Rightarrow_* B) \rightarrow [D_*(A), D_*(B)]$ to be the restriction of fun to states of $D_*(A \Rightarrow_* B)$ and $D_*(A)$ — i.e. $\text{fun}_*(\sigma)(x) =$

$$\{(c, a) \in E(B)_\bullet \mid \exists x' \subseteq x. ((x', c), a) \in \sigma \vee \exists c'. (c', \bullet) \in x \wedge ((x', c), c') \in \sigma\}$$

Proof that this is well-defined and monotone stable follows that for fun (it is also straightforward to check that if σ and x are finite-branching then $\text{fun}_*(\sigma)(x)$ is finite-branching.) So we need to show that for any finite-branching σ , $\text{fun}_*(\sigma)$ is \leq_S -continuous — i.e. if $\Delta \subseteq D_*(A)$ is stably directed then $\text{fun}_*(\sigma)(\bigvee \Delta) = \bigvee \{\text{fun}_*(\sigma)(x) \mid x \in \Delta\}$. Since $\bigvee_{x \in \Delta} \text{fun}_*(\sigma)(x) \leq_S \text{fun}_*(\sigma)(\bigvee \Delta)$ by stability of f , it suffices to show that for any cell c in B , if $(c, \bullet) \in \bigvee_{x \in \Delta} \text{fun}_*(\sigma)(x)$ then $(c, \bullet) \in \text{fun}_*(\sigma)(\bigvee \Delta)$.

Lemma 6.14. For all n , either $(c, \bullet) \in \text{fun}_*(\sigma)(\bigvee \Delta)$ or there exists a set of events $\{((y_1, c), c_1), \dots, ((y_n, c), c_n)\} \subseteq \sigma$ such that $y_i \subseteq \bigvee \Delta$ and $c_i = c_j$ implies $i = j$ for each i .

Proof. By induction on n . Suppose we have $\{((y_1, c), c_1), \dots, ((y_n, c), c_n)\} \subseteq \sigma$ satisfying the above conditions. If any of the c_i is filled with \bullet in $\bigvee \Delta$, then $(c, \bullet) \in f(\bigvee \Delta)$, so suppose that for each $i \in \{1, \dots, n\}$, $(c_i, \bullet) \notin \bigvee \Delta$ and so there exists $x_i \in \Delta$ such that $(c_i, \bullet) \notin x_i$. By directedness of Δ , there exists $x \in \Delta$ such that none of the c_i are filled with \bullet in x . Since $(c, \bullet) \in \text{fun}_*(\sigma)(x)$, there exists an event $((y, c), a) \in \sigma$ such that $y \subseteq x$, and either $a = \bullet$ or $a = c'$, where $y + (c', \bullet) \subseteq x$. If $y \subseteq \bigvee \Delta$, then in the first case $(c, \bullet) \in \text{fun}_*(\sigma)(\bigvee \Delta)$, and in the second case we may set $y_{n+1} = y$ and $c_{n+1} = c'$, satisfying the induction hypothesis.

If $y \not\subseteq \bigvee \Delta$, let P be a proof of $((y, c), a)$ in σ , and let $((z, d), a)$ be the least element of this proof such that $z \not\subseteq \bigvee \Delta$. Then $((z, d), a)$ is immediately preceded in P by an event $((z', d), d')$ such that $z' \subseteq \bigvee \Delta$ and $z = z' + (d', V)$, for some set of values V . So there exists $v \in V$ such that $(d', v) \notin \bigvee \Delta$ and hence $(d', \bullet) \in x$, since $z \subseteq x \leq_S \bigvee \Delta$. $((z', c), d') \geq (z', d'), d) \in \sigma$ and so we may set $y_{n+1} = z'$ and $c_{n+1} = d'$, satisfying the induction hypothesis. \square

We can now prove that $\text{fun}_*(\sigma)$ is \leq_S -continuous.

Proposition 6.15. $\text{fun}_*(\sigma)(\bigvee \Delta) = \bigvee_{x \in \Delta} \text{fun}_*(\sigma)(x)$.

Proof. By Lemma 6.14, if $(c, \bullet) \in \bigvee_{x \in \Delta} \text{fun}_*(\sigma)(x)$, either $(c, \bullet) \in \text{fun}_*(\sigma)(\bigvee \Delta)$ or there is an infinite set of events $\{((y_i, c), c_i) \in \sigma \mid i \in \omega\}$ such that each $y_i \subseteq \bigvee \Delta$ and $c_i = c_j$ implies $i = j$. In the former case we are done, so assume the latter and consider the tree of proofs of these events in σ . If this is infinite branching, then either there is some $((y_i, c), c_i)$ which enables infinitely many cells (y_j, c) — in which case c_i contains infinitely many values in $\bigvee \Delta$ and so $(c_i, \bullet) \in \bigvee \Delta$ and $(c, \bullet) \in f(\bigvee \Delta)$ as required — or else there is some $y \subseteq \bigvee \Delta$

such that $y_i = y$ for infinitely many i — i.e. (y, c) contains infinitely many c_i in σ , and so $((y, c), \bullet) \in \sigma$ and $(c, \bullet) \in f(\bigvee \Delta)$ as required.

If the tree of proofs is finite-branching then by König's lemma it contains an infinite branch — i.e. there exists $z \subseteq \bigvee \Delta$ such that (z, c) has an infinite proof in σ , and by the infinite livelock condition $((z, c), \bullet) \in \sigma$ and $(c, \bullet) \in f(\bigvee \Delta)$ as required. \square

The inverse to fun_* is the restriction of strat to stably continuous functions on finitely branching states — i.e. for stably continuous $f : D_*(A) \rightarrow D_*(B)$, $\text{strat}_*(f) =$

$$\begin{aligned} & \{((x, c), a) \in C(A \Rightarrow_* B) \times V(B)_\perp \mid (c, a) \in f(x)\} \\ & \cup \{((x, c), c') \in C(A \Rightarrow_* B) \times C(A) \mid (c, \bullet) \in f(x + (c', \bullet))\} \end{aligned}$$

Proof that this satisfies the safety condition and furnishes an inverse to fun_* follows the proof of Lemma 6.8. We need to show in addition that if $f : D_*(A) \rightarrow D_*(B)$ is stably continuous then $\text{strat}_*(f)$ is finite-branching.

Lemma 6.16. *$\text{strat}_*(f)$ satisfies the infinite livelock property.*

Proof. Suppose $((x, c), \bullet)$ has an infinite proof $(x_\alpha, c_\alpha)_{\alpha \leq \kappa}$ in $\text{strat}_*(f)$. If the set $\{c_\alpha \mid \alpha \leq \kappa\}$ is infinite then (c, \bullet) has an infinite proof in $f(x)$ and so $(c, \bullet) \in f(x)$ and $((x, c), \bullet) \in \text{strat}_*(f)$ as required.

Otherwise, the proof contains an infinite substring $((x_{\alpha+i}, c_{\alpha+i}), a_{\alpha+i})_{i < \omega}$, where $c_{\alpha+i} = c_\alpha$ and $a_{\alpha+i}$ is a cell in A enabled by $x_{\alpha+i}$ for each $i < \omega$. Let $y_i = x_{\alpha+i} + \{(a_{\alpha+j}, \bullet) \mid j > i \wedge a_j \in A(x_{\alpha+i})\}$ for each $i \in \omega$. The y_i form a stable ω -chain such that $(c_\alpha, \bullet) \in f(y_i)$ for each $i \in \omega$ and hence by stable continuity $(c, \bullet) \geq (c_\alpha, \bullet) \in \bigvee_{i \in \omega} f(y_i) = f(\bigcap_{i \in \omega} y_i) \subseteq f(x)$ and so $((x, c), \bullet) \in \text{strat}(f)$ as required. \square

Lemma 6.17. *$\text{strat}_*(f)$ is finite-branching.*

Proof. Suppose (x, c) is an enabled cell in $A \Rightarrow_* B$ such that there are infinitely many values a such that $((x, c), a) \in \text{strat}_*(f)$. If there are infinitely many values of B such that $((x, c), v) \in \text{strat}_*(f)$ then $(c, \bullet) \in f(x)$ by the finite branching condition on $f(x)$, and so $((x, c), \bullet) \in \text{strat}_*(f)$ as required.

Otherwise there are infinitely many cells c' of A such that $((x, c), c') \in \text{strat}(f)$. Let Y be the set of states y which extend x by filling all but a finite number of its accessible cells with \bullet — i.e. such that $y \leq_S x$ and $A(x) - F(y)$ is finite. Y is \leq_S -directed — if $y, y' \in Y$ then $y \cap y' \in Y$ — and has supremum x .

For any $y \in Y$ there exists a cell $c' \in A(x)$ such that $(c', \bullet) \in y$ and $((x, c), c') \in \text{strat}(f)$ (since there are infinitely many such cells, but only finitely many are not in x). By definition $(c, \bullet) \in f(x + (c', \bullet)) \subseteq f(y)$. Hence $(c, \bullet) \in f(y)$ for all $y \in Y$. and thus $(c, \bullet) \in f(x)$ by continuity, and so $((x', c), \bullet) \in \text{strat}_*(f)$ as required. \square

Finally, we observe that $A \Rightarrow_* B$ is also the internal hom in the category of ocds and continuously stable functions, based on the following observations:

Lemma 6.18. *For any ocds A , $D^*(A)$ is algebraic.*

Proof. Let $\mathcal{K}(A)$ be the set of finitely generated, finitely safe states of A — i.e. $x \in \mathcal{K}(A)$ if and only if it is the upper set of the union of a finite set of proofs. It is straightforward to show that every such state is \subseteq -compact and that every $y \in D^*(A)$ is the union of $\mathcal{K}(y) = \{x \in \mathcal{K}(A) \mid x \subseteq y\}$. \square

Thus $f : D^*(A) \rightarrow D^*(B)$ is continuous if and only if for any $y \in D^*(A)$ and $e \in f(y)$ there exists $x \in \mathcal{K}(y)$ such that $(c, a) \in f(x)$, and we may define an internal hom $(A \Rightarrow_* B)$

by restriction of $A \Rightarrow B$ to cells (x, c) such that x is finitely generated, and $\text{fun}^* : D^*(A \Rightarrow_* B) \rightarrow [D^*(A), D^*(B)]_{\mathcal{B}_\varepsilon}$ by restriction of fun to finitely safe states — i.e. $\text{fun}^*(\sigma)(x) =$

$$\{(c, a) \in E(B)_\bullet \mid \exists x' \in \mathcal{K}(x). ((x', c), a) \in \sigma \vee \exists c'. (c', \bullet) \in x \wedge ((x', c), c') \in \sigma\}$$

for which stability follows as for fun , and continuity follows from proposition 6.18. Its inverse, $\text{strat}^* : [D^*(A), D^*(B)]_{\mathcal{B}_\varepsilon} \rightarrow D^*(A \Rightarrow_* B)$ is similarly defined by restriction of strat to events $((x, c), a)$ such that x is finitely generated, which will have a finite proof in $\text{strat}^*(f)$ provided (c, a) has a finite proof in $f(x)$. The composition of strat^* with fun^* in each direction is the identity on all compact elements, and is therefore the identity on all elements. Finally, we may show (by induction on proof-length) that:

Lemma 6.19. $((x, c), a) \in E(A \Rightarrow_* B)$ has a finite proof if and only if $x \in \mathcal{K}(A)$ and c has a finite proof in B .

Hence the ocds $A \Rightarrow_* B$ and $A \Rightarrow_* B$ are isomorphic in \mathcal{OC}_ε .

7. FULL ABSTRACTION AND UNIVERSALITY

We now return to the may and must testing semantics of bounded and unbounded $\Lambda^+(N)$. In each case we now have an intensional representation of each program of pointed type $t : P$ as a state $\llbracket t \rrbracket$ of an ordered concrete data structure $\llbracket P \rrbracket$, since the biorder denoted by a type of the form $N \rightarrow T$ is isomorphic to $\llbracket T \rrbracket^\omega$. (Formally, $\Lambda^+(N)$ may be interpreted by applying the *families construction* (biproduct completion) to our categories of ordered concrete data structure — see [Abr98].)

First, our aim is to show that our models of bounded $\Lambda^+(N)$ are fully abstract with respect to may testing and must testing, respectively, and characterize their $\Lambda^+(N)$ -definable elements (a form of *universality* result). Full abstraction for models in which terms denote continuous morphisms may be established by proving that each type-object has a basis of elements which are definable (i.e. denotations of terms). In the case of bounded $\Lambda^+(N)$ one may derive such a basis from the fact — already remarked (Proposition 2.10) — that all elements of finite types are definable in Λ^+ , since every pointed $\Lambda^+(N)$ type T is a limit of a chain of finite products of such types (see [Lai09] or [McC96] for details of such a proof). However, this rather cheap reduction to universality for Λ^+ is not very informative about the definability properties of infinite elements — such as the difference between the may and must testing models. We shall prove a stronger result extending Proposition 2.10 — that the type $U \triangleq N \rightarrow o \rightarrow o$ is *universal* in our models of bounded $\Lambda^+(N)$ — i.e. $T \trianglelefteq U$ for all pointed $\Lambda^+(N)$ types T . In particular $U \rightarrow U \trianglelefteq U$ — i.e. U denotes a *reflexive object* in the categories of concrete data structures/complete biorders and continuous functions.

Remark 7.1. This is similar to results for the sequential algorithms model of deterministic SPCF [Lai07a], which are refined in [Lai07b] to show that all retractions are definable in a linearly typed version of SPCF without nesting of function calls (and fixed point combinators replaced with iteration). The retractions given here can be defined within a similar typing system for non-deterministic SPCF

Remark 7.2. We have already observed that $\llbracket U \rrbracket \cong D^*(\mathcal{U})$, where \mathcal{U} is the ocds with a countable set of initial cells $\{c.i \mid i \in \omega\}$ which may each be filled by a single value, and that this is a universal object in the categories of countable ocds and stably continuous/continuously stable functions. Thus every pointed type of $\Lambda^+(N)$ denotes a retract of \mathcal{U} . It remains to show that each such a retraction is definable.

The proofs that there are $\Lambda^+(N)$ -definable retractions from pointed $\Lambda^+(N)$ types into U are similar for the may testing and must testing models — we give the former (slightly simpler) case first.

Observe that $N \rightarrow T$ denotes a \mathbb{N} -indexed product (or infinite list) of copies of $\llbracket T \rrbracket$. Thus we may define the following list operations on terms:

Given $t_1 : T$ and $t_2 : N \rightarrow T$, let $t_1 :: t_2 : N \rightarrow T \triangleq \lambda u. (\mathbf{Eq}(u, 0) t_1) (\mathbf{pred}(u) t_2)$ and given $t : N \rightarrow T$, define $\mathbf{hd}(t) : T \triangleq t 0$ and $\mathbf{tl}(t) \triangleq \lambda v. t \mathbf{succ}(v)$, so that $\llbracket \mathbf{hd}(t_1 :: t_2) \rrbracket = t_1$ and $\llbracket \mathbf{tl}(t_1 :: t_2) \rrbracket = \llbracket t_2 \rrbracket$. Thus, in particular, since $\mathbf{nat} \triangleq (N \rightarrow o) \rightarrow o$

Lemma 7.3. $o \rightarrow \mathbf{nat} \trianglelefteq \mathbf{nat}$.

For any type T , define the pairing and projections:

$\mathbf{pair} : (N \rightarrow T) \rightarrow (N \rightarrow T) \rightarrow N \rightarrow T \triangleq \mathbf{Y} \lambda F. \lambda x. \lambda y. \mathbf{hd}(x) :: (F y \mathbf{tl}(x))$, with $\mathbf{fst} : (N \rightarrow o) \rightarrow (N \rightarrow o) \triangleq \mathbf{Y} \lambda F. \lambda x. (\mathbf{hd}(x) :: (F \mathbf{tl}(\mathbf{tl}(x))))$ and $\mathbf{snd} : (N \rightarrow o) \rightarrow (N \rightarrow o) \triangleq \mathbf{Y} \lambda F. \lambda x. \mathbf{hd}(\mathbf{tl}(x)) :: (F \mathbf{tl}(\mathbf{tl}(x)))$, so that $\llbracket \mathbf{fst} \mathbf{pair}(t_1, t_2) \rrbracket = \llbracket t_1 \rrbracket$ and $\llbracket \mathbf{snd} \mathbf{pair}(t_1, t_2) \rrbracket = \llbracket t_2 \rrbracket$.

Lemma 7.4. For any type T , $N \rightarrow N \rightarrow T \trianglelefteq N \rightarrow T$.

Proof. Let:

$\mathbf{merge} : (N \rightarrow N \rightarrow T) \rightarrow (N \rightarrow T) \triangleq \lambda f. (\mathbf{Y} \lambda F. \lambda u. \mathbf{pair}((f u), (F f) \mathbf{succ}(u))) 0$,
 $\mathbf{unmerge} : (N \rightarrow T) \rightarrow N \rightarrow N \rightarrow T \triangleq (\mathbf{Y} \lambda F. \lambda u. \lambda x. \lambda v. (\mathbf{Eq}(u, v) \mathbf{fst}(x)) ((F \mathbf{succ}(u)) \mathbf{snd}(x))) 0$. \square

So in particular, $N \rightarrow U \trianglelefteq U$. Recall that the denotation of the type \mathbf{nat} is isomorphic to the biorder $\mathcal{P}(\mathbb{N})_{\uparrow}$.

Lemma 7.5. $\mathbf{nat} \trianglelefteq U$.

Proof. Observe that we may faithfully represent any subset $X \subseteq \mathbb{N}$ as a function $b \in \llbracket U \rrbracket$ such that $b(i) = I$ if $i \in X$ and $b(i) = \perp_E$ otherwise (where $I \in [\Sigma, \Sigma]$ is the identity function). This corresponds to the following $\Lambda^+(N)$ -definable retraction:

$$\mathbf{inj} : \mathbf{nat} \rightarrow U \triangleq \lambda f. \lambda u. \lambda x. f \lambda v. (\mathbf{Eq}(u, v) x) \Omega$$

$$\mathbf{proj} : U \rightarrow \mathbf{nat} \triangleq \lambda f. \lambda g. ?_{\mathbf{nat}} \lambda w. (f w) (g w)$$

For any $g : \tilde{\mathbb{N}} \rightarrow \Sigma$, the set $\{\top_S[g(n)]_n \mid n \in \mathbb{N}\}$ is stably bounded above by the constantly \top_S function, and $g = \bigwedge \{\top_S[g(n)]_n \mid n \in \mathbb{N}\}$. Thus $(\mathbf{proj} \mathbf{inj}(f))(g) = \bigwedge_{n \in \mathbb{N}} f(\top_S[g(n)]_n) = f(\bigwedge_{n \in \mathbb{N}} \top_S[g(n)]_n) = f(g)$ by stability. \square

The key retraction reducing the *order* of types (depth to which the arrow is nested on the left) is based on the observation that we may represent a continuously stable function f from $\llbracket U \rrbracket \cong [\Sigma, \Sigma]^\omega$ into Σ via the following information:

- Its set of sequentiality indices,
- For each sequentiality index i , the function $\lambda x. f(x[I]_i)$.

Lemma 7.6. For all $g \in \llbracket N \rightarrow o \rightarrow o \rrbracket$, $f(g) = \bigsqcup_{i \in \mathbf{si}(f)} (g(i) f(x[I]_i))$.

Proof. If $g(i) = \perp_S = \top_E$ for some sequentiality index $i \in \mathbf{si}(f)$ then $f(g) = \perp_S = g(i) f(g[I]_i) = \bigsqcup_{i \in \mathbf{si}(f)} (\pi_i(x) f(x[I]_i))$. So suppose that $g(i) \in \{\perp_E, I\}$ for all $i \in \mathbf{si}(f)$.

If $g(i) = \perp_E$ then $g(i) f(g[I]_i) = \perp_E$ and if $g(i) = I$ then $g(i) f(g[I]_i) = I f(g) = f(g)$.

So if $g(i) = \perp_E$ for every sequentiality index i , then $\bigsqcup_{i \in \mathbf{si}(f)} (g(i) f(x[I]_i)) = \perp_E = f(g)$ by

Lemma 3.14.

Otherwise there is some $i \in \text{si}(f)$ such that $g(i) = I$ and so $\bigsqcup_{i \in \text{si}(f)} g(i) f(g[I]_i) = f(g)$ as required. \square

Thus, extracting sequentiality indices for $f, \{\lambda x.f(x[I]_i) \mid i \in \text{si}(f)\}, \{\lambda x.f(x[I]_i[I]_j) \mid i \in \text{si}(f), j \in \text{si}(\lambda x.f(x[I]_i))\}, \dots$ generates a set of finite sequences of indices from which f may be recovered.

Lemma 7.7. $U \rightarrow o \leq N \rightarrow \text{nat}$.

Proof. Let:

$$\begin{aligned} \text{inj} : (U \rightarrow o) \rightarrow N \rightarrow \text{nat} &\triangleq \lambda f. \mathbf{Y}_{N \rightarrow \text{nat}} \lambda F. (\lambda x.f \lambda y.x) :: (\text{merge}(\lambda u.F \lambda x.f x[I]_u)). \\ \text{proj} : (N \rightarrow \text{nat}) \rightarrow U \rightarrow o &\triangleq \mathbf{Y}_U \lambda F. \lambda x. \lambda y. \text{hd}(x) \lambda v. (y v) (F (\text{unmerge}(\text{tl}(x)) v)). \end{aligned}$$

We prove by induction on $k \in \mathbb{N}$ that for all $e \in \llbracket N \rightarrow o \rightarrow o \rrbracket$, if $|\{n \in \text{nat} \mid e(n) = I\}| \leq k$ then for all $f \in \llbracket U \rightarrow o \rrbracket$, $\text{inj}(\text{proj}(f))(e) = f(e)$. For the induction step, unfold the fixed point, so that: $\text{inj}(\text{proj}(f))(e) = \bigsqcup_{i \in \text{si}(f)} e(i) \text{inj}(\text{proj}(\lambda x.f(x[I]_i)))(e)$. Thus, by Lemma 7.6 it is sufficient to show that if $i \in \text{si}(f)$ then $e(i) \text{inj}(\text{proj}(\lambda x.f(x[I]_i)))(e) = e(i) f(e[I]_i)$. This is evident if $e(i) \neq I$ (i.e. $e(i)$ is the constant function \perp_S or \perp_E). If $e(i) = I$ then

$$\begin{aligned} f(e[I]_i) &= \text{inj}(\text{proj}(\lambda x.f(x[I]_i)))(e[\perp_E]_i) \\ &\sqsubseteq_E \text{inj}(\text{proj}(\lambda x.f(x[I]_i)))(e) \\ &\sqsubseteq_E \text{inj}(\text{proj}(\lambda x.f(x[I]_i)))(e[\top_E]_i) = f(e[I]_i) \text{ by induction hypothesis applied to } e[\perp_E]_i \text{ and } e[\top_E]_i. \end{aligned}$$

So $e(i) \text{inj}(\text{proj}(\lambda x.f(x[I]_i)))(e) = e(i) f(e[I]_i)$, completing the induction.

Since any element $e \in \llbracket U \rrbracket$ is the \sqsubseteq_E -supremum of the approximants e_k such that $e_k(i) = e(i)$ if $i \leq k$ and $e_k(i) = \perp_E$ otherwise (to which the above induction applies), for any $f \in \llbracket U \rightarrow o \rrbracket$, $\text{inj}(\text{proj}(f))(e) = f(e)$ by continuity of f . \square

By composing definable retractions we may now show that U is a (definably) reflexive object (i.e. $U \rightarrow U$ is a definable retract of U).

Proposition 7.8. $U \rightarrow U \leq U$.

$$\begin{aligned} \text{Proof.} \text{ We have } U \rightarrow U &= U \rightarrow N \rightarrow o \rightarrow o \cong N \rightarrow o \rightarrow U \rightarrow o \\ &\leq N \rightarrow o \rightarrow N \rightarrow \text{nat} \text{ by Lemma 7.7} \\ &\cong N \rightarrow N \rightarrow o \rightarrow \text{nat} \leq N \rightarrow N \rightarrow \text{nat} \text{ by Lemma 7.3} \\ &\leq N \rightarrow N \rightarrow U \text{ by Lemma 7.5} \\ &\leq N \rightarrow U \leq U \text{ by Lemma 7.4.} \end{aligned} \quad \square$$

7.1. Full Abstraction. Remarking that \mathcal{U} also corresponds to a universal object in the category of countable ocds and continuous functions, we may apply largely the same argument to establish that U is also a reflexive type in our must testing model of bounded $\Lambda^+(N)$: most of the terms used in the proof of Proposition 7.10 also denote retractions in \mathcal{B}_S , and the same proofs apply. The exception is the retraction of nat into U (Lemma 7.5), which requires the use of unbounded nondeterminism. Recall that in our stably continuous model, nat corresponds to a powerdomain of *finite* sets of natural numbers: by using a different representation of this biorder in \mathcal{U} , we may construct an alternative retraction.

Lemma 7.9. $\text{nat} \leq U$ in the must testing model.

Proof. As in Lemma 7.5 we may represent a (finite) set X of integers as an element $b \in \llbracket U \rrbracket$ such that $\pi_i(b) = I$ if $i \in X$ and $\pi_i(b) = \top_E$ otherwise. The problem is that the projection which sends $b \in \llbracket U \rrbracket$ to the set of all natural numbers i such that $b(i) = I$ exhibits unbounded non-determinism (since there may be infinitely many such i) and is therefore not stably continuous. Our solution is to use a second value $c \in \llbracket U \rrbracket$ to represent an upper bound for X — i.e. $c(i) = I$ if there exists $c \in X$ such that $i \leq j$, and $c(i) = \top_E$ otherwise. We can then return all natural numbers i such that $b(i) = I$ and $\forall j \leq i. c(j) = I$ by stepping through values of i in sequence.

Thus we have a definable retraction of \mathbf{nat} into U^2 , given by the terms:

$$\begin{aligned} \mathbf{inj}_1 &: \mathbf{nat} \rightarrow U \triangleq \lambda x : \mathbf{nat}. \lambda u : N. \lambda y : o. x \lambda v. (\mathbf{Eq}(u, v) y) \mathcal{U} \\ \mathbf{inj}_2 &: \mathbf{nat} \rightarrow U \triangleq \lambda x : \mathbf{nat}. \lambda u : N. \lambda y. \mathbf{leq}(u, v) y \mathcal{U}. \\ \mathbf{proj} &: U \rightarrow U \rightarrow \mathbf{nat} \triangleq \lambda y. \lambda z. \lambda k. (\mathbf{Y}_{N \rightarrow o} \lambda F. \lambda v : N. (z v) (((y v) (k v)) + (F \mathbf{suc}(v)))) 0 \end{aligned}$$

Where:

$$\mathbf{leq} : N \rightarrow N \rightarrow \mathbf{bool} \triangleq \lambda u. \lambda v. \lambda x. \lambda y. (\mathbf{Y} \lambda f. \lambda w. (\mathbf{Eq}(u, w) x) ((\mathbf{Eq}(w, v) y) (f \mathbf{suc}(w)))) 0. \quad \square$$

Thus $U \leq U \rightarrow U$ in the must testing model, exactly as in Proposition 7.8, and we may prove by structural induction:

Proposition 7.10. *For every pointed type P , $P \leq U$ in the may testing and must testing semantics of bounded $\Lambda^+(N)$.*

It remains to prove full abstraction for both models. Say that completeness holds at type T if for all closed $s, t : T$, if $s \lesssim^{\text{may}} t$ then $\llbracket s \rrbracket_{\text{may}} \sqsubseteq_E \llbracket t \rrbracket_{\text{may}}$ and if $s \lesssim^{\text{must}} t$ then $\llbracket t \rrbracket_{\text{must}} \sqsubseteq_E \llbracket s \rrbracket_{\text{must}}$.

Lemma 7.11. *Completeness holds at the type $U \triangleq N \rightarrow o \rightarrow o$.*

Proof. Suppose e.g. $s \lesssim^{\text{must}} t$. Then by soundness and adequacy, for any $d \in \mathbb{N}$ and $e \in \{\top_S, \perp_S\}$ we have $(\llbracket s \rrbracket_{\text{must}} d) e = \top_S = \perp_E$ implies $(\llbracket t \rrbracket_{\text{must}} d) e = \top_S = \perp_E$, and so $\llbracket t \rrbracket_{\text{must}} \sqsubseteq_E \llbracket s \rrbracket_{\text{must}}$. \square

Thus completeness holds at all pointed types (it also holds at N , since e.g. if $t \lesssim^{\text{may}} s$ then $((\mathbf{Eq} t) \mathbf{n}) \mathcal{U} \Omega \Downarrow^{\text{may}}$ implies $((\mathbf{Eq} s) \mathbf{n}) \mathcal{U} \Omega \Downarrow^{\text{may}}$, and hence $\llbracket t \rrbracket_{\text{may}} = \llbracket s \rrbracket_{\text{may}}$).

Theorem 7.12. *For all terms s, t , $s \lesssim^{\text{may}} t$ if and only if $\llbracket s \rrbracket_{\text{may}} \sqsubseteq_E \llbracket t \rrbracket_{\text{may}}$ and $s \lesssim^{\text{must}} t$ if and only if $\llbracket s \rrbracket_{\text{must}} \sqsubseteq \llbracket t \rrbracket_{\text{must}}$.*

Remark 7.13. Universality of U has several other applications. For example, it allows us to characterize the $\Lambda^+(N)$ -definable infinitary elements of a (pointed) type-object $\llbracket T \rrbracket$ as those elements $e \in \llbracket T \rrbracket$ such that $\mathbf{inj}(e) : U$ corresponds to a computable (partial) function from \mathbb{N} into $\{I, \top\}$. As \mathcal{U} is a reflexive object, endomorphisms on \mathcal{U} form a C -monoid, providing an interpretation of the untyped λ -calculus, similar to other graph models, but with a weakly sequential, intensional interpretation as sequential algorithms. Since $[\mathcal{U}, \mathcal{U}_\uparrow] \leq \mathcal{U}$ and $[\mathcal{U}, \mathcal{U}]_\uparrow \leq \mathcal{U}$, the C -monoid of endomorphisms on \mathcal{U} also furnishes weakly sequential semantics for the untyped call-by-value and lazy λ -calculi, to which we can apply the argument in [Lai03] to prove full abstraction.

7.2. Unbounded Nondeterminism. The situation in the non-continuous models is different. Although finitary elements of these models are still definable, in the absence of continuity this does not imply full abstraction, which in fact does not hold (contradicting [Lai06]). We focus on the model of must testing (since this requires non-continuity). We have shown by example that some events with infinite proofs are required to distinguish inequivalent terms in this model, but hinted that there are many more which may distinguish denotations of terms which cannot be separated operationally — we now give an example of a stable and monotone function which can be used to distinguish two must equivalent terms, establishing that the model is not fully abstract. Consider the function $h : \llbracket \mathbf{nat} \rrbracket \rightarrow \llbracket o \rightarrow o \rrbracket$ such that:

- $h(x) = \top_E$ if $x = \top_E = \perp_S$ ($\text{si}(x)$ is empty),
- $h(x) = \perp_E$, if $\text{si}(x)$ is finite but non-empty,
- $h(x) = I$, otherwise ($\text{si}(x)$ is infinite).

This is monotone and stable, since $\text{si}(x) \subseteq \text{si}(y)$ and $\text{si}(y)$ finite implies $\text{si}(x)$ finite, and $x \leq_S y$ in $\llbracket \mathbf{nat} \rrbracket$ implies $x = y$ or $x = \perp_S = \top_E$. However, it cannot be computed in $\Lambda^+(N)$, since there is no test which can check for unboundedness and then return, rather than diverging. Note that $g : \llbracket \mathbf{nat} \rrbracket \rightarrow \llbracket o \rrbracket$ such that $g(x) = \perp_E$ if $\text{si}(x)$ is finite, and $g(x) = \top_E$ otherwise, is definable as $\lambda x. (\mathbf{Y}_{N \rightarrow o} \lambda f. \lambda u. x \lambda v. (\mathbf{1eq}(u, v) \mathbf{U}) (f \mathbf{suc}(v))) 0$.

We may prove that h is not definable in $\Lambda^+(N)$ by showing that all definable functions have the following property:

Definition 7.14. A monotone stable function of extensionally complete biorders $f : D \rightarrow E$ is *weakly continuous* if for any \sqsubseteq_E -directed set X , $f(\bigsqcup X) \leq_S \bigsqcup f(X)$.

The function h is not weakly continuous: let X be the \sqsubseteq_E -directed set $\{\bigsqcup_{i \in J} \pi_i : \Sigma^\omega \rightarrow \Sigma \mid J \subseteq_{\text{fin}} \mathbb{N}\}$. Then $\text{si}(\bigsqcup X) = \mathbb{N}$ and so $h(\bigsqcup X) = I \not\leq_S \perp_E = \bigsqcup_{x \in X} h(x)$. However, all $\Lambda^+(N)$ terms denote weakly continuous functions. We prove this using a logical relation — for each type T , we define a set $H(T) \subseteq \llbracket T \rrbracket$ of *hereditarily weakly continuous* elements of $\llbracket T \rrbracket$ as follows:

- $H(N) = \mathbb{N}$ and $H(o) = \Sigma$
- $f \in H(S \rightarrow T)$ if f is weakly continuous and for any $x \in H(S)$, $f(x) \in H(T)$.

A function $f : \llbracket S_1 \rrbracket \times \dots \times \llbracket S_n \rrbracket \rightarrow \llbracket T \rrbracket$ is hereditarily weakly continuous if it is weakly continuous and for any $\langle x_1, \dots, x_n \rangle \in H(S_1) \times \dots \times H(S_n)$, $f(x_1, \dots, x_n) \in H(T)$.

Proposition 7.15. *Every term of $\Lambda^+(N)$ denotes a hereditarily weakly continuous function.*

Proof. A straightforward structural induction. For the fixed point operator, we note that if F is a \leq_S -directed set of weakly continuous functions from D to (bicomplete) E , then $\bigvee F$ is weakly continuous, since for any \sqsubseteq_E -directed set X , $(\bigvee F)(\bigsqcup X) = \bigvee \{f(\bigsqcup X) \mid f \in F\} \leq_S \bigvee \{\bigsqcup_{x \in X} f(x) \mid f \in F\}$. \square

Hence we have established that:

Proposition 7.16. *There is no term $t : \mathbf{nat} \rightarrow o \rightarrow o$ which denotes the function h .*

Moreover, the weak continuity of definable functions may be used to prove that the model is not fully abstract: we now give a $\Lambda^+(N)$ -definable test which distinguishes h from all weakly continuous functions, and therefore from all definable elements of the model. Define $? : N \rightarrow \mathbf{nat} \triangleq \lambda u. \lambda f. ?_N \lambda v. (\mathbf{1eq}(u, v) (f v)) \mathbf{U}$ so that $?(n)$ returns a bounded choice over all values less than n . Clearly, $?(n)_{n \in \omega}$ is a \sqsubseteq_E -chain with $\bigsqcup_{n \in \omega} ?(n) = ?_N$.

Lemma 7.17. *For any weakly continuous $f : \llbracket \mathbf{nat} \rrbracket \rightarrow \llbracket o \rightarrow o \rrbracket$, $f(?_N) \sqcup_{i \in \omega} (f(?_N)(i)(\perp_S)) = f(?_N)(\perp_S)$.*

Proof. If $f(?_N)$ is a constant function, this is evident. Otherwise $f(?_N) = I$. By weak continuity $I = f(?_N) \leq_S \sqcup_{n \in \omega} f(?_N)(n)$, and by extensional monotonicity $\sqcup_{i \in \omega} f(?_N)(i) \sqsubseteq_E f(?_N) = I$, and so $\sqcup_{n \in \omega} f(?_N)(n) = I$. So $f(?_N) \sqcup_{n \in \omega} (f(?_N)(n)(\perp_S)) = I(I(\perp_S)) = \perp_S = f(?_N)(\perp_S)$ as required. \square

We now give $\Lambda^+(N)$ terms of type $(\mathbf{nat} \rightarrow o \rightarrow o) \rightarrow o$ corresponding to this test. Let $t \triangleq \lambda f.(f ?_N) (?_N \lambda u.(f ?(u)) \Omega)$ and $s = \lambda f.(f ?_N) \Omega$.

Lemma 7.18. $\llbracket t \rrbracket_{must} \neq \llbracket s \rrbracket_{must}$.

Proof. Observe that $h(?_N) = I$ but $h(?_N)(n)$ is the constantly \perp_E function for each n , and hence $\llbracket t \rrbracket(h) = h(?_N)(\sqcup_{n \in \omega} h(?_N)(n)(\perp_S)) = \perp_E$ but $\llbracket s \rrbracket(h) = h(?_N)(\perp_S) = \perp_S$. \square

Lemma 7.19. t and s are *must* equivalent in $\Lambda^+(N)$.

Proof. By extensionality of $\Lambda^+(N)$, it suffices to show that if $r : \mathbf{nat} \rightarrow o \rightarrow o$ is any closed term of $\Lambda^+(N)$ then tr is equivalent to sr . By Proposition 7.15 r denotes a weakly continuous function and so $\llbracket tr \rrbracket = \llbracket r \rrbracket (?_N) (\sqcup_{i \in \omega} \llbracket r \rrbracket (?_N)(i)(\perp_S)) = \llbracket r \rrbracket (?_N) (\perp_S) = \llbracket sr \rrbracket$, and so by computational adequacy, $tr \Downarrow_{must} \mathcal{U}$ if and only if $sr \Downarrow_{must} \mathcal{U}$ as required. \square

Thus we have shown that:

Theorem 7.20. *The *must* testing semantics of $\Lambda^+(N)$ in \mathcal{B} is not fully abstract.*

7.3. Further Directions. We have established new extensional and intensional representations of higher-order functional computation with bounded and unbounded nondeterminism, and shown how they can be used to interpret a simple but expressive programming language. Further study of this model may shed light on the ongoing debate over the relevance of the concept of *fairness* [Dji76], or bear relevance to abstract interpretation of functional languages.

In this paper, we have considered only simply typed languages, but a more general domain theoretic approach would extend our results to recursive types. As well as fixed points for morphisms, we may also construct fixed points (minimal invariants) for functors on our categories of (complete) biorders, as coinciding limit/colimits of ordinal chains of embedding projection pairs (in the continuous cases, these will be ω -chains) obtained by standard constructions. We refer to [Lai03] for one particular application — the construction of fully abstract models of call-by-value and lazy call-by-name λ -calculi by solving the domain equations $D \cong [D, D_\dagger]$ and $D \cong [D, D]_\dagger$ (respectively) in the category of stably complete biorders and stably continuous functions (the models described [Lai03] are in the category of bicpos [Ber78], but all go through in \mathcal{B}_S).

Theorem 7.21 ([Lai03]). *The semantics of λ_v in \mathcal{B}_S is fully abstract.*

By our results relating stably continuous functions to sequential algorithms we may therefore construct fully abstract intensional models of λ_v and λ_l which (unlike e.g. the games model of the lazy λ -calculus in [AM95]) is fully abstract without any “intrinsic preorder quotient”. We leave as further work the extension of our semantics of unbounded non-determinism to recursive types.

Further questions suggested by our semantics include:

- Can the notion of weak continuity be completed to give a characterization of the functions and/or sequential algorithms which are definable in unbounded $\Lambda^+(N)$? We have also left open some important questions regarding the algebraicity properties of biorders arising from ocds, which might suggest further constraints on them.
- Can we interpret nondeterministic languages with richer type systems (for example, explicitly controlling the use of effects such as nondeterminism). In particular, can we relate powerdomain constructions on biorders to ordered concrete data structures.
- Can we describe models of unbounded non-determinism which make the “games-like” structure of ordered concrete data structures more explicit — for example, using graph games [HS02] or concurrent games [RW11]? Is there a corresponding linear decomposition of our model?
- Jones et al. [JB08] have described powerful (but necessarily incomplete) methods for deciding properties of terms in weakly sequential languages such as the untyped call-by-value λ -calculus — can we relate these to our intensional and extensional models?

REFERENCES

- [Abr90] S. Abramsky. The lazy λ -calculus. In D. Turner, editor, *Research Topics in Functional Programming*, pages 65–117. Addison Wesley, 1990.
- [Abr98] S. Abramsky and G. McCusker. Call-by-value games. In M. Nielsen and W. Thomas, editors, *Proceedings of CSL '97*, pages 1–17. Springer-Verlag, 1998.
- [AM95] S. Abramsky and G. McCusker. Games and full abstraction for the lazy λ -calculus. In *Proceedings of the Tenth Annual IEEE Symposium on Logic in Computer Science*, pages 234–243. IEEE Computer Society Press, 1995.
- [AP86] K. R. Apt and G. D. Plotkin. Countable nondeterminism and random assignment. *Journal of the ACM*, 33(4):724–767, 1986.
- [BC82] G. Berry and P.-L. Curien. Sequential algorithms on concrete data structures. *Theoretical Computer Science*, 20:265–321, 1982.
- [Ber78] G. Berry. Stable models of typed λ -calculi. In *Proceedings of the 5th International Colloquium on Automata, Languages and Programming*, number 62 in LNCS, pages 72–89. Springer, 1978.
- [CCF94] R. Cartwright, P.-L. Curien, and M. Felleisen. Fully abstract semantics for observably sequential languages. *Information and Computation*, 1994.
- [CF92] R. Cartwright and M. Felleisen. Observable sequentiality and full abstraction. In *Proceedings of POPL '92*, 1992.
- [CWP97] P.-L. Curien, G. Winskel, and G. Plotkin. Bistructures, bidomains and linear logic. In *Milner Festschrift*. MIT Press, 1997.
- [Dji76] E. W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
- [HA80] M. Hennessy and E. Ashcroft. A mathematical semantics for a non-deterministic typed λ -calculus. *Theoretical Computer Science*, 11:227–245, 1980.
- [Har99] R. Harmer. *Games and Full Abstraction for Nondeterministic Languages*. PhD thesis, Imperial College London, 1999.
- [HM98] R. Harmer and G. McCusker. A fully abstract games semantics for finite non-determinism. In *Proceedings of the Fourteenth Annual Symposium on Logic in Computer Science, LICS '99*. IEEE Computer Society Press, 1998.
- [HS02] Martin Hyland and Andrea Schalk. Games on graphs and sequentially realizable functionals. In *Proceedings of LICS '02*. IEEE Press, 2002.
- [JB08] N. D. Jones and N. Bohr. Call-by-value termination in the untyped lambda-calculus. *Logical Methods in Computer Science*, 4(1), 2008.
- [KP93] G. Kahn and G. Plotkin. Concrete domains. *Theoretical Computer Science*, Böhm Festschrift special issue, 1993. First appeared as technical report 338 of INRIA-LABORIA, 1978.
- [Lai03] J. Laird. A fully abstract bidomain model of unary FPC. In *Proceedings of TLCA '03*, number 2701 in LNCS, 2003.

- [Lai05] J. Laird. Locally Boolean domains. *Theoretical Computer Science*, 342:132–148, 2005.
- [Lai06] J. Laird. Bidomains and full abstraction for countable non-determinism. In *Proceedings of FoS-SaCS'06*, number 3921 in LNCS, pages 352–366. Springer, 2006.
- [Lai07a] J. Laird. Bistability: A sequential domain theory. *Logical Methods in Computer Science*, 3, 2007.
- [Lai07b] J. Laird. On the expressiveness of affine programs with non-local control: The elimination of nesting in SPCF. *Fundamenta Informaticae*, 4(4):511–532, 2007.
- [Lai09] J. Laird. Nondeterministic sequential algorithms. In *Proceedings of CSL '09*, number 5771 in LNCS. Springer, 2009.
- [Lai15] J. Laird. Sequential algorithms for unbounded nondeterminism. In *Proceedings of MFPS '15*, number 319 in Electronic notes in Theoretical Computer Science, pages 271–287, 2015.
- [Lev04] P. B. Levy. *Call-By-Push-Value*. Semantic Structures in Computation. Kluwer, 2004.
- [Lev08] P. B. Levy. Infinite trace equivalence. *Annals of Pure and Applied Logic*, 151(2–3):170–198, 2008.
- [Loa98] R. Loader. Unary PCF is decidable. *Theoretical Computer Science*, 206, 1998.
- [Loa00] R. Loader. Finitary PCF is not decidable. *Theoretical Computer Science*, 266(1–2):341–364, 2000.
- [McC96] G. McCusker. *Games and full abstraction for a functional metalanguage with recursive types*. PhD thesis, Imperial College London, 1996. Cambridge University Press.
- [Mil77] R. Milner. Fully abstract models of typed lambda-calculi. *Theoretical Computer Science*, 4:1–22, 1977.
- [Plo77] G. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5:223 – 255, 1977.
- [Plo82] G.D. Plotkin. A powerdomain for countable nondeterminism. In *Proceedings of ICALP '82*, number 140 in LNCS. Springer, 1982.
- [Plo85] G. Plotkin. Lectures on predomains and partial functions, 1985. Notes for a course given at the Center for the study of Language and Information, Stanford.
- [Ros93a] A. W. Roscoe. An alternative order for the failures model. *Journal of Logic and Computation*, 2(5):557–577, 1993.
- [Ros93b] A. W. Roscoe. Unbounded nondeterminism in CSP. *Journal of Logic and Computation*, 3(2):131–172, 1993.
- [RW11] S. Rideau and G. Winskel. Concurrent strategies. In *Proceedings of LICS '11*. IEEE press, 2011.
- [TO15] T. Tsukada and C.-H. L. Ong. Nondeterminism in game semantics via sheaves. In *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2015)*, pages 220–231, 2015.