

Full Abstraction for Set-Based Models of the Symmetric Interaction Combinators

Damiano Mazza¹ and Neil J. Ross²

¹ LIPN, UMR 7030 CNRS-Université Paris 13
Damiano.Mazza@lipn.univ-paris13.fr
<http://www-lipn.univ-paris13.fr/~mazza>

² Department of Mathematics and Statistics, Dalhousie University
neil.jr.ross@dal.ca

Abstract. The symmetric interaction combinators are a model of distributed and deterministic computation based on Lafont's interaction nets, a special form of graph rewriting. The interest of the symmetric interaction combinators lies in their universality, that is, the fact that they may encode all other interaction net systems; for instance, several implementations of the lambda-calculus in the symmetric interaction combinators exist, related to Lamping's sharing graphs for optimal reduction. A certain number of observational equivalences were introduced for this system, by Lafont, Fernandez and Mackie, and the first author. In this paper, we study the problem of full abstraction with respect to one of these equivalences, using a class of very simple denotational models based on pointed sets.

Keywords: Interaction nets, Observational equivalence, Denotational semantics.

1 Introduction

The symmetric interaction combinators. Interaction nets are a model of distributed, deterministic computation introduced by Lafont [7]. By *distributed* we mean that computation, which is based on graph rewriting, is such that elementary rewriting steps may be applied at different places in the graph, in a completely asynchronous way. By *deterministic* we mean that such elementary steps never overlap, so that the order in which they are executed does not matter, and the computation is essentially unique.

In terms of expressiveness, interaction nets are extremely versatile; for instance, Turing machines may be seen as special interaction nets, in which parallelism is absent. Also, the optimal implementation of λ -calculus evaluation uses certain systems of interaction nets, known as *sharing graphs* [9, 10].

Among all systems of interaction nets, the *interaction combinators* [8] are of special interest, because of their simplicity and universality: in spite of being composed of only 3 combinators with 6 rewriting rules, any system of interaction nets may be encoded in the interaction combinators, preserving the parallelism

of computations. Additionally, the λ -calculus under certain reduction strategies (such as head reduction) may be directly represented in the interaction combinators [11].

The symmetric interaction combinators are a variant of the interaction combinators, also introduced by Lafont [8], having essentially the same expressiveness (for instance, the results of [11] may be immediately transported to the symmetric interaction combinators). The advantage of considering this system is that it is even simpler: although there are still 3 combinators and 6 rules, these latter may actually be arranged in just 2 patterns (annihilations and commutations), and the systems lends itself to a simpler study from the point of view of denotational semantics [12].

Contextual observational equivalences. As an abstract programming language, the symmetric interaction combinators are not too far from the λ -calculus: computation is based on a confluent rewriting relation (yielding the analog of β -equivalence), there is a notion of normal form, a relation analogous to η -equivalence, and so on. Additionally, one may also easily formulate the concept of *context*, from which a rich class of *observational equivalences* may be defined.

As a matter of fact, after Morris [14], we have a general way of defining observational equivalences in a language with an internal notion of context: given a set of syntactic objects S , one defines, for any two syntactic objects t, u ,

$$t \simeq_S u \text{ iff, for every context } C, C[t] \in S \text{ iff } C[u] \in S.$$

In languages like the λ -calculus, one usually considers S to be closed under β -equivalence, so that this latter is automatically contained in \simeq_S . Morris himself introduced the first, still widely used, observational equivalence for the λ -calculus, by taking S as the set of normalizable λ -terms; we denote such equivalence by \simeq_n . Following this pattern, other interesting equivalences may be defined; we refer the reader to [3] for a detailed survey.

The presence of contexts allows more generally to consider the notion of *congruence* on the syntax, *i.e.*, an equivalence relation \sim such that $t \sim u$ implies $C[t] \sim C[u]$, for all objects t, u and context C . Obviously, a Morris observational equivalence is a congruence. Given a congruence \sim , it is natural to ask whether and how it relates to a particular observational equivalence \simeq . We say that \sim is an *abstraction* of \simeq when $\simeq \subseteq \sim$; it is *adequate* with respect to \simeq when $\sim \subseteq \simeq$. These two properties are most frequently investigated when \sim is induced by a denotational model; in particular, one seeks a model whose induced congruence enjoys both of them, in which case the model is said to be *fully abstract*.

The observational equivalence \simeq_n on λ -terms has been widely studied and characterized in several different ways: it is the congruence $=_{\mathfrak{B}\mathfrak{T}_\eta}$ induced by equality of η -normal Böhm trees [6]; and it is the congruence induced by equality in Coppo, Dezani-Ciancaglini and Zacchi's filter model defined in [2].

Back to the symmetric interaction combinators. It is fairly natural to attempt reformulating all of the above congruences, and the questions concerning them, in the symmetric interaction combinators. After Lafont's initial work [8], Fernández

and Mackie were the first to formulate a notion of observational equivalence for interaction nets [4]. More recently, the first author [13] introduced a notion of solvability for nets of symmetric interaction combinators, together with the concept of *observable axiom*, corresponding to a “head variable”. With these, comes a notion of *edifice*, which is a sort of infinite normal form for nets, analogous to a Böhm tree. According to these notions, Fernández and Mackie’s observational equivalence turns out to be quite strong: it is not *sensible*, *i.e.*, it does not identify all unsolvable nets. Alternatively, in analogy with the λ -calculus, it is possible to consider the following two congruences on nets: one which we still denote by \simeq_n , the Morris observational equivalence induced by normalizable nets; and $=_{\varepsilon_\eta}$, the congruence induced by equality of η -expanded edifices, analogous to $=_{\mathfrak{B}\mathfrak{T}_\eta}$.

Perhaps surprisingly, it turns out that these two congruences are distinct: in fact, in the symmetric interaction combinators, \simeq_n is not even *semi-sensible*, *i.e.*, it equates a solvable and an unsolvable net; this is due to the purely local nature of reduction in interaction nets, in which diverging computations cannot be erased.

Interestingly, equality of η -expanded edifices does correspond to an observational equivalence on nets, called *finitary axiom-equivalence*, also introduced in [13]. Rephrased in terms of Böhm trees, this would correspond to the Morris equivalence \simeq_{B_f} , where B_f is the set of λ -terms whose Böhm tree is finite. In the λ -calculus, it coincides with \simeq_n (to prove $\simeq_n \subseteq \simeq_{B_f}$, simply observe that $=_{\mathfrak{B}\mathfrak{T}_\eta}$ obviously discriminates between terms in and not in B_f ; to prove $\simeq_{B_f} \subseteq \simeq_n$, consider the contrapositive, and use Böhm’s theorem to find a context discriminating with respect to B_f); in the symmetric interaction combinators, it is a wholly different equivalence.

The contribution of this paper. Our present objective is to study the question of full abstraction for finitary axiom-equivalence in terms of set-based models, which have a more “abstract” flavor than edifices (just like, say, Plotkin’s model $P\omega$ [1] is more “abstract” than Böhm trees). These models, which were introduced by the first author [12], interpret nets as pointed sets, and are based on the notion of *experiment*, first used by Girard for linear logic [5]. The interest of set-based models lies in their simplicity and concreteness: in many cases, the denotational interpretation of a net may be explicitly computed with ease.

Technically, these models are given by an *interaction set*, which is a pointed set X together with two pointed bijections onto $X \oplus X$ (the product of X with itself), satisfying a certain commutation property. We give sufficient conditions for an interaction set to be fully abstract with respect to finitary axiom-equivalence, and we provide a concrete example of such an interaction set, which thus plays the role, in interaction nets, that Coppo, Dezani-Ciancaglini and Zacchi’s filter model [2] plays in the λ -calculus. The conditions we give are all related to some kind of *approximation property*, *i.e.*, the fact that the denotation of a net is entirely determined by the denotation of its approximations.




Fig. 1. A net.

Acknowledgments. This work was partially supported by ANR projects PANDA (09-BLAN-0169-02) and LOGOI (10-BLAN-0213-02).

2 The Symmetric Interaction Combinators

2.1 Nets

Computation in the symmetric interaction combinators (or, henceforth, simply the symmetric combinators) is a special form of graph rewriting. The graphs on which computation is performed are called *nets*, which are composed by *cells* and *wires*.

Each cell carries a *symbol*, and has a number of *ports*, exactly one of which is *principal*, the other being *auxiliary*. There are three symbols: δ , ε , and ζ . We shall use the letters α, β to range over symbols. Cells of type δ or ζ have two auxiliary ports, hence are called *binary*, and are represented by a triangle; the principal port is represented by one of the “tips” of the triangle, while the auxiliary ports are on the opposite edge. Auxiliary ports are numbered: port 1 (resp. 2) is the “left” (resp. “right”) auxiliary port when the cell is represented with its principal port pointing “down”, and the numbering is preserved by rotations. Cells of type ε have no auxiliary port, hence are called *nullary*, and are represented by a circle.

A *wire* has exactly two extremities, which may be connected to the ports of cells. A *loop* is a wire with its extremities connected together. A wire which is not a loop is called *proper*.

A *net* is a finite (possibly empty) collection of cells and wires, such that each port of each cell is attached to the extremity of a wire. Nets will be ranged over by μ, ν . A net may contain proper wires with one or both extremities not connected to any cell; these are called the *free ports* of the net. If a net has n free ports, they are supposed to be numbered by the integers $1, \dots, n$. As an example, the net in Fig. 1 has 11 cells, of which 4 nullary, 1 loop, 16 proper wires, and 7 free ports, assumed to be numbered increasingly from “left” to “right”.

Let us introduce some remarkable nets, which will be useful in the sequel. A *wiring* is a net containing no cell and no loop. Wirings are permutations of free ports; they are ranged over by ω . We shall often use ω also to denote a single wire.

The net with n free ports consisting of n ε cells is denoted by \mathbf{E}_n .


A *tree* is a net defined by induction as follows. A single ε cell is a tree with no leaf, denoted by ε ; a proper wire is a tree with one leaf, denoted by \bullet ; if τ_1, τ_2 are two trees with resp. n_1, n_2 leaves, and if α is a binary symbol, the net obtained by plugging the root of τ_i into the auxiliary port i of an α cell, with $i \in \{1, 2\}$, is a tree with $n_1 + n_2$ leaves, denoted by $\alpha(\tau_1, \tau_2)$. Trees are represented adopting the same graphical notations as cells. We shall avoid possible ambiguities by never using $\delta, \varepsilon, \zeta$ to denote trees, and by using α, β exclusively to range over cell symbols, so that a triangle annotated with α or β will unambiguously represent a single cell.

An *active pair* is a net consisting of two cells whose principal ports are connected by a wire.

It is also useful to define some special sorts of wires in nets. An *axiom* is a proper wire such that none of its extremities is the principal port of a cell. A *cut* is a proper wire connecting two principal ports. An *axiom-cut* is either a loop, or a proper wire connecting the root of a tree to one of its own leaves. We say that a net is *cut-free* if it contains no cuts and no axiom-cuts. Note that cuts are in one-to-one correspondence with active pairs. As an example, consider the net in Fig. 1, in which the reader should find 7 axioms, 2 cuts (or active pairs), and 2 axiom-cuts.

The following result is proved by induction on the number of cells.

Lemma 1 (Shape). *Let ν be a cut-free net with n free ports. Then, for each $1 \leq i \leq n$ there exist a unique tree τ_i , and there exists a unique wiring ω such that the equality below on the left holds. Let μ be a net with n free ports and k cuts and axiom-cuts. Then, there exists a cut-free net ν with $n + 2k$ free ports such that the equality below on the right holds.*



Note that we use rectangles to represent generic nets, including wirings. However, ω will always denote a wiring. Observe that all wires in the wiring ω of the left equality are axioms; in fact, that is the shape of a generic cut-free proof net of multiplicative linear logic [5], with the axiom links in ω and the logical links in τ_1, \dots, τ_n , whence our terminology. Also observe that the net ν of the right equality is unique as soon as μ does not contain axiom-cuts.

We conclude this section with the essential notion of *context*:

Definition 1 (Context, test). *Let μ be a net with n free ports. A context for μ is a net C with at least n free ports. We denote by $C[\mu]$ the application of C to μ , which is the net obtained by plugging the free port i of μ to the free port i of C , with $i \in \{1, \dots, n\}$. A test for μ is a particular context consisting of n trees τ_1, \dots, τ_n such that the root of each τ_i is the free port i .*




Fig. 2. The rules and equations defining β -reduction and η -equivalence. We assume $\alpha \neq \beta$. In the left β -rule, the right member is empty in case $\alpha = \varepsilon$. In the left η -equation, α is binary.

In the sequel, when we use the notation $C[\mu]$ we implicitly assume that C has enough ports so that μ can be plugged into it. Moreover, we shall say that μ' is a *subnet* of μ if there exists C such that $\mu = C[\mu']$.

2.2 Reductions and equivalences

Computation in the symmetric combinators is performed by rewriting active pairs. We define \rightarrow_β as the contextual closure of the rules of Fig. 2, and denote by \rightarrow_β^* its reflexive-transitive closure. Since active pairs are always disjoint, the relation \rightarrow_β trivially satisfies the diamond property, *i.e.*, reduction is strongly confluent. A confluent rewriting relation always induces an equivalence relation (in this case, a congruence): we define $\mu \simeq_\beta \mu'$ iff there is ν such that $\mu \rightarrow_\beta^* \nu$ and $\mu' \rightarrow_\beta^* \nu$. As usual, a net is *normal* when no β -reduction applies to it. Note that cut-free nets are always normal, but normal nets need not be cut-free.

There is also a notion of η -equivalence, first introduced by Lafont [8] and Fernández and Mackie [4]. Unlike in the λ -calculus, it cannot be presented as the symmetrization of a rewriting relation, because the rightmost equation of Fig. 2 when both α, β are binary is intrinsically non orientable. We define \simeq_η as the reflexive, transitive, and contextual closure of the equations of Fig. 2. Then, we define $\beta\eta$ -equivalence as $\simeq_{\beta\eta} = (\simeq_\beta \cup \simeq_\eta)^+$.

The following definitions were introduced by the first author in [13], and are inspired by similar notions in the λ -calculus [1].

Definition 2 (Solvability). A quasi-wire is a net of the following shape:



where μ_0 is any net with no free ports (including the empty net). A net μ is solvable if there exists a test θ such that $\theta[\mu] \rightarrow_\beta^* W$, where W is a quasi-wire. A net is unsolvable if it is not solvable.

We write $\mu \rightarrow_\varepsilon \mu'$ if $\mu = C[\mu_0]$, $\mu' = C[\mathbf{E}_n]$, and μ_0 is an unsolvable net with n free ports different from \mathbf{E}_n . We then define $\rightarrow_{\beta\varepsilon}$ as the union of \rightarrow_β and \rightarrow_ε , and denote by $\rightarrow_{\beta\varepsilon}^*$ its reflexive-transitive closure. The relation $\rightarrow_{\beta\varepsilon}^*$ too may be proved to enjoy the diamond property [13], so we may consider its induced congruence $\simeq_{\beta\varepsilon}$, and set $\simeq_{\beta\eta\varepsilon} = (\simeq_{\beta\varepsilon} \cup \simeq_\eta)^+$ (the transitive closure of union).

The following Morris observational equivalence was also introduced in [13].

Definition 3 (Finitary axiom-equivalence). Let μ be a net with n free ports. We write $\mu \Downarrow$ if μ is $\beta\varepsilon$ -normalizable; otherwise, we write $\mu \Uparrow$. Let μ, μ' be two nets with the same number of free ports. We say that they are finitarily axiom-equivalent, and we write $\mu \cong \mu'$, if, for every context C , $C[\mu] \Downarrow$ iff $C[\mu'] \Downarrow$.

3 Denotational Semantics

Informally, a denotational semantics of a programming language with a notion of evaluation (denoted by \rightarrow_β) and a notion of context is an interpretation $\llbracket \cdot \rrbracket$ of the syntax into some kind of mathematical structure (which might be the syntax itself) which satisfies the following, for all syntactic objects t, u :

Invariance: $t \rightarrow_\beta u$ implies $\llbracket t \rrbracket = \llbracket u \rrbracket$;

Congruence: for every context C , $\llbracket t \rrbracket = \llbracket u \rrbracket$ implies $\llbracket C[t] \rrbracket = \llbracket C[u] \rrbracket$.

By definition, a denotational semantics induces a congruence on the syntax, by setting $t \sim u$ iff $\llbracket t \rrbracket = \llbracket u \rrbracket$. We make the following important remark, concerning the adequacy property. If \simeq_S is the Morris observational equivalence based on a set S , it is easy to see that it is the greatest congruence contained in $(S \times S) \cup (\mathbb{C}S \times \mathbb{C}S)$; therefore, the semantic congruence is adequate w.r.t. \simeq_S iff $\llbracket t \rrbracket = \llbracket u \rrbracket$ implies that either both t, u are in S , or neither is.

3.1 Interaction sets

The first set-based denotational semantics for the symmetric combinators was introduced in [12]. It is based on *pointed sets*, i.e., sets with a distinguished element, which we denote by $\mathbf{0}$. A morphism of pointed sets, or *pointed function*, is a function between the underlying sets which preserves the distinguished element. Pointed sets and their morphisms form a category, which is equivalent to the category of sets and *partial* functions. This category has a zero object (the pointed set $\{\mathbf{0}\}$) and biproducts: given a family of pointed sets $(X_i)_{i \in I}$, their biproduct $\bigoplus_{i \in I} X_i$ is the pointed set whose underlying set is the product of the underlying sets of the family, and whose distinguished element is the everywhere-zero sequence.

Definition 4 (Interaction set). An interaction set is a triple $(X, \langle \cdot, \cdot \rangle, [\cdot, \cdot])$ where X is a pointed set, and $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ are isomorphisms between $X \oplus X$ and X satisfying

$$\langle [x, y], [z, w] \rangle = [\langle x, z \rangle, \langle y, w \rangle],$$

for all $x, y, z, w \in X$. An interaction set is non-trivial if $X \neq \{\mathbf{0}\}$.

Let $\varphi, \psi : X \rightarrow X \oplus X$ be the inverses of $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$, respectively, and let $\pi_1, \pi_2 : X \oplus X \rightarrow X$ be the projections associated with the biproduct. Then, for $i \in \{1, 2\}$, we define $\delta_i = \pi_i \circ \varphi$ and $\zeta_i = \pi_i \circ \psi$. In other words, for all $x \in X$, $\delta_1(x)$ and $\delta_2(x)$ are the unique elements of X such that $\langle \delta_1(x), \delta_2(x) \rangle = x$, and similarly for ζ_i and $[\cdot, \cdot]$.

In the sequel, we shall denote interaction sets by specifying only their underlying pointed set, leaving the bijections implicit.

Every interaction set X induces a denotational semantics of the symmetric combinators. A net with n free ports is interpreted as a pointed subset of $X^n = X \oplus \cdots \oplus X$ (n times), as follows.

Definition 5 (Experiment, interpretation). *Let X be an interaction set, and let μ be a net with n free ports. An experiment on μ in X is a function e from the ports of μ (including the free ports) to X such that:*

- if p, q are two ports connected by a wire, $e(p) = e(q)$;
- if p_1, p_2 are auxiliary ports number 1 and 2 of a δ (resp. ζ) cell whose principal port is q , then $e(q) = \langle e(p_1), e(p_2) \rangle$ (resp. $e(q) = [e(p_1), e(p_2)]$);
- if q is the principal port of a ε cell, then $e(q) = \mathbf{0}$.

If $p_1 < \dots < p_n$ are the free ports of μ , the sequence $(e(p_1), \dots, e(p_n))$ is said to be the result of e . We use the notation $e : \mu$ as a shorthand for “ e is an experiment on μ ” (the interaction set will always be clear from the context), and we denote by $|e|$ the result of e .

The interpretation of μ in X , denoted by $\llbracket \mu \rrbracket$, is the set containing all results of all experiments on μ in X .

Note that $\llbracket \mathbf{E}_n \rrbracket = \{(\mathbf{0}, \dots, \mathbf{0})\}$. The fact that $\llbracket \mu \rrbracket$ is a pointed set is then immediate:

Lemma 2. *For every net μ with n free ports, $\llbracket \mathbf{E}_n \rrbracket \subseteq \llbracket \mu \rrbracket$.*

Proof. The function assigning $\mathbf{0}$ to all ports of μ is always an experiment. \square

Proposition 1. *For all nets μ, μ' , $\mu \simeq_{\beta\eta} \mu'$ implies $\llbracket \mu \rrbracket = \llbracket \mu' \rrbracket$. Moreover, $\llbracket - \rrbracket$ enjoys the congruence property.*


Proof. To prove invariance under reduction, given $\mu \rightarrow_{\beta} \mu'$, it is enough to show that, for every experiment on μ , there is an experiment on μ' yielding the same result, and vice versa. Invariance under η -equivalence is proved in the same way. The congruence property is an immediate consequence of the definition of experiment. Details may be found in [12]. \square

3.2 Edifices

The following notion, introduced in [13], is analogous to that of a head variable in the λ -calculus.

Definition 6 (Observable axiom). *Let μ be a net. An observable axiom of μ is an axiom connecting the leaves of two trees τ_i, τ_j whose respective roots i, j are both free ports of μ . We say that such observable axiom is based at i, j .*

It is perhaps useful to visualize observable axioms. A net μ contains an observable axiom ω iff it is of the following shape:



If $i = j$, then $\tau_i = \tau_j$, and ω connects two leaves of the same tree. Note also that one or both of τ_i, τ_j may be equal to a wire; in particular, a wire whose both extremities are free is an observable axiom.

The following result is proved in [13]. It is analogous to the λ -calculus result stating that solvability is equivalent to having a head normal form.

Proposition 2. *A net μ is solvable iff there exists a net μ' containing an observable axiom such that $\mu \rightarrow_{\beta}^* \mu'$.*

The observable axioms appearing during the reduction of a net may be collected, just like the head variables of a λ -term, to form the analogous of a Böhm tree. We first need to assign a unique identifier to each observable axiom within a net. In what follows, \mathbb{W} denotes the set of finite words over $\{1, 2\}$. We use a, b to range over \mathbb{W} and denote the empty word by ϵ . Concatenation of words is denoted by juxtaposition.

Definition 7 (Address of a leaf). *Let τ be a tree of cells, and let l be a leaf of τ . We define the δ -address and ζ -address of l in τ , denoted by $\text{addr}_{\delta}^{\tau}(l)$ and $\text{addr}_{\zeta}^{\tau}(l)$, respectively, by induction on τ :*

- if $\tau = \bullet$, then $\text{addr}_{\delta}^{\tau}(l) = \text{addr}_{\zeta}^{\tau}(l) = \epsilon$;
- if $\tau = \delta(\tau_1, \tau_2)$, and l belongs to τ_i (with $i \in \{1, 2\}$), then $\text{addr}_{\delta}^{\tau}(l) = i\text{addr}_{\delta}^{\tau_i}(l)$, and $\text{addr}_{\zeta}^{\tau}(l) = \text{addr}_{\zeta}^{\tau_i}(l)$;
- if $\tau = \zeta(\tau_1, \tau_2)$, and l belongs to τ_i (with $i \in \{1, 2\}$), then $\text{addr}_{\delta}^{\tau}(l) = \text{addr}_{\delta}^{\tau_i}(l)$, and $\text{addr}_{\zeta}^{\tau}(l) = i\text{addr}_{\zeta}^{\tau_i}(l)$.

Definition 8 (Arch). *A pillar is an element of $\mathbb{W} \times \mathbb{W} \times \mathbb{N}$, denoted by $(a, b)@i$; an arch is a set containing exactly two pillars, denoted by $(a, b)@i \frown (a', b')@i'$ (which is the same as $(a', b')@i' \frown (a, b)@i$). Arches are ranged over by \mathbf{a} .*

Let ω be an observable axiom of the net μ . By definition, ω connects two leaves l_i, l_j of two trees τ_i, τ_j whose roots i, j are free ports of μ . Then, we represent ω by the arch

$$(\text{addr}_{\delta}^{\tau_i}(l_i), \text{addr}_{\zeta}^{\tau_i}(l_i))@i \frown (\text{addr}_{\delta}^{\tau_j}(l_j), \text{addr}_{\zeta}^{\tau_j}(l_j))@j.$$

Note that different observable axioms are necessarily represented by different arches; even more, two different observable axioms have no pillar in common. This is because each leaf of a tree in a net may only be connected to one extremity of one observable axiom.

Definition 9 (Edifice). Let μ be a net. We denote by $\text{ax}(\mu)$ the set of all arches representing all observable axioms of μ . Then, we define the edifice of μ to be the following set of arches:

$$\mathfrak{E}(\mu) = \bigcup_{\mu \rightarrow_{\beta}^* \mu'} \text{ax}(\mu').$$

The η -expanded edifice of μ is defined by

$$\mathfrak{E}_{\eta}(\mu) = \{(ac, bd)@i \frown (a'c, b'd)@i' \mid (a, b)@i \frown (a', b')@i' \in \mathfrak{E}(\mu), c, d \in \mathbb{W}\}.$$

In other words, $\mathfrak{E}(\mu)$ is the set of all arches representing all observable axioms appearing during the reduction of μ , and $\mathfrak{E}_{\eta}(\mu)$ is obtained by “ η -expanding” the arches in all possible ways.

As mentioned above, edifices are to nets what Böhm trees are to λ -terms. Indeed, they too yield a denotational semantics, whose induced congruence contains $\simeq_{\beta\eta\varepsilon}$ and is fully abstract with respect to finitary axiom-equivalence:

Proposition 3 (Full abstraction for edifices). For all nets $\mu, \mu', \mu \cong \mu'$ iff $\mathfrak{E}_{\eta}(\mu) = \mathfrak{E}_{\eta}(\mu')$.

Proof. The result is non-trivial; we refer the reader to [13]. \square

Finally, the following result justifies the name given to \cong : $\mu \Downarrow$ means that μ generates a finite number of observable axioms.

Proposition 4. For all μ , $\mu \Downarrow$ iff $\mathfrak{E}(\mu)$ is finite.

Proof. Suppose $\mu \Downarrow$; by definition $\mu \rightarrow_{\beta\varepsilon}^* \nu$ with ν $\beta\varepsilon$ -normal. Now, it is easy to verify that $\beta\varepsilon$ -normal nets are all cut-free, hence the result follows immediately from Lemma 1. For the converse, $\mathfrak{E}(\mu)$ finite means that all observable axioms are generated in a finite number of steps, i.e., $\mu \rightarrow_{\beta}^* C[\mu_0]$, with μ_0 unsolvable by Proposition 2; but then μ is $\beta\varepsilon$ -normalizable to $C[\mathbf{E}_k]$. \square

4 Approximations and Full Abstraction

Definition 10 (Approximation). An approximation of a net μ is a cut-free net ν of the form $C[\mathbf{E}_k]$ such that $\mu \rightarrow_{\beta}^* C[\mu_0]$ for some net μ_0 . In that case, we write $\nu \sqsubseteq \mu$, and we denote by $\text{apx}(\mu)$ the set of all approximations of μ .

Note that $\text{apx}(\mu)$ is never empty; in fact, if μ has n free ports, we always have $\mathbf{E}_n \sqsubseteq \mu$. Moreover, it is not hard to see that $\text{apx}(\mu)$ is a directed set w.r.t. \sqsubseteq . It is then natural to look for interaction sets X in which the interpretation enjoys the following *approximation property*, inspired by algebraicity in domain theory: for every net μ , we ask

$$\llbracket \mu \rrbracket = \bigcup_{\nu \sqsubseteq \mu} \llbracket \nu \rrbracket,$$

that is, the interpretation of μ in X is the supremum of the interpretations of its approximations.

The congruence induced on nets by an interaction set X , which we denote by $=_X$, may be “located” quite precisely in case X satisfies the approximation property (we mention this result without proof, as we shall not need it in the sequel). Indeed, if we let \simeq denote the Morris equivalence induced by unsolvable nets (*i.e.*, $\mu \simeq \mu'$ iff, for every C , $C[\mu], C[\mu']$ are either both solvable, or both unsolvable), which is the greatest semi-sensible theory, we have:

Proposition 5. *Let X be an interaction set enjoying the approximation property. Then, $\cong \subseteq =_X \subseteq \simeq$.*

Observe that the approximation property is trivially enjoyed by edifices (it is an immediate consequence of the definition). The key to full abstraction will be another approximation property, this time connected to the way the set-based semantics “sees” observable axioms.

If X is an interaction set, recall the “projections” $\delta_1, \delta_2, \zeta_1, \zeta_2 : X \rightarrow X$ introduced in Definition 4. If $w = i_1 \cdots i_n \in \mathbb{W}$, with $n > 0$, and if $\alpha \in \{\delta, \zeta\}$ and $x \in X$, in the sequel we shall use the notation $\alpha_w(x) = \alpha_{i_1}(\dots \alpha_{i_n}(x)\dots)$.

Definition 11 (Semantic axioms). *Let X be an interaction set; we denote the diagonal of X^2 by Δ_X . Let $z \in X^n$ with $n > 0$; we denote by z_j the j -th component of z . Given the arch $\mathbf{a} = (a, b)@i \frown (a', b')@i'$, with $1 \leq i, i' \leq n$, we define the projection of z onto \mathbf{a} as*

$$\pi_{\mathbf{a}}(z) = (\delta_a(\zeta_b(z_i)), \delta_{a'}(\zeta_{b'}(z_{i'}))).$$

If $U \subseteq X^n$, we denote by $\pi_{\mathbf{a}}(U)$ the set resulting from the pointwise application of $\pi_{\mathbf{a}}$ to the elements of U . Then, given a net μ , we define

$$\text{sax}(\mu) = \{\mathbf{a} \mid \pi_{\mathbf{a}}(\llbracket \mu \rrbracket) = \Delta_X\}.$$

We say that an interaction set enjoys the *axiom approximation property* if, for every net μ ,

$$\text{sax}(\mu) = \bigcup_{\nu \sqsubseteq \mu} \text{sax}(\nu).$$

Approximation and axiom approximation do not coincide; we know of interaction sets satisfying the first but not the second, and we believe the converse implication to be false as well. However, we shall now establish that an interaction set satisfying both is fully abstract with respect to finitary axiom-equivalence.

Lemma 3. *Let ν be a cut-free net with 2 free ports such that $\llbracket \nu \rrbracket = \Delta_X$, where X is a non-trivial interaction set. Then, ν is η -equivalent to a wire.*

Proof. Observe that any cut-free net $\nu' \simeq_{\eta} \nu$ decomposes, by Lemma 1, into two trees τ_1, τ_2 and a wiring ω . Now, suppose none of these nets (including ν) is such that $\tau_1 = \tau_2 = \tau$ with ω connecting exactly the matching occurrences of the leaves of the two copies of τ (i -th leaf with i -th leaf). Then, since X

is non-trivial, we may easily find an experiment showing that $(x, x') \in \llbracket \nu \rrbracket$ for some $x \neq x'$, a contradiction. Hence, there is an η -equivalent form of ν which is as above; now, if τ contained any ε cell, it is easy to see that we would have $\llbracket \nu \rrbracket \subsetneq \Delta_X$, another contradiction. But a net as the one obtained may be shown to be η -equivalent to a wire by a straightforward induction on τ . \square

Lemma 4. *In a non-trivial interaction set satisfying the axiom approximation property, we have $\text{sax}(\mu) = \mathfrak{E}_\eta(\mu)$, for every net μ .*

Proof. The inclusion $\mathfrak{E}_\eta(\mu) \subseteq \text{sax}(\mu)$ always holds, with no need of the axiom approximation property. In fact, if $\mathfrak{a} \in \mathfrak{E}_\eta(\mu)$ by definition $\mathfrak{a} = (ac, bd)@i \frown (a'c, b'd)@i'$ for some $c, d \in \mathbb{W}$ and with $(a, b)@i \frown (a', b')@i'$ being the arch of an observable axiom ω of a reduct μ' of μ . Then, it is enough to consider experiments on μ' which label every port with $\mathbf{0}$ except those “descending” from ω to see that $\mathfrak{a} \in \text{sax}(\mu)$, by invariance under reduction.

For what concerns the converse, we first consider a cut-free net ν , with n free ports. Let $\mathfrak{a} = (a, b)@i \frown (a', b')@i' \in \text{sax}(\nu)$. It can be shown [13] that, for every sequence of trees τ_1, \dots, τ_n , there exists a cut-free $\nu' \simeq_{\beta\eta} \nu$ such that, for all $1 \leq i \leq n$, we find τ_i rooted at the free port i of ν' . We may then choose $\tau_i, \tau_{i'}$ (the case $i = i'$ changes nothing to the argument) so that two of their resp. leaves l, l' have δ -addresses and ζ -addresses given by a, a' , and b, b' . Now, by invariance under $\simeq_{\beta\eta}$, we still have $\mathfrak{a} \in \text{sax}(\nu')$, which means that every experiment on ν' is forced to assign the same element of X to l, l' , and all elements of X may be assigned to them. But ν' must therefore contain a subnet ν_0 whose free ports coincide with l, l' , and which is detached from the rest of the net, for otherwise, by the non-triviality of X , it would be easy to define an experiment which sets one of l, l' to $\mathbf{0}$ and the other to a non-zero value. Then, we apply Lemma 3 to ν_0 and infer $\mathfrak{a} \in \mathfrak{E}_\eta(\nu') = \mathfrak{E}_\eta(\nu)$. The general statement follows immediately from the axiom approximation property, and from the approximation property for edifices. \square

Lemma 5. *In a non-trivial interaction set satisfying both the approximation and axiom approximation property, we have $\llbracket \mu \rrbracket = \llbracket \mu' \rrbracket$ iff $\text{sax}(\mu) = \text{sax}(\mu')$, for all nets μ, μ' .*

Proof. The implication from left to right is obvious and does not depend on any approximation property. It is then enough to show that $\text{sax}(\mu) \subseteq \text{sax}(\mu')$ implies $\llbracket \mu \rrbracket \subseteq \llbracket \mu' \rrbracket$. By Lemma 4 and the approximation property, this amounts to show that, under the hypothesis $\mathfrak{E}_\eta(\mu) \subseteq \mathfrak{E}_\eta(\mu')$, $\llbracket \nu \rrbracket \subseteq \llbracket \mu' \rrbracket$ for all $\nu \sqsubseteq \mu$. So let $\nu \sqsubseteq \mu$ and $z \in \llbracket \nu \rrbracket$. By Lemma 2, we may suppose $z \neq (\mathbf{0}, \dots, \mathbf{0})$. Then, z is the result of an experiment assigning non-null points to $p > 0$ axioms of ν . These axioms induce p edifices $\mathfrak{A}_1, \dots, \mathfrak{A}_p \subseteq \mathfrak{E}_\eta(\mu) \subseteq \mathfrak{E}_\eta(\mu')$, which means that the “same” axioms (modulo \simeq_η) appear during the reduction of μ' , from which we immediately infer $z \in \llbracket \mu' \rrbracket$, as desired. \square

As announced, combining Proposition 3, Lemma 4, and Lemma 5, we obtain

Theorem 1 (Full abstraction for interaction sets). *If X is a non-trivial interaction set satisfying the approximation and axiom approximation properties, then for all nets $\mu, \mu', \mu \cong \mu'$ iff $\llbracket \mu \rrbracket = \llbracket \mu' \rrbracket$.*

4.1 A fully abstract model

We now proceed to give an example of interaction set satisfying both the approximation and axiom approximation properties.

Let \mathbb{W}^∞ be the set of infinite words over $\{1, 2\}$. We shall consider the pointed set $X = \mathcal{P}_{\text{fin}}(\mathbb{W}^\infty \times \mathbb{W}^\infty)$ of finite subsets of pairs of infinite words, with $\mathbf{0} = \emptyset$. If $x \in X$ and $(a, b) \in \mathbb{W} \times \mathbb{W}$, we set $(a, b) \cdot x = \{(au, bv) \in \mathbb{W}^\infty \times \mathbb{W}^\infty \mid (u, v) \in x\}$, and define the pointed bijections $\langle x, y \rangle = (1, \epsilon) \cdot x \cup (2, \epsilon) \cdot y$ and $[x, y] = (\epsilon, 1) \cdot x \cup (\epsilon, 2) \cdot y$. It is easy to see that $(X, \langle \cdot, \cdot \rangle, [\cdot, \cdot])$ is an interaction set.

In the sequel, all denotational interpretations are assumed to be in X . We shall keep denoting by $\delta_1, \delta_2, \zeta_1, \zeta_2$ the “projections” associated with $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$, and we also use the generalized notations δ_w, ζ_w , with $w \in \mathbb{W}$, introduced just before Definition 11.

Definition 12 (Mean and nice elements). *We say that $x \in X$ is δ -mean (resp. ζ -mean), if the set $\pi_1(x)$ (resp. $\pi_2(x)$), i.e., the set of first (resp. second) projections of the elements of x , is empty or a singleton. If $\alpha \in \{\delta, \zeta\}$, and if $\bar{\alpha}$ is “the other” binary symbol, we introduce the following notations and terminology:*

- M_α is the set of α -mean elements;
- $M = M_\delta \cap M_\zeta$ is the set of mean elements (which are singletons or empty);
- $N = X \setminus (M_\delta \cup M_\zeta)$ is the set of nice elements
- $M_\alpha^* = M_\alpha \setminus M_{\bar{\alpha}}$ is the set of strictly α -mean elements.

In the sequel, if $x \in X$, we shall denote by $\|x\|$ its cardinality. Moreover, we will always have $\alpha \in \{\delta, \zeta\}$, with $\bar{\alpha}$ denoting “the other” binary symbol.

Lemma 6. *For all $x \in X$, the following properties hold:*

Positivity: $\|\mathbf{0}\| = 0$, and $\|x\| \neq 0$ for all $x \neq \mathbf{0}$.

Additivity: $\|\delta_1(x)\| + \|\delta_2(x)\| = \|x\| = \|\zeta_1(x)\| + \|\zeta_2(x)\|$.

Heredity: For all $\beta \in \{\delta_1, \delta_2, \zeta_1, \zeta_2\}$, $x \in M_\alpha$ implies $\beta(x) \in M_\alpha$.

Proof. A simple verification. □

In the following, we write $e : \mu$ to mean “ e is an experiment on μ ”.

Lemma 7. *Let τ be a tree entirely composed of α cells (resp. an arbitrary tree). If $e : \tau$ assigns $x \in M_\alpha$ (resp. $x \in M$) to the root of τ , then e assigns $x' \in M_\alpha$ (resp. $x' \in M$) with $\|x'\| = \|x\|$ to exactly one leaf of τ , and $\mathbf{0}$ to all remaining leaves.*

Proof. By heredity, additivity, and the fact that, if $x \in M_\alpha$ and $i \in \{1, 2\}$, then $\alpha_i(x)$ is either empty, or has the same cardinality as x . □

We say that an experiment assigns an element x to an active pair when it assigns x to both extremities of the cut associated with the active pair.

Lemma 8. *Let μ be a net with $n > 0$ free ports. If $e : \mu$ is such that the labels assigned by e to the active pairs of μ all belong to M then there exists $\nu \sqsubseteq \mu$ and $e_0 : \nu$ such that $|e_0| = |e|$.*

Proof. The idea of the proof is to take each active pair labelled by some $x \in M$, and replace the active pair with a net admitting an experiment of identical result. The case $x = \mathbf{0}$ is easy; otherwise, Lemma 7 is used to extract only two leaves of subtrees of μ on which e assigns non-zero points; then, the net formed by the two trees rooted at the cut is wholly reduced, and the zero-labelled subnets in the result are carefully removed. We omit the technical details. \square

Proposition 6. *X enjoys the approximation property.*

Proof. The right-to-left inclusion holds for all interaction sets, as a corollary of the fact that $\nu \sqsubseteq \mu$ implies $\llbracket \nu \rrbracket \subseteq \llbracket \mu \rrbracket$. Indeed, by definition, $\nu = C[\mathbf{E}_n]$ and $\mu \rightarrow_{\beta}^* C[\mu_0]$. Then, by Lemma 2 and the congruence property, $\llbracket \nu \rrbracket \subseteq \llbracket C[\mu_0] \rrbracket$, so we conclude by invariance.

For the converse, we give the idea of the proof, which is fairly straightforward. We start by defining two functions $\#_{\delta}, \#_{\zeta} : X \rightarrow \mathbb{N}$. Let $x \in X$, and let n_{δ} (resp. n_{ζ}) be the length of the longest common prefix between two different words in $\pi_1(x)$ (resp. $\pi_2(x)$); then, we define $\#_{\alpha}(x)$ to be 0 if $x = \mathbf{0}$, 1 if $x \in M_{\alpha} \setminus \{\mathbf{0}\}$, or $n_{\alpha} + 2$ otherwise. We define the *measure* of x as the integer $m(x) = \|x\| \cdot \#_{\delta}(x) \cdot \#_{\zeta}(x)$. Now, by a simple case inspection, one checks that:

- (i) if $x \in N$, then for all $\beta \in \{\delta_1, \delta_2, \zeta_1, \zeta_2\}$, $m(\beta(x)) < m(x)$;
- (ii) if $x \in M_{\alpha}^*$, then for all $i \in \{1, 2\}$, $m(\bar{\alpha}_i(x)) < m(x)$.

We now show in three steps that by obstinately reducing the active pairs of μ labelled by elements belonging to N first, M_{δ}^* then, and M_{ζ}^* finally, we obtain a net μ' and $e' : \mu'$ such that $\mu \rightarrow_{\beta}^* \mu'$, $|e'| = |e|$, and all the labels assigned to the active pairs of μ' by e' belong to M . The result then follows by Lemma 8.

1. The first step is shown by induction on the multiset $H(\mu)$ (under the usual well-ordering of multisets of integers) containing the measures of all nice elements labelling the active pairs of μ in e , using property (i). At the end of this step we obtain a reduct μ_1 of μ and an experiment $e_1 : \mu_1$ such that $|e_1| = |e|$, and no active pair is labelled by a nice element.
2. The second step is shown by induction on the multiset $J(\mu_1)$ containing the measures of all strictly δ -mean elements labelling the active pairs of μ_1 . Here, we use Lemma 6, Lemma 7 and property (ii). At the end of this step we obtain a reduct μ_2 of μ_1 and an experiment $e_2 : \mu_2$ such that $|e_2| = |e_1|$, and no active pair is labelled by a strictly δ -mean element.
3. The third step is similar to the second.

\square

The fact that X enjoys the axiom approximation property is a consequence of the following

Lemma 9. *Any net μ such that $\llbracket \mu \rrbracket = \Delta_X$ is $\beta\eta$ -equivalent to a wire.*

Proof. It is enough to show that $\text{apx}(\mu)$ contains a net η -equivalent to a wire. Suppose, for the sake of contradiction, that this is not the case. If μ is $\beta\varepsilon$ -normalizable (necessarily to a cut-free form), we may immediately conclude by Lemma 3; therefore, we suppose that μ is not $\beta\varepsilon$ -normalizable. By Proposition 4, $\text{apx}(\mu)$ is infinite; since it is directed w.r.t. \sqsubseteq , we may take an ascending chain $\nu_0 \sqsubseteq \nu_1 \sqsubseteq \dots$ of approximations of μ . By the same arguments given in the proof of Lemma 3, every ν_k decomposes into two copies of a tree τ_k , whose matching occurrences of leaves are joined by axioms forming a wiring ω_k . Now, by hypothesis, every ν_k is not η -equivalent to a wire, which implies that every ν_k contains at least one ε cell. Then, we may build a sequence of finite pairs of words (a_k, b_k) describing a leaf of τ_k where we find an ε cell, such that (a_k, b_k) is a prefix (not necessarily strict) of (a_{k+1}, b_{k+1}) . This increasing sequence defines some $(u, v) \in \mathbb{W}^\infty \times \mathbb{W}^\infty$ such that, by construction, $z = (\{(u, v)\}, \{(u, v)\}) \notin \llbracket \nu_k \rrbracket$ for all $k \in \mathbb{N}$. Then, by the approximation property, $z \notin \llbracket \mu \rrbracket$, in contradiction with $\Delta_X \subseteq \llbracket \mu \rrbracket$. \square

Proposition 7. *X enjoys the axiom approximation property.*

Proof. Let μ be a net, and suppose $\mathbf{a} \in \text{sax}(\mu)$. By the approximation property, using the same argument given in the second part of the proof of Lemma 4, we have that the arch \mathbf{a} defines two leaves l, l' of two trees rooted at two conclusions of a β -reduct μ' of μ , such that l, l' are the free ports of a subnet μ'_0 of μ' whose interpretation is exactly Δ_X . Then, if we apply Lemma 9, reduce μ'_0 , and replace everything else in μ' by ε cells, we obtain an approximation ν of μ such that $\mathbf{a} \in \text{sax}(\nu)$. \square

References

1. Barendregt, H.P.: The Lambda Calculus. North Holland, revised edn. (1984)
2. Coppo, M., Dezani-Ciancaglini, M., Zacchi, M.: Type theories, normal forms, and D_∞ -lambda-models. *Information and Computation* 72(2), 85–116 (1987)
3. Dezani-Ciancaglini, M., Giovannetti, E.: From Bohm's theorem to observational equivalences: an informal account. *Electronic Notes in Theoretical Computer Science* 50(2), 85–118 (2001)
4. Fernández, M., Mackie, I.: Operational equivalence for interaction nets. *Theoretical Computer Science* 297(1–3), 157–181 (2003)
5. Girard, J.Y.: Linear logic. *Theoretical Computer Science* 50(1), 1–102 (1987)
6. Hyland, M.: A syntactic characterization of the equality in some models of the lambda calculus. *Journal of the London Mathematical Society* 2(12), 361–370 (1976)
7. Lafont, Y.: Interaction nets. In: *Conference Record of POPL '90*. pp. 95–108. ACM Press (1990)

8. Lafont, Y.: Interaction combinators. *Information and Computation* 137(1), 69–101 (1997)
9. Lamping, J.: An algorithm for optimal lambda calculus reduction. In: *Conference Record of POPL '90*. pp. 16–30. ACM Press (1990)
10. Mackie, I.: Efficient lambda evaluation with interaction nets. In: *Proceedings of RTA '04*. pp. 155–169. LNCS, Springer (2004)
11. Mackie, I., Pinto, J.S.: Encoding linear logic with interaction combinators. *Information and Computation* 176(2), 153–186 (2002)
12. Mazza, D.: A denotational semantics for the symmetric interaction combinators. *Mathematical Structures in Computer Science* 17(3), 527–562 (2007)
13. Mazza, D.: Observational equivalence and full abstraction in the symmetric interaction combinators. *Logical Methods in Computer Science* 5(4:6) (2009)
14. Morris, J.H.: *Lambda calculus models of programming languages*. Ph.D. Thesis, Massachusetts Institute of Technology (1968)