

Edifices and Full Abstraction for the Symmetric Interaction Combinators

Damiano Mazza

Laboratoire d'Informatique de Paris Nord
damiano.mazza@lipn.univ-paris13.fr
<http://www-lipn.univ-paris13.fr/~mazza>

Abstract. The symmetric interaction combinators are a variant of Lafont's interaction combinators. They are a graph-rewriting model of parallel deterministic computation. We define a notion analogous to that of head normal form in the λ -calculus, and make a semantical study of the corresponding observational equivalence. We associate with each net a compact metric space, called *edifice*, and prove that two nets are observationally equivalent iff they have the same edifice. Edifices may therefore be compared to Böhm trees in infinite η -normal form, or to Nakajima trees, and give a precise topological account of phenomena like infinite η -expansion.

1 Introduction

Lafont's interaction nets [1] are a powerful and versatile model of parallel deterministic computation, derived from the proof-nets of Girard's linear logic [2, 3]. Interaction nets are characterized by the atomicity and locality of their rewriting rules. They can be seen as "parallel Turing machines": computational steps are elementary enough to be considered as executable in constant time, but several steps can be done at the same time.

Several interesting applications of interaction nets exist. The most notable ones are implementations of optimal evaluators for the λ -calculus [4, 5], but efficient evaluation of other functional programming languages using richer data structures is also possible with interaction nets [6].

However, so far the practical aspects of this computational model have arguably received much more attention than the strictly theoretical ones. With the exception of Lafont's work on the interaction combinators [7] and Fernández and Mackie's work on operational equivalence [8], no foundational study of interaction nets can be found in the existing literature. For example, until very recently [9], no denotational semantics had been proposed for interaction nets.

This work aims precisely at studying and expanding the theory of interaction nets, in particular of the symmetric interaction combinators. These latter are particularly interesting because of their *universality*: any interaction net system can be translated in the symmetric interaction combinators [7]. Therefore, a semantical study of the symmetric combinators applies, modulo a translation, to all interaction net systems.

We introduce *observable nets*, which are analogous to head normal forms in the λ -calculus, and we define an observational equivalence based on them. This equivalence is different from the one introduced by Fernández and Mackie: the latter is in fact based on *interface normal forms*, which appear to be related to λ -calculus *weak* head normal forms.

In the λ -calculus, head normal form equivalence (hnf-equivalence) was semantically characterized in the early '70s by the independent results of Wadsworth and Hyland [10, 11]: two λ -terms are hnf-equivalent iff their Böhm tree has the same infinite η -normal form. Shortly after, Nakajima introduced a similar characterization in terms of what are now known as Nakajima trees [12].

In the present work we introduce *edifices*, which play the same rôle as Böhm or Nakajima trees, in that they provide a fully abstract model of the symmetric combinators. Edifices are compact (hence complete) metric spaces, related to Cantor spaces. When nets are interpreted as edifices, phenomena similar to infinite η -expansion, which are also present in the symmetric combinators, receive a precise topological explanation.

Apart from characterizing the interactive behavior of nets, edifices show other interesting aspects, not developed in this paper. They have many common features with the *strategies* of game semantics, and are related to the Geometry of Interaction interpretation of nets [13, 7]. They may be of help in improving the theory of interaction nets, for example by serving as the base for a typed semantics, or by suggesting additive (or non-deterministic) extensions of interaction nets; they may also turn out to be useful in defining alternative models of other systems, like proof-nets, or the λ -calculus itself, as these can all be encoded in the symmetric combinators.

2 The Symmetric Interaction Combinators

2.1 Nets

The symmetric interaction combinators, or, more simply, the symmetric combinators, are an interaction net system [1, 7]. An interaction net is the union of two structures: a labelled, directed hypergraph, and an undirected multigraph:

Definition 1 (Net). A net μ is a triple $(\text{Ports}(\mu), \text{Cells}(\mu), \text{Wires}(\mu))$, where:

- $\text{Ports}(\mu)$ is a finite set, the elements of which are called the ports of μ ;
- $\text{Cells}(\mu)$ is a set of cells, which are tuples of the form $(\alpha, i_0, i_1, \dots, i_n)$, where $\alpha \in \{\delta, \varepsilon, \zeta\}$, and i_0, i_1, \dots, i_n are pairwise distinct ports, such that $n = 2$ if $\alpha = \delta$ or $\alpha = \zeta$, and $n = 0$ if $\alpha = \varepsilon$;
- $\text{Wires}(\mu)$ is a multiset of wires, which are unordered pairs of distinct ports.

$\text{Cells}(\mu)$ and $\text{Wires}(\mu)$ must satisfy the following constraints:

- each port appears in at least one wire;
- each port appears at most twice in $\text{Cells}(\mu) + \text{Wires}(\mu)$ ($\text{Cells}(\mu)$ is considered as a multiset in this union).




Fig. 1. A net (left) and its port graph (right, internal edges dotted).

The ports of μ appearing only once in $\text{Cells}(\mu) + \text{Wires}(\mu)$ are called free; the set of the free ports of μ is referred to as its interface. In a cell $(\alpha, i_0, i_1, \dots, i_n)$, the port i_0 is called principal, and the ports i_1, \dots, i_n are called auxiliary.

Most of the time, it is convenient to present a net graphically, as in Fig. 1. In these representations, only cells and wires are drawn, and ports are left implicit. For a binary cell (i.e., of type δ or ζ), the principal port is represented by one of the “tips” of the triangle representing it. A wire is represented as... a wire, and the free ports appear as extremities of “pending” wires. For example, the net in Fig. 1 has 7 free ports. In the rest of the paper, we shall always assume that if a net has n free ports, then they are labelled by the integers in $\{1, \dots, n\}$. Note also that graphical representations equate nets differing only modulo an injective renaming of ports and a collapse/extension of wires (a sort of α -equivalence).

Each net μ determines an undirected multigraph $\text{PG}(\mu)$, which will be useful to speak of paths in μ (see Fig. 1):

Definition 2 (Port graph). The port graph of a net μ , denoted $\text{PG}(\mu)$, is the undirected multigraph whose vertices are the elements of $\text{Ports}(\mu)$ and such that there is an edge between two ports i, j iff one of the following (non mutually exclusive) conditions hold:

External edges: $\{i, j\} \in \text{Wires}(\mu)$ (multiplicities are counted here, i.e., if $\{i, j\}$ appears twice in $\text{Wires}(\mu)$, there will be two edges relating i and j in $\text{PG}(\mu)$);

Internal edges: i and j are principal and auxiliary ports of a cell of μ .

2.2 Interaction Rules

An *active pair* is a net consisting of two cells whose principal ports are connected by a wire. Active pairs may be reduced according to the *interaction rules* (Fig. 2): the annihilations, concerning the interaction of two cells of the same type, and the commutations, concerning the interaction of two cells of different type.

Reducing an active pair inside a net means removing it and replacing it with the net given by the corresponding rule. If a net μ is transformed into μ' after such an operation, we write $\mu \rightarrow \mu'$. We define $\mu \simeq_{\beta} \nu$ if there exists o such that $\mu \rightarrow^* o$ and $\nu \rightarrow^* o$. It is easy to show that the relation \rightarrow^* is (strongly) confluent, so \simeq_{β} is an equivalence relation (indeed a congruence).

The interest of the symmetric combinatorics is given by the following result:




Fig. 2. The interaction rules: annihilation (left) and commutation (right). In the annihilation, the right member is empty in case $\alpha = \varepsilon$.


Theorem 1 (Lafont [7]). *Any interaction net system can be translated in the symmetric combinators.*

The definitions of interaction net system and of the notion of translation are out of the scope of this paper. We shall only say that, modulo an encoding, Turing machines, cellular automata, and the **SK** combinators are all examples of interaction net systems [7, 9]. An example of encoding of linear logic and the λ -calculus in the symmetric combinators¹ is given by Mackie and Pinto [14]. We refer the reader to Lafont's paper [7] for a proper formulation and proof of Theorem 1.

2.3 Basic Nets

Wirings. A net containing no cell and no cyclic wire is called a *wiring*. Wirings are permutations of free ports; they are ranged over by ω .

Trees. A single ε cell is a tree with no leaf, denoted by ε ; a single wire is a tree with one leaf (it is arbitrary which of the two extremities is the root and which is the leaf), denoted by \bullet ; if τ_1, τ_2 are two trees with resp. n_1, n_2 leaves, and if $\alpha \in \{\delta, \zeta\}$, the net



is a tree with $n_1 + n_2$ leaves, denoted by $\alpha(\tau_1, \tau_2)$.

Trees annihilate in a way which generalizes the annihilation of binary cells:

Lemma 1. *Let τ be a tree. Then, we have*



Proof. By induction on τ . □

¹ Actually these encodings use the interaction combinators, but they can be adapted with very minor changes to the symmetric combinators.




Fig. 3. The equations defining η -equivalence ($\alpha \in \{\delta, \zeta\}$).

3 Observational Equivalence

3.1 Eta Equivalence and Internal Separation

As in the λ -calculus, if reduction is extended by adding other suitable rewriting rules, a result similar to Böhm's theorem can be proved [15].

Definition 3 (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 forest of n trees τ_1, \dots, τ_n such that the root of each τ_i is labelled by i . A test θ is therefore a particular context, and we denote by $\theta[\mu]$ its application to μ .

Definition 4 (η - and $\beta\eta$ -equivalence). η -equivalence is the reflexive, transitive, and contextual closure of the equations of Fig. 3. $\beta\eta$ -equivalence is defined as $\simeq_{\beta\eta} = (\simeq_{\beta} \cup \simeq_{\eta})^+$.

In the following, W and E denote the nets with two free ports consisting resp. of a single wire and of two ε cells.

Theorem 2 (Separation [15]). Let μ, ν be two total² nets with the same interface, such that $\mu \not\simeq_{\beta\eta} \nu$. Then, there exists a test θ such that $\theta[\mu] \rightarrow^* W$ and $\theta[\nu] \rightarrow^* E$, or vice versa.

The following result is the analogous of Lemma 1 for η -equivalence, and will be used in Sect. 5 (like Lemma 1, the proof is a straight-forward induction):

Lemma 2. Let τ be a tree without ε cells. Then, we have

$$\begin{array}{c} \text{---} \langle \tau \rangle \text{---} \\ \vdots \\ \text{---} \langle \tau \rangle \text{---} \end{array} \simeq_{\eta} \text{---} \text{---}$$

² Total means admitting a normal form without vicious circles. A vicious circle is either a cyclic wire, or a configuration consisting of n binary cells c_1, \dots, c_n such that, for all $i \in \{1, \dots, n-1\}$, the principal port of c_i is connected to an auxiliary port of c_{i+1} , and the principal port of c_n is connected to an auxiliary port of c_1 . Such configurations are stable under reduction, because cells can interact only through their principal port. Totality will not be relevant to the main definitions and results of this paper.





Fig. 4. An observable path.

Corollary 1. *For any net ν and for any trees without ε cells τ_1, \dots, τ_n , there exists a net ν' such that*



Proof. Simply “ η -expand” the wires connected to the free ports of ν as in Lemma 2. \square

3.2 Path-based Observational Equivalence

The Separation Theorem distinguishes two nets by sending one to a net presenting a direct connection between its free ports, and the other to a net in which no such direct connection will ever form. This inspires the following definitions.

Definition 5 (Straight path, Danos and Regnier [13]). *Let μ be a net, and $i, j \in \text{Ports}(\mu)$. We say that there is a straight path between i and j in μ iff there is a path (not necessarily simple) connecting i and j in $\text{PG}(\mu)$ and alternating between internal and external edges (see Definition 2). We say that a straight path crosses an active pair iff it contains an edge connecting two principal ports. A maximal path is a non-empty straight path connecting two free ports of μ .*

Definition 6 (Observable path). *Let μ be a net. An observable path of μ is a maximal path crossing no active pair. We denote by $\text{op}(\mu)$ the set of observable paths of μ , and we set*

$$\text{op}^*(\mu) = \bigcup_{\mu \rightarrow^* \mu'} \text{op}(\mu').$$

It is perhaps useful to visualize observable paths. A net μ contains an observable path between its free ports i and j if it is of the shape given in Fig. 4. If $i = j$, then $\tau_1 = \tau_2$, and the wire shown connects two leaves of the same tree. The actual observable path, if seen from i to j , takes the branch of τ_1 leading to the leaf connected by the wire shown, follows this wire, and descends to the root of τ_2 through the only possible branch.

Proposition 1. *Let $\mu \rightarrow^* \mu'$. Then, $\text{op}(\mu) \subseteq \text{op}(\mu')$, and $\text{op}^*(\mu) = \text{op}^*(\mu')$.*

Proof. An immediate consequence of the locality of interaction rules. \square

Note that, for any net μ , $\text{op}(\mu)$ is always finite; then, by Proposition 1, $\text{op}^*(\mu)$ is finite whenever μ has a normal form. The stability of observable paths under reduction is the main reason for considering them as the base of observational equivalence.

Definition 7 (Observability predicates). *We say that μ is immediately observable, and we write $\mu \downarrow$, iff $\text{op}(\mu) \neq \emptyset$. We say that μ is observable, and we write $\mu \Downarrow$, iff $\text{op}^*(\mu) \neq \emptyset$, or, equivalently, $\mu \rightarrow^* \mu' \downarrow$. If $\text{op}^*(\mu) = \emptyset$, we say that μ is blind, and we write $\mu \uparrow$.*

Definition 8 (Observational equivalence). *Two nets μ, ν with the same interface are observationally equivalent, and we write $\mu \simeq \nu$, iff for all contexts C , $C[\mu] \Downarrow$ iff $C[\nu] \Downarrow$.*

It helps thinking of an immediately observable net as a head normal form in the λ -calculus. As a matter of fact, it is possible to extend our definition of observable path to any interaction net system, in particular to sharing graphs [4]. In these latter, observable paths can be seen to be related to *persistent* paths [13]. Then, one can adapt the definition of observable net so as to obtain that a λ -term is in head normal form iff its corresponding net is immediately observable. This adaptation, which we do not detail here, takes into account only the observable paths starting from the free port representing the “root” of the term, and iteratively using the “root” of each subterm.

The existence of a “root” (i.e., a distinguished free port in sharing graphs) is what allows one to define the notion of *principal* head normal form, of which no meaningful equivalent exists for nets. This is because nets, like proof-nets, are “classical”, as opposed to λ -terms, which are “intuitionistic”. This is also the reason why the symmetric combinators equivalent of Böhm trees will not be trees (cf. Sect. 4).

Following the analogy with the λ -calculus, blind nets correspond to unsolvable terms. If we deem *semi-sensible* a congruence on nets including \simeq_β and not equating a blind and an observable net, then it is not hard to show that \simeq is the greatest semi-sensible congruence, just like the corresponding theory \mathcal{H}^* in the λ -calculus.

We also have that, if μ is a blind net with n free ports, then $\mu \simeq E_n$, where E_n is the net consisting of n cells of type ε . Thus, each equivalence class of blind nets (for any interface) has a representative which is normal, in sharp contrast with the λ -calculus. In this respect, one may consider ε cells as the “reification” of unsolvability. Additionally, it can be shown that $\simeq_{\beta\eta}$ is a semi-sensible congruence, so that $\simeq_{\beta\eta} \subseteq \simeq$. Therefore, by Theorem 2, $\simeq_{\beta\eta}$ and \simeq coincide on total nets; in particular, two normal nets without vicious circles are observationally equivalent iff they are $\beta\eta$ -equivalent.³ These results are all consequences of Theorem 3 (Sect. 5), but can also be proved independently.

We conclude by stating an important Context Lemma, saying that tests suffice to discriminate between nets (the proof is technical, and is omitted here):

Lemma 3 (Context). *$\mu \simeq \nu$ iff, for every test θ , $\theta[\mu] \Downarrow$ iff $\theta[\nu] \Downarrow$.*

³ See footnote 2 for the definition of vicious circle and total net.

4 Edifices

We shall now introduce the main mathematical objects of our paper, namely *edifices*. These will be used to develop a denotational semantics for nets, borrowing ideas from the path semantics of linear logic, i.e., Girard's Geometry of Interaction as formulated by Danos and Regnier [13]. Although edifices and Böhm trees are technically quite different, there are strong analogies between the two. Also, the topology used to define edifices is the same used by Kennaway et al. to define the infinitary λ -calculus [16].

In what follows, $\mathcal{C} = \{\mathbf{p}, \mathbf{q}\}^{\mathbb{N}}$ is the set of infinite binary words, ranged over by x, y , equipped with the Cantor topology. We remind that \mathcal{C} is metrizable, with the distance defined for example by $d_{\mathcal{C}}(x, y) = 2^{-k}$, where k is the length of the longest common prefix of x, y . We denote by $\mathcal{B}_{x,r}^{\circ}$ the open ball of center x and radius r . The elements of $\mathcal{C} \times \mathcal{C}$, which is also a Cantor space, will be denoted by $x \otimes y$, and ranged over by u, v, w . Below, the set \mathbb{N} of non-negative integers, ranged over by i, j , will be considered equipped with the discrete topology.

Definition 9 (Pillar). *Given $I \subseteq \mathbb{N}$, set $\mathcal{P}_I = \mathcal{C} \times \mathcal{C} \times I$, equipped with the product topology. A pillar is an element of $\mathcal{P} = \mathcal{P}_{\mathbb{N}}$. Pillars are denoted by $u @ i$, and are ranged over by ξ, v . The pillar $u @ i$ is said to be based at i .*

Observe that \mathcal{P} is also metrizable; if $\xi = x \otimes y @ i$ and $v = x' \otimes y' @ i'$, we shall consider the distance $d(\xi, v) = \max\{d_{\mathcal{C}}(x, y), d_{\mathcal{C}}(x', y'), d_{\text{disc}}(i, i')\}$, where d_{disc} is the discrete metric, defined as $d_{\text{disc}}(i, j) = 0$ if $i = j$, and $d_{\text{disc}}(i, j) = 2$ if $i \neq j$. Therefore, to be "close", two pillars must be based at the same integer.

Definition 10 (Arch). *Given $I \subseteq \mathbb{N}$, pose $\vec{\mathcal{A}}_I = \mathcal{P}_I \times \mathcal{P}_I$, equipped with the product topology, and set $(\xi, v) \sim (\xi', v')$ iff $\xi' = v$ and $v' = \xi$, or $\xi' = \xi$ and $v' = v$. We then define $\mathcal{A}_I = \vec{\mathcal{A}}_I / \sim$, equipped with the quotient topology. An arch is an element of $\mathcal{A} = \mathcal{A}_{\mathbb{N}}$. Arches are denoted by $\xi \frown v$ (which is the same as $v \frown \xi$), and ranged over by \mathbf{a} ; sets of arches are ranged over by \mathfrak{E} . An arch is said to be based at the unordered pair where its two pillars are based.*

The following helps understanding the topology given to \mathcal{A} :

Proposition 2. *The space \mathcal{A} is metrizable; if $\mathbf{a} = \xi \frown v$ and $\mathbf{a}' = \xi' \frown v'$, the function $D(\mathbf{a}, \mathbf{a}') = \min\{\max\{d(\xi, \xi'), d(v, v')\}, \max\{d(\xi, v'), d(v, \xi')\}\}$ is a distance inducing its topology.*

In other words, to compare two arches, we overlap them in both possible ways, and we take the way that "fits best". The distance D is in fact the standard quotient metric; in this case, it collapses to this simple form.

The space \mathcal{A} is not compact. In fact, we can give a characterization of its compact subsets:

Proposition 3. *\mathfrak{E} is compact iff it is a closed subset of \mathcal{A}_I for some finite I .*

Proof. If \mathfrak{E} is compact, then it must be closed; suppose however that $\mathfrak{E} \not\subseteq \mathcal{A}_I$ for any finite I . Then, let $\mathbf{a}_{i,j}$ be a sequence of arches spanning all of the i, j where

the arches of \mathfrak{E} are based, and set $U_{i,j} = \mathfrak{E} \cap \mathcal{B}_{\mathbf{a}_{i,j},2}^\circ$. These are all open sets in the relative topology, and since, for all i,j , $D(\mathbf{a}_{i,j}, \mathbf{a}) < 2$ if \mathbf{a} is based at i,j , they form an open cover of \mathfrak{E} . Now observe that, by the same remark on the distance, if we remove any $U_{m,n}$ we loose all arches of \mathfrak{E} based at m,n . But we have supposed the sequence $\mathbf{a}_{i,j}$ to be infinite, so $U_{i,j}$ is an infinite open cover of \mathfrak{E} admitting no finite subcover, in contradiction with the compactness of \mathfrak{E} .

For the converse, I being finite, it is not hard to show that \mathcal{P}_I is homeomorphic to \mathcal{C} . Therefore, \mathcal{P}_I is a Cantor space, hence compact. So \mathcal{A}_I is compact, because it is the quotient of a product of compact spaces. But a closed subset of a compact space is compact, hence the result. \square

It can be shown that each \mathcal{A}_I is also perfect and totally disconnected, which means that actually these are all Cantor spaces whenever I is finite. What really matters to us though is compactness, which implies completeness (with respect to the metric D of Proposition 2): when I is finite, there is identity between closed, compact, and complete subsets of \mathcal{A}_I .

Definition 11 (Edifice). *An edifice is a compact set of arches.*

5 Nets as Edifices

The basic idea to assign an edifice to a net is that arches model observable paths.⁴ These latter in fact can be seen as unordered pairs of addresses in trees; now, in a pillar $x \otimes y @ i$, any pair of finite prefixes of x, y may be seen as an address, and the base i identifies the tree (a net may have several free ports, and each may be the root of a tree). The need for infinite words arises from η -expansion (the $\alpha\alpha$ equation at left in Fig. 3), which can be applied indefinitely, as in the pure λ -calculus.

In the following, we let a, b range over the set $\{\mathbf{p}, \mathbf{q}\}^*$ of finite binary words, and we denote by $\mathbf{1}$ the empty word. Pairs of finite words are denoted by $a \otimes b$, and ranged over by s, t . The concatenation of two finite words a, b or of a finite word a and an infinite word x are denoted by simple juxtaposition, i.e., as ab and ax respectively. The concatenation of two pairs of finite words $a \otimes b, a' \otimes b'$ or of a pair of finite words $a \otimes b$ and a pair of infinite words $x \otimes y$ are defined resp. as $aa' \otimes bb'$ and $ax \otimes by$, and are also denoted by juxtaposition. If u is a pair of infinite words, when we say that s is a prefix of u we mean that $u = su'$ for some u' , and we always implicitly assume that $s = a \otimes b$ with a, b of equal length, which is also said to be the length of s .

Definition 12 (Address of a leaf). *Let τ be a tree, and l a leaf of τ . The address of l in τ , denoted by $\text{addr}_\tau(l)$, is a pair of finite binary words defined by induction on τ :⁵*

$$- \tau = \bullet: \text{addr}_\tau(l) = \mathbf{1} \otimes \mathbf{1};$$

⁴ Graphically (Fig. 4), observable paths look like arches, hence the terminology.

⁵ For the acquainted reader, $\text{addr}_\tau(l)$ is nothing but the GoI weight of the path going down from l to the root of τ [7]. This justifies our notations for binary words.

- $\tau = \delta(\tau_1, \tau_2)$: $\text{addr}_\tau(l) = (\mathbf{p} \otimes \mathbf{1})\text{addr}_{\tau_1}(l)$ if l is a leaf of τ_1 , $\text{addr}_\tau(l) = (\mathbf{q} \otimes \mathbf{1})\text{addr}_{\tau_2}(l)$ if l is a leaf of τ_2 ;
- $\tau = \zeta(\tau_1, \tau_2)$: $\text{addr}_\tau(l) = (\mathbf{1} \otimes \mathbf{p})\text{addr}_{\tau_1}(l)$ if l is a leaf of τ_1 , $\text{addr}_\tau(l) = (\mathbf{1} \otimes \mathbf{q})\text{addr}_{\tau_2}(l)$ if l is a leaf of τ_2 .

Definition 13 (Edifice of an observable path). Let μ be a net, and let ϕ be an observable path of μ connecting the free ports i and j . By definition, ϕ is completely described by the free ports i, j and the leaves l_i, l_j of the two trees τ_i, τ_j rooted at i, j which are connected in ϕ (cf. Fig. 4). Therefore, if we put $s = \text{addr}_{\tau_i}(l_i)$ and $t = \text{addr}_{\tau_j}(l_j)$, we define

$$\phi^\bullet = \{sw @ i \frown tw @ j ; \forall w \in \mathcal{C} \times \mathcal{C}\}.$$

It is not hard to check that the set defined above is indeed an edifice:

Proposition 4. If μ is a net with n free ports and ϕ an observable path of μ , ϕ^\bullet is a closed subset of $\mathcal{A}_{\{1, \dots, n\}}$.

Definition 14 (Edifice of a net). Let μ be a net. The pre-edifice of μ is the set

$$\mathfrak{E}_0(\mu) = \bigcup_{\phi \in \text{op}^*(\mu)} \phi^\bullet.$$

The edifice of μ is the closure of its pre-edifice: $\mathfrak{E}(\mu) = \overline{\mathfrak{E}_0(\mu)}$.

The soundness of the above definition can be checked as follows: by Proposition 4, all of the ϕ^\bullet are subsets of \mathcal{A}_I for some finite I ; arches based outside I are “too far” to be adherent to $\mathfrak{E}_0(\mu)$, therefore its closure is still in \mathcal{A}_I . By Proposition 3, this is enough to ensure the compactness of $\mathfrak{E}(\mu)$.

Observe that if μ is normalizable, then $\text{op}^*(\mu)$ is finite, hence by Proposition 4 $\mathfrak{E}_0(\mu)$ is already closed. It is however possible to find non-normalizable nets whose pre-edifice is not an edifice (e.g. the net of Fig. 5 discussed below).

The closure is in fact essential for yielding a fully-abstract denotational semantics of nets. It is crucial in the proof of the following result:

Lemma 4. Let μ, ν be two nets with n free ports. Then, $\mathfrak{E}(\mu) \neq \mathfrak{E}(\nu)$ implies that there exist $i, j \in \{1, \dots, n\}$, two pairs of finite words s, t , and two observable paths $\phi \in \text{op}^*(\mu)$ and $\psi \in \text{op}^*(\nu)$ such that, if we put $\mathbf{a}_w = sw @ i \frown tw @ j$, either for all w , we have $\mathbf{a}_w \in \phi^\bullet \setminus \mathfrak{E}(\nu)$, or for all w , we have $\mathbf{a}_w \in \psi^\bullet \setminus \mathfrak{E}(\mu)$.


Proof. Suppose, without loss of generality, that there exists $\mathbf{a} \in \mathfrak{E}(\mu) \setminus \mathfrak{E}(\nu)$, based at $i, j \in \{1, \dots, n\}$. Remember that $\mathfrak{E}(\mu)$ and $\mathfrak{E}(\nu)$ are defined as the closures of resp. $\mathfrak{E}_0(\mu)$ and $\mathfrak{E}_0(\nu)$, and that by Proposition 3 they are both compact, hence complete. Then, if $\mathbf{a} \in \mathfrak{E}(\mu) \setminus \mathfrak{E}_0(\mu)$, \mathbf{a} must be a “missing limit” of a Cauchy sequence $\mathbf{a}_n \in \mathfrak{E}_0(\mu)$. Since a subsequence of a Cauchy sequence is still a Cauchy sequence, there must exist an integer m such that, for all $n \geq m$, $\mathbf{a}_n \in \mathfrak{E}_0(\mu) \setminus \mathfrak{E}(\nu)$, otherwise \mathbf{a} would belong to $\mathfrak{E}(\nu)$ because of its completeness. Therefore, modulo replacing it by one of these \mathbf{a}_n , we can always assume that $\mathbf{a} \in \mathfrak{E}_0(\mu)$. If it is so, then by definition there exists an observable

path $\phi \in \text{op}^*(\mu)$ such that $\mathbf{a} \in \phi^\bullet$, which means that $\mathbf{a} = sw_0 @ i \frown tw_0 @ j$ and, for every $w \in \mathcal{C} \times \mathcal{C}$, $sw @ i \frown tw @ j \in \phi^\bullet$, where s and t are the addresses of two leaves in the reduct(s) of μ in which ϕ appears. Now let s'_1, \dots, s'_n, \dots be a sequence of prefixes of increasing length of w_0 , and set, for all n , $s_n = ss'_n$ and $t_n = ts'_n$. Suppose that, for all n , there exist two pairs of infinite words u_n, v_n such that $\mathbf{a}_n = s_n u_n @ i \frown t_n v_n @ j \in \mathfrak{E}(\nu)$; it is not hard to verify that the arches \mathbf{a}_n would form a Cauchy sequence of limit \mathbf{a} , and thus, by the completeness of $\mathfrak{E}(\nu)$, we would obtain $\mathbf{a} \in \mathfrak{E}(\nu)$, a contradiction. Therefore, there must exist an integer n such that, for all w , $s_n w @ i \frown t_n w @ j \in \phi^\bullet \setminus \mathfrak{E}(\nu)$. \square

Lemma 5. $\mu \simeq_\eta \nu$ and $\mu \rightarrow^* \mu'$ implies that there exist $\mu'' \simeq_\eta \nu''$ such that $\mu' \rightarrow^* \mu''$ and $\nu \rightarrow^* \nu''$.

Proof. Omitted (see [15]). \square

Definition 15 (η -equivalent observable paths). Let $\tau_1, \tau_2, \tau'_1, \tau'_2$ be trees, with $\tau_1 = \tau_2$ iff $\tau'_1 = \tau'_2$, and let ϕ, ϕ' be two observable paths, such that in ϕ there is a connection between two leaves l_1, l_2 of τ_1 and τ_2 , and in ϕ' there is a connection between two leaves l'_1, l'_2 of τ'_1 and τ'_2 . We say that ϕ is η -equivalent to ϕ' iff



Lemma 6. Let $\mu \simeq_\eta \nu$, and let $\phi \in \text{op}^*(\mu)$. Then, there exists $\psi \in \text{op}^*(\nu)$ such that ϕ and ψ are η -equivalent.


Proof (sketch). By definition, $\phi \in \text{op}^*(\mu)$ means that ϕ is an observable path of a reduct μ' of μ . By Lemma 5, $\mu' \rightarrow^* \mu''$ and $\nu \rightarrow^* \nu''$ such that $\mu'' \simeq_\eta \nu''$. But observable paths are preserved under reduction, so ϕ is also present in μ'' . Now if, in rewriting μ'' to ν'' , no active pair is introduced to alter the observability of ϕ , then clearly ν'' contains an observable path η -equivalent to ϕ . Otherwise, it is easy to check that an $\alpha\alpha$ equation must have been used. In this case, one can prove that the active pairs introduced can be reduced to obtain $\nu'' \rightarrow^* o$ such that o contains an observable path ψ η -equivalent to ϕ . But the reducts of ν'' are also reducts of ν , so $\psi \in \text{op}^*(\nu)$. \square

Lemma 7. If $\mu \simeq_\eta \nu$, then $\mathfrak{E}_0(\mu) = \mathfrak{E}_0(\nu)$ (hence $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$).

Proof (sketch). By Lemma 6, it is enough to check that, whenever ϕ and ψ are η -equivalent observable paths, $\phi^\bullet = \psi^\bullet$. The η -equations concerning ε cells need not be considered; in the case of the $\delta\zeta$ equation, the fact that in pillars δ and ζ cells are treated by separate words makes their relative order irrelevant, and thus accounts for their commutation; the $\alpha\alpha$ equations, which in this case may only be applied to the wire connecting the two leaves of an observable path, are modelled by the fact that all possible "uniform completions" of the addresses of the leaves are taken in the edifice of an observable path. \square

We now prove that $\mathfrak{E}(\cdot)$ induces a congruence with respect to tests:

Lemma 8. 1. Let τ be a tree, and let




Then, $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$ iff $\mathfrak{E}(\mu_0) = \mathfrak{E}(\nu_0)$.

2. Let μ, ν be two nets with the same interface such that $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$, and let τ be a tree without ε cells. Then, if we pose



we have $\mathfrak{E}(\mu') = \mathfrak{E}(\nu')$.

3. Let μ, ν be two nets with the same interface such that $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$, and let



Then, $\mathfrak{E}(\mu') = \mathfrak{E}(\nu')$.

Proof. 1. Easy.


2. Simply consider the nets μ'', ν'' obtained from μ', ν' by adding a copy of τ to the one already existing in the two nets, so that each leaf l in one copy is connected to the same leaf l in the other copy. By Lemma 2, we have that $\mu'' \simeq_{\eta} \mu$ and $\nu'' \simeq_{\eta} \nu$; by point 1, we have $\mathfrak{E}(\mu') = \mathfrak{E}(\nu')$ i $\mathfrak{E}(\mu'') = \mathfrak{E}(\nu'')$; but by Lemma 7, and by hypothesis, $\mathfrak{E}(\mu'') = \mathfrak{E}(\mu) = \mathfrak{E}(\nu) = \mathfrak{E}(\nu'')$.
3. Call k the free port of μ to which the ε cell is connected in μ' . Observe that such ε cell can either disappear, or be duplicated, and that, in any case, ε cells cannot be used by observable paths. Hence, $\phi \in \text{op}^*(\mu')$ i $\phi \in \text{op}^*(\mu)$ and ϕ connects two free ports of μ both different than k . Therefore, $\mathfrak{E}(\mu') = \{u @ i \frown u @ j \in \mathfrak{E}(\mu) ; j, k \neq i\}$. The same holds for ν , so from $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$ it easily follows that $\mathfrak{E}(\mu') = \mathfrak{E}(\nu')$. \square

Corollary 2. Let μ, ν be two nets with the same interface, and let θ be a test. Then, $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$ implies $\mathfrak{E}(\theta[\mu]) = \mathfrak{E}(\theta[\nu])$.


To prove full abstraction, we first need the following separation result:

Lemma 9. Let W be a net with two free ports connected by a wire, and let μ be a net with two free ports, such that $\phi \in \text{op}^*(\mu)$ implies that ϕ does not connect the port 1 to the port 2. Then, there exists a test θ such that $\theta[W] \Downarrow$ and $\theta[\mu] \Uparrow$.

Proof. If $\mu \Uparrow$, the identity test succeeds, so suppose $\mu \Downarrow$. By hypothesis, all observable paths appearing in the reducts of μ connect one of the free ports to itself. Therefore, there exists μ' such that $\mu \rightarrow^* \mu'$, and



In the above picture, we have supposed that the observable path connects the free port 1 to itself, and that the leaves connected in the path are the two “leftmost” leaves of τ . These are just graphically convenient assumptions, causing no loss of generality: the observable path may as well connect port 2 to itself, and the leaves connected may be any two leaves of τ . Now, if we define




we have that, thanks to Lemma 1, $\theta[W] \rightarrow^* W$, while $\theta[\mu]$ reduces to a net whose free port 1 is connected to an ε cell. If this net is blind, we are done; otherwise, there is a reduct of $\theta[\mu]$ containing an observable path between the free port 2 and itself. This observable path can be “eliminated” with the same technique, while the ε cell on port 1 will “eat” any tree fed to it, so in the end we obtain a test θ' such that $\theta'[W] \rightarrow^* W \downarrow$, while $\theta'[\mu] \uparrow$, as desired. \square

We are now ready to prove our main result:

Theorem 3 (Full abstraction). $\mu \simeq \nu$ iff $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$.

Proof. Consider first the backward implication (also known as the adequacy property). We start by observing that, for any net o , $o \downarrow$ i $\text{op}^*(o) \neq \emptyset$ i $\mathfrak{E}(o) \neq \emptyset$. Now, suppose $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$, and let θ be a test. By Corollary 2, we have $\mathfrak{E}(\theta[\mu]) = \mathfrak{E}(\theta[\nu])$, so following the above remark $\theta[\mu] \downarrow$ i $\mathfrak{E}(\theta[\mu]) \neq \emptyset$ i $\mathfrak{E}(\theta[\nu]) \neq \emptyset$ i $\theta[\nu] \downarrow$. Then $\mu \simeq \nu$ follows from the Context Lemma 3.

Now we turn to the actual full abstraction property. For this, we consider the contrapositive statement, and assume $\mathfrak{E}(\mu) \neq \mathfrak{E}(\nu)$. Let I be the interface of μ and ν . By Lemma 4, we know that there exist $i, j \in I$, $\phi \in \text{op}^*(\mu)$, and two leaves in a reduct of μ of addresses s, t such that, for all $w, sw @ i \sim tw @ j \in \phi \bullet \setminus \mathfrak{E}(\nu)$ (it could actually be that these arches belong to $\psi \bullet \setminus \mathfrak{E}(\mu)$, where $\psi \in \text{op}^*(\nu)$, but obviously our assumption causes no loss of generality). We shall suppose $i \neq j$; the reader is invited to check that the argument can be adapted to the case $i = j$. By Definition 13, and by the fact that $\phi \in \text{op}^*(\mu)$, we have



where we have explicitly drawn the connection between the two leaves of resp. addresses s and t . On the other hand, by Corollary 1, we have






Fig. 5. A non-normalizable net observationally equivalent to a wire.




where we have called k and l the two free ports of ν' corresponding resp. to the addresses t and s in τ_i and τ_j . Observe that, by Lemma 5, the edifice of the net on the right is still $\mathfrak{E}(\nu)$. Now if, in any reduct of ν' , there appeared an observable path between k and l , then we would contradict the fact that, for all w , $sw @ i \frown tw @ j \notin \mathfrak{E}(\nu)$. Therefore, no observable path ever appears between k and l in any reduct of ν' .

Consider then the test



where we have left free only the leaves corresponding to the addresses s and t of τ_i and τ_j . Now, by Lemma 1, $\theta[\mu] \rightarrow^* W$, where W is a wire plus a net with no interface; on the other hand, we have



But ν' never develops observable paths between k and l , so Lemma 9 applies, and we obtain $\mu \not\sim \nu$. \square

As an immediate application of Theorem 3, we give an example of a net which is not normalizable, and yet is observationally equivalent to a wire; this is analogous to Wadsworth's "infinitely η -expanding" term $J = RR$, where $R = \lambda xzy.z(xxy)$, which is well known to be hnf-equivalent to $\lambda z.z$.

Consider a net ι containing no observable paths, and reducing as in Fig. 5. Such a net exists, although its description is not as concise as that of J . We see that $\phi \in \text{op}^*(\iota)$ i $\phi^\bullet = \{\mathbf{q}^n \mathbf{p}x \otimes y @ 1 \frown \mathbf{q}^n \mathbf{p}x \otimes y @ 2 ; \forall x, y \in \mathcal{C}\}$ for some non-negative integer n . On the other hand, if W denotes a wire, $\mathfrak{E}(W) = \mathfrak{E}_0(W) = \{u @ 1 \frown u @ 2 ; \forall u \in \mathcal{C} \times \mathcal{C}\}$. Now, if \mathbf{q}^∞ denotes an infinite

sequence of \mathbf{q} 's, all arches of the form $\mathbf{a}_y = \mathbf{q}^\infty \otimes y @ 1 \frown \mathbf{q}^\infty \otimes y @ 2$ are missing from $\mathfrak{E}_0(\iota)$, hence $\mathfrak{E}_0(\iota) \subsetneq \mathfrak{E}_0(W)$. But these arches are all adherent to $\mathfrak{E}_0(\iota)$: in fact, it is very easy to construct a Cauchy sequence in $\mathfrak{E}_0(\iota)$ of limit \mathbf{a}_y , for any y . Therefore, $\mathfrak{E}(\iota) = \mathfrak{E}(W)$, and $\iota \simeq W$.

Notice that the reducts of ι are "almost" η -equivalent to a wire: there is just one missing connection. We can say that this connection forms "in the limit", when the reduction is carried on forever. When one interprets nets as edifices, this informal remark becomes a precise topological fact, i.e., we have a true limit.

Compactness is crucial for obtaining full abstraction. Notice in fact that $\mathfrak{E}_0(\cdot)$ already gives an adequate semantics of nets, which however fails to be fully abstract, as the above example itself shows.

References

1. Lafont, Y.: Interaction nets. In: Conference Record of POPL'90, ACM Press (1990) 95–108
2. Girard, J.Y.: Proof-nets: The parallel syntax for proof-theory. In Ursini, Agliano, eds.: Logic and Algebra. Marcel Dekker, Inc. (1996)
3. Lafont, Y.: From proof nets to interaction nets. In Girard, J.Y., Lafont, Y., Regnier, L., eds.: Advances in Linear Logic. Cambridge University Press (1995) 225–247
4. Gonthier, G., Abadi, M., Lévy, J.J.: The geometry of optimal lambda reduction. In: Conference Record of POPL 92, ACM Press (1992) 15–26
5. Mackie, I.: Efficient lambda evaluation with interaction nets. In: Proceedings of RTA '04. LNCS, Springer (2004) 155–169
6. Mackie, I.: An interaction net implementation of additive and multiplicative structures. Journal of Logic and Computation **15**(2) (2005) 219–237
7. Lafont, Y.: Interaction combinators. Information and Computation **137**(1) (1997) 69–101
8. Fernández, M., Mackie, I.: Operational equivalence for interaction nets. Theoretical Computer Science **297**(1–3) (2003) 157–181
9. Mazza, D.: A denotational semantics for the symmetric interaction combinators. To appear in Mathematical Structures in Computer Science **17**(3) (2007)
10. Wadsworth, C.: The relation between computational and denotational properties for Scott's D_∞ models. Siam J. Comput. **5**(3) (1976) 488–521
11. Hyland, M.: A syntactic characterization of the equality in some models of the lambda calculus. J. London Math. Society **2**(12) (1976) 361–370
12. Nakajima, R.: Infinite normal forms for the λ -calculus. In Böhm, C., ed.: Lambda-Calculus and Computer Science Theory. LNCS, Springer (1975) 62–82
13. Danos, V., Regnier, L.: Proof nets and the Hilbert space. In Girard, J.Y., Lafont, Y., Regnier, L., eds.: Advances in Linear Logic, Cambridge University Press (1995) 307–328
14. Mackie, I., Pinto, J.S.: Encoding linear logic with interaction combinators. Information and Computation **176**(2) (2002) 153–186
15. Mazza, D.: Observational equivalence for the interaction combinators and internal separation. In Mackie, I., ed.: Proceedings of TERMGRAPH 2006. ENTCS, Elsevier (2006) 7–16
16. Kennaway, R., Klop, J.W., Sleep, R., de Vries, F.J.: Infinitary lambda calculus. Theoretical Computer Science **175**(1) (1997) 93–125