

A denotational semantics for the symmetric interaction combinators

DAMIANO MAZZA

*Institut de Mathématiques de Luminy (UMR 6206),
Campus de Luminy, Case 907 – 13288 Marseille Cedex 9, France*
Email: mazza@iml.univ-mrs.fr
Website: <http://iml.univ-mrs.fr/~mazza>.

Received 25 April 2006; revised 19 August 2006

The symmetric interaction combinators are a variant of Lafont's interaction combinators. They enjoy a weaker universality property with respect to interaction nets, but are equally expressive. They are a model of deterministic distributed computation and share the good properties of Turing machines (elementary reductions) and of the λ -calculus (higher-order functions and parallel execution). We introduce a denotational semantics for this system, which is inspired by the relational semantics for linear logic, and prove an injectivity and full completeness result for it. We also consider the algebraic semantics defined by Lafont, and prove that the two are strongly related.

1. Introduction

Interaction nets (Lafont 1990) are a model of distributed deterministic computation. Based on graph-rewriting, they can be seen as a generalisation of multiplicative linear logic proof-nets (Girard 1987b; Danos and Regnier 1989; Lafont 1995).

Interaction nets are interesting for several reasons:

- They are highly expressive: Turing machines, cellular automata, and a number of word or term rewriting systems can all be seen as special instances of interaction nets. When we say this we mean more than just the existence of some (possibly impractical) encoding: in many cases, interaction nets are able to 'implement' a computational model preserving its fundamental properties (such as sequentiality, parallelism and complexity). Two interesting examples of such translations (Turing machines and cellular automata) are given in Lafont (1997).
- They have been shown to be the 'true' syntax underlying Girard's Geometry of Interaction (Girard 1989; Gonthier *et al.* 1992). This is potentially of great theoretical interest, since the geometry of interaction is a semantics that attempts to give a mathematical meaning to the *execution* of a program, and not just to its final result; therefore, it may offer a powerful instrument to study the dynamics of computation, such as its complexity.
- From a more applied point of view, they can be seen as a programming paradigm, and can be turned into a practical (typed or untyped) programming language in which important properties (like deadlock-freeness) are checked automatically using similar techniques to those used in linear logic (that is, correctness criteria (Lafont 1990)).

Amongst all interaction net systems, the *interaction combinators* (Lafont 1997) stand out as particularly interesting because they are *universal*, in the sense that all other interaction net systems can be ‘compiled’ into them; again, this compilation process preserves the basic properties of parallelism and complexity of the original system.

As a consequence, the interaction combinators can be seen as a computational model of its own, combining in some sense the good properties of Turing machines (local execution, transitions of elementary complexity and strong determinism) with those of the λ -calculus (higher-order functional programming, and the possibility of ‘parallelising’ the execution).

From this it would appear that studying a denotational semantics for the interaction combinators may be, at least in principle, as interesting as studying the denotational semantics of the λ -calculus. And yet to this day there have been very few efforts in this direction; the only work dealing directly with the semantics of the interaction combinators is Lafont’s original paper, in which a path semantics for nets of combinators is introduced, and an interpretation in terms of stack automata is given. Another contribution with a semantical flavour is Fernández and Mackie (2003), in which the fundamental operational equivalences for the interaction combinators are obtained as an application of more general results.

Our work aims precisely at deepening the semantical study of the interaction combinators. In a previous paper (Mazza 2006), we analysed the notion of observational equivalence for nets of combinators. A congruence analogous to $\beta\eta$ -equivalence is defined, and an internal separation result similar to Böhm’s Theorem is proved for it. This paper takes this result as a basis for developing a denotational semantics, that is, we find a mathematical structure in which $\beta\eta$ -equivalence becomes an equality.

More precisely, our semantics is inspired by the relational semantics for linear logic: nets are interpreted as subsets of a certain domain \mathcal{D} , called *interaction sets*, which do not need to have any particular structure apart from the existence of two bijections between $\mathcal{D} \times \mathcal{D}$ and \mathcal{D} itself that satisfy a certain condition. As expected, this semantics can be proved to be injective with respect to $\beta\eta$ -equivalence, that is, two nets are $\beta\eta$ -equivalent if and only if they have the same semantical interpretation. Moreover, we prove a full-completeness result with respect to a certain class of subsets, called *balanced* subsets, which is reminiscent of a similar result proved by Michele Pagani for multiplicative proof-nets (Pagani 2007).

We also consider interaction sets with a minimum of algebraic structure, namely that of a monoid. These structures, which are called *interaction monoids*, have the property of naturally inducing an algebraic semantics for the combinators, which is a model of the geometry of interaction described by Lafont. In this semantics, a net μ is interpreted as a pair of monoid endomorphisms (u, σ) , where $\sigma = \mathbf{0}$ (the everywhere-zero endomorphism) means that μ is in normal form. When $\sigma \neq \mathbf{0}$, and if μ does have a normal form, the endomorphism interpreting it can be computed by means of Girard’s *execution formula* $\text{Ex}(u, \sigma)$.

The denotational and algebraic semantics are tightly connected to each other: in fact we prove that if μ is a net admitting a normal form, given (u, σ) , the denotational semantics of μ is equal to the submonoid of the fixpoints of $\text{Ex}(u, \sigma)$; conversely, the denotational semantics of μ defines the endomorphism interpreting its normal form.

However, there is an important technical point to be clarified first: the semantics we discuss here does not deal with the interaction combinators, but with a slightly different variant, which we call the *symmetric combinators*. This interaction net system, also introduced by Lafont (Lafont 1997), is not universal in the same sense as the interaction combinators, but is just as expressive. In particular, every application of the interaction combinators found so far (for instance Mackie and Pinto's encoding of linear logic and the λ -calculus (Mackie and Pinto 2002)) can be reformulated with virtually no change using the symmetric combinators. Of course, our work on observational equivalence mentioned above also applies *mutatis mutandi* to this system.

The symmetric combinators are also tightly connected to the *directed combinators* (Lafont 1997), which are an extension of multiplicative linear logic proof-structures, and may therefore have interesting logical properties.

Contents of the paper

Section 2 contains introductory material required for the development in the rest of the paper. The exposition is as self-contained as possible, so even a reader completely unfamiliar with interaction nets should be able to follow the technical contents. In particular, in Section 2.4 we give an explicit proof of the expressiveness of the symmetric combinators, by encoding the **SK** combinators in them, and in Section 2.5 we briefly recall the main results of the above-mentioned paper (Mazza 2006), which we use later in one of our proofs.


Section 3 is the heart of the paper, and contains the definition of our denotational semantics, together with the injectivity and full completeness proofs. In Section 4 we introduce interaction monoids, and develop the algebraic semantics described in Lafont's original paper, proving the relationship between it and our denotational semantics.

Section 5 concludes the paper with a discussion of the technical reasons behind our choice of the symmetric combinators instead of the 'standard' interaction combinators, and gives some hints on future work.

2. The symmetric interaction combinators

2.1. Cells, wires, nets


The *symmetric interaction combinators*, or, more simply, the *symmetric combinators*, are the following three *cells*:



Each cell has a number of *ports*: δ and ζ have three ports; ε has only one. The fundamental property of cells is that *exactly one* of their ports is *principal* (drawn at the bottom in the above graphical representation), the others are *auxiliary*.

The auxiliary ports of the two binary combinators are ordered: to distinguish them, we call one the *left* port and the other the *right* port. Of course it does not matter which one we choose to call ‘left’ and which we call ‘right’, as long as once the convention is set, we stick to it. In this paper, we use this terminology in reference to the picture above: when cells are drawn like this, left ports are actually on the left, and right ports actually on the right. Note, however, that if a cell is drawn ‘upside-down’ with respect to the above representation, its left port will be on the right of the picture, and *vice versa*.

Ports may be used to plug cells[†] together by means of *wires* to form *nets*, as in the following example:




Wires can have one or both of their extremities not connected to any cell, in which case the net has a *free port*, which may be principal or auxiliary (or neither), depending on the nature of the port of the cell connected to the other extremity of the wire. For example, the net above has 7 free ports, of which 1 is principal and 4 are auxiliary. The free ports of a net are referred to as its *interface*. The set of all ports of all cells contained in a net μ , with the addition of its free ports, is denoted $\text{Ports}(\mu)$.

2.2. Interaction rules

The distinction between principal and auxiliary ports comes into play when defining the dynamics of nets. As a matter of fact, when two cells are connected through their principal ports, they form an *active pair*, and they may be replaced by another subnet according to the appropriate *interaction rules*.


There are six interaction rules, one for each possible active pair. Interaction rules are divided into two groups: the *annihilations*, describing what happens when two cells of the same kind form an active pair, and the *commutations*, describing what happens if the active pair is composed of two cells of different kinds.

If we put $\alpha \in \{\delta, \zeta\}$, the six rules can be condensed into four basic schemes: the annihilations



[†] Here, and throughout the rest of the paper, we shall systematically confuse cells and *occurrences* of cells.

and the commutations



When a net μ' is obtained from μ by the application of one of the above rules, we say that μ *reduces in one step* to μ' , and we write $\mu \rightarrow \mu'$. We can then define the *reduction relation* \rightarrow^* as the reflexive-transitive closure of \rightarrow . We write $\mu \simeq_\beta \mu'$ if and only if there exists μ'' such that $\mu \rightarrow^* \mu''$ and $\mu' \rightarrow^* \mu''$.

Notice that interaction rules are purely local; if we add to this the fact that cells have exactly one principal port, we immediately obtain the following proposition.


Proposition 2.1 (Strong confluence). The relation \rightarrow is confluent, that is, if μ, μ_1, μ_2 are three nets such that $\mu \rightarrow \mu_1$ and $\mu \rightarrow \mu_2$, then there exists μ' such that $\mu_1 \rightarrow \mu'$ and $\mu_2 \rightarrow \mu'$.

Proposition 2.1 means that the reduction process, that is, the relation \rightarrow^* , is *strongly* confluent. Confluence implies that \simeq_β is an equivalence relation, and that the system is deterministic in the sense that each net has at most one normal form. Strong confluence reinforces this determinism, because it implies that the computation is also unique, up to permutation of rules.

Note that there is here a substantial difference with respect to the λ -calculus, which is the absence of a meaningful concept of *strategy*, at least as far as the efficiency of reduction is concerned. All reductions leading from a net to its normal form have the same length and require the same work; in particular, if a net is normalisable, it is strongly so.

2.3. Basic nets

Vicious circles and cut-free nets. A net may contain configurations that cannot be removed through interaction, such as



in which clearly no cell can interact first (there is a sort of deadlock). The following case is even simpler:



Deadlocked configurations like these are called *vicious circles*.

Definition 2.1 (Straight path). Let μ be a net, $i, j \in \text{Ports}(\mu)$, and p be a path from i to j in the graph-theoretical sense. We say that p is *straight* if and only if whenever p enters a cell through one of its auxiliary ports, it exits it by its principal port, and whenever p enters a cell through its principal port, it exits it by one of its auxiliary ports.

Notice that the last requirement implies that a straight path can never ‘bounce back’ out of an ε cell.


Definition 2.2 (Vicious circle). A *vicious circle* is a cyclic straight path that never crosses two principal ports in a row.

A net containing no active pair and no vicious circle is said to be *cut-free* (or *reduced* in Lafont’s terminology). A net admitting a cut-free form through reduction (which by confluence is necessarily unique) is said to be *total*.


Cut-free nets are the ‘true’ normal forms of the reduction; they can be thought of as the final result of a computation. On the other hand, non-total nets represent error-bound computations, which are either diverging or leading to a deadlock.

Usually, when one translates classical computational models (for example, Turing machines and the λ -calculus) into the interaction combinators (or more generally into interaction nets), the encoding is such that vicious circles never arise through reduction. An example is the encoding of the **SK**-combinators we give in Section 2.4. Those acquainted with linear logic will see that vicious circles are a legacy of proof-structures: incorrect proof-structures may reduce to nets containing cuts that are not reducible according to the standard cut-elimination procedure (for example, an axiom whose conclusions are premises of a cut rule). Indeed, this is the reason behind our ‘cut-free’ terminology.


Wirings. A net containing no cells but just wires will be called a *wiring*. We shall represent the generic wiring as



The following is an example of wiring:




We also allow the free ports of a wiring to belong to ε cells, in which case we speak of an ε -wiring and use the notation $\tilde{\omega}$. The following is an example of an ε -wiring:



Trees. Trees are defined inductively as follows. The wiring




is a tree with one leaf (it is arbitrary which of the two extremities is the root and which is the leaf). If τ_1 and τ_2 are two trees with n_1 and n_2 leaves, respectively, then we can define a tree τ with $n_1 + n_2$ leaves as




where $\alpha \in \{\delta, \zeta\}$.

It is not hard to verify that any cut-free net ν with n free ports can be decomposed in terms of trees and ε -wirings as follows:




Principal nets, tests and packages. A principal net of arity n is either a single wire (in which case $n = 1$), or a cut-free net with n free auxiliary ports and 1 free principal port. If $n = 0$ (respectively, $n = 1$), we say that the net is a *package* (respectively, a *test*). Principal nets can be thought of as ‘compound’ cells, and will be drawn just like ordinary cells.

Notice that trees are special examples of principal nets. A particular family of principal nets of arity $n \geq 0$, the members of which are denoted \mathbf{Z}_n , and which are trees for $n \geq 1$, is defined as follows:



One can check that Z_n trees have the following annihilation property:



2.4. Expressive power

In spite of their great simplicity, the symmetric combinators are Turing-complete. This is a consequence of the following more general result.


Theorem 2.2 (Universality (Lafont 1997)). Any polarised interaction net system can be translated into the symmetric combinators.

The symmetric combinators are a particular example of *non-polarised* interaction net system; their polarised version are what Lafont calls the *directed combinators* (Lafont 1997).

The definition of generic and polarised interaction net systems, and of the translation from an interaction net system to another are beyond the scope of this paper, but to understand the importance of Theorem 2.2, it is enough to know that Turing machines, one-dimensional cellular automata, the **SK** combinators, linear logic proof-nets and the λ -calculus can all be seen as polarised interaction net systems (for the latter two, see, for example, the survey Lafont (1995), Mackie and Pinto (2002), and Lippi (2002)).


Even though Theorem 2.2 will not be proved here, and its meaning will be left to the intuition of the reader, we shall, nevertheless, give a direct proof of the expressive power of the symmetric combinators. In fact, in the rest of this section we shall show how the call-by-name **SK** combinators can be encoded within this system.

We begin by introducing a fundamental construction, which is due to Lafont (Lafont 1997). We take a generic package π containing n δ cells. We can always write π as




where π' contains no δ cell. We want to ‘abstract’ the δ cells contained in π to form a package $!\pi$ that does not contain δ cells, but from which π can be recovered.


Using the notation introduced above, we define



Then, we put




One can then check that we have



The package $!\pi$ is called the *code* of π , and the principal net \mathbf{D} is called the *decoder*.


It is not hard to prove the following lemma.

Lemma 2.3 (Erasure and duplication). Let π be a package and π' be a package containing no δ cells. Then, we have




As a consequence, a package of the form $!\pi$ can be erased and duplicated.


Now we build the two packages



and



and define the translation $[\cdot]$ from **SK**-terms to nets:



If x, y are **SK**-terms, we write $x > y$ if x reduces to y through call-by-name reduction. For example, $\mathbf{K}(\mathbf{KSK})\mathbf{S} > \mathbf{KSK}$, but $\mathbf{K}(\mathbf{KSK})\mathbf{S} \not> \mathbf{KSS}$.

The translation defined above has the following property.

Theorem 2.4. Let x, y be **SK**-terms. Then, $x >^* y$ implies $[x] \rightarrow^* [y]$.

Proof. It is enough to check that, for all terms x, y, z ,


$$[\mathbf{K}xy] \rightarrow^* [x]$$

and


$$[\mathbf{S}xyz] \rightarrow^* [xz(yz)].$$

The first verification is easy and is left to the reader (Lemma 2.3 is needed for erasing). Here we shall concentrate on the second verification, which is more complex.


We have




The \mathbf{Z}_4 tree annihilates with the \mathbf{Z}_4 tree contained in \mathbf{S} (see the end of Section 2.3), and we obtain




Now the decoder extracts $[x]$ from its code, and by Lemma 2.3 the code of $[z]$ is duplicated, so we get



At this point, the ‘topmost’ Z_4 tree annihilates with the Z_4 tree inside $![y]$, and we are left with



Remember that $![z]$ does not contain any δ cell, which means that




Therefore, the final net obtained is indeed equal to $[xz(yz)]$. □

2.5. Observational equivalence

We showed recently in Mazza (2006) that there is a notion of observational equivalence for the symmetric combinators[†] that can be defined directly on the syntax, and is maximal on total nets with at least one free port.

We define the relation \simeq_η as the reflexive, transitive and contextual closure of the following equations:



As usual, $\alpha \in \{\delta, \zeta\}$. The top-right and bottom rules were already known to Lafont (Lafont 1997) and the top-left one to Fernández and Mackie (Fernández and Mackie 2003). What is interesting is that they define an equivalence relation that is very like η -equivalence in the λ -calculus (this is the reason for our notation).

We define $\simeq_{\beta\eta}$ as the transitive closure of $\simeq_\beta \cup \simeq_\eta$. If θ is a test and π a package, we write $\theta[\pi]$ for the net with one free port obtained by plugging π into the only free principal of θ (or to any free port if θ is a wire).

In the following, ϵ stands for the package consisting of a single ϵ combinator, while δ stands for the the package containing one δ cell whose auxiliary ports are connected by a wire.


Theorem 2.5 (Internal separation (Mazza 2006)). Let π, π' be two packages such that $\pi \not\approx_{\beta\eta} \pi'$. Then, there exists a test θ such that $\theta[\pi] \rightarrow^* \epsilon$ and $\theta[\pi'] \rightarrow^* \delta$, or vice versa.

[†] In fact, our results were originally formulated for the interaction combinators, but they apply equally well to the symmetric combinators.

The above result, which is similar to Böhm’s Theorem for the λ -calculus, is of fundamental importance for the theory of observational equivalence. In particular, it implies the following proposition.

Proposition 2.6 (Maximality). $\simeq_{\beta\eta}$ is the greatest non-trivial congruence on total nets respecting reduction, that is, if \approx is a congruence on total nets such that $\simeq_{\beta} \subseteq \approx$, then either $\approx \subseteq \simeq_{\beta\eta}$, or $\mu \approx \mu'$ for all total nets μ, μ' with the same number of free ports.

Proof. Suppose that $\approx \not\subseteq \simeq_{\beta\eta}$. This means that there exist two total nets μ, μ' , both with $n \geq 1$ free ports, such that $\mu \approx \mu'$ but $\mu \not\approx_{\beta\eta} \mu'$. First, take any tree τ with n leaves and build the packages π and π' by ‘closing’ the cut-free forms of μ and μ' , respectively. Since \approx is a congruence, we have $\pi \approx \pi'$; moreover, it can be proved that just adding a tree ‘below’ μ and μ' does not affect their differences with respect to $\simeq_{\beta\eta}$, so we still have $\pi \not\approx_{\beta\eta} \pi'$. Therefore, Theorem 2.5 applies, giving us a test θ such that, for example, $\theta[\pi] \rightarrow^* \varepsilon$ and $\theta[\pi'] \rightarrow^* \delta$. Since \approx is a congruence, and since $\simeq_{\beta} \subseteq \approx$, we have $\varepsilon \approx \delta$. But then consider the principal net



where v is any total net with $m \geq 1$ free ports. If we plug ε into the only free principal port of v' , we obtain a net that is β -equivalent to a net consisting of m ε cells, which we call ε_m ; on the other hand, if we perform the same operation with the package δ , we obtain a net β -equivalent to v . Hence, again by the fact that \approx is a congruence containing \simeq_{β} , for any total v with $m \geq 1$ free ports, $v \approx \varepsilon_m$, and by transitivity of \approx , we get the thesis. □

The results of this section will guide us in our search for a denotational semantics for the symmetric combinators. In fact, once a good observational equivalence like $\simeq_{\beta\eta}$ is found, the goal of denotational semantics can be thought of as the individuation of a mathematical structure in which $\simeq_{\beta\eta}$ becomes an equality.

3. Denotational semantics

As we mentioned earlier, interaction nets are a generalisation of multiplicative proof-nets. Thus, in seeking a denotational semantics for the symmetric combinators, it seems natural to draw inspiration from linear logic.

The simplest denotational semantics of linear logic is the *relational semantics*. As far as we are aware, there is no formal introduction to relational semantics in the literature. The best way to think of it is, perhaps, as ‘coherent spaces without coherence’; it has

been considered by many as the starting point for building other denotational semantics of linear logic (semantics = relational semantics + structure), as, for example, in the work of Thomas Ehrhard (Ehrhard 2005). In the multiplicative case, the relational semantics is obtained from coherent spaces by simply ignoring the coherence relation: a formula A is interpreted by a set $|A|$, and a proof of A by a subset of $|A|$; no particular structure is attached to or required for $|A|$. In categorical terms, the denotational interpretation takes place in the category **Rel** of sets and relations.

The denotational interpretation of a sequent calculus proof is, as usual, defined by induction. More interestingly, the interpretation can also be defined directly on proof-structures (and thus, in particular, on proof-nets) by means of *experiments* (Girard 1987a). This will be our main source of inspiration, as a sequent calculus is obviously not available in our framework.

3.1. Companion bijections

The relational semantics of linear logic is typed. In particular, the two multiplicative connectives (which are the only ones of interest to us) are interpreted by the *cartesian product*: if the formulas A and B are interpreted by $|A|$ and $|B|$, respectively, then $A \otimes B$ and $A \wp B$ are interpreted by $|A| \times |B|$. Our nets are not typed, which means that if we think of our binary combinators as multiplicative rules, the natural thing to do would be to consider a set \mathcal{D} in bijection with $\mathcal{D} \times \mathcal{D}$, that is, any infinite set.

The presence of *two* combinators actually requires two bijections, and the $\delta\zeta$ commutation inspires the following definition.

Definition 3.1 (Companion bijections). Let \mathcal{D} be an infinite set and

$$\langle \cdot, \cdot \rangle, [\cdot, \cdot] : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{D}$$

be two bijections. $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ are said to be *companions* if and only if, for all $a, b, c, d \in \mathcal{D}$,

$$\langle [a, b], [c, d] \rangle = [\langle a, c \rangle, \langle b, d \rangle].$$

Companion bijections do exist, as proved by the following example. Let \mathcal{A} be any infinite set, and $\beta : \mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}$ be a bijection. We use $\mathcal{S}(\mathcal{A})$ to denote the set of all infinite sequences of elements of \mathcal{A} . Then we define two functions from $\mathcal{S}(\mathcal{A}) \times \mathcal{S}(\mathcal{A})$ to $\mathcal{S}(\mathcal{A})$ as follows:

$$\langle d, e \rangle_n = \begin{cases} d_k & \text{if } n = 2k \\ e_k & \text{if } n = 2k + 1 \end{cases}$$

$$[d, e]_n = \beta(d_n, e_n).$$

In other words, $\langle \cdot, \cdot \rangle$ takes two sequences and builds a new one by interleaving them, while $[\cdot, \cdot]$ simply superposes the two sequences using the bijection β . The two functions are clearly bijections; moreover, given four sequences a, b, c, d , for even indexes we get

$$\langle [a, b], [c, d] \rangle_{2k} = [a, b]_k = \beta(a_k, b_k) = \beta(\langle a, c \rangle_{2k}, \langle b, d \rangle_{2k}) = [\langle a, c \rangle, \langle b, d \rangle]_{2k},$$

and similarly for odd indexes, which proves that the two bijections are indeed companions.

If \mathcal{A} is a denumerably infinite pointed set, whose distinguished element we call zero, the same example can be built on $\Phi(\mathcal{A})$, the set of infinite sequences of elements of \mathcal{A} that are almost everywhere zero. So companion bijections also exist on countable sets, in particular on $\Phi(\mathbb{N})$, which by the fundamental theorem of arithmetic is isomorphic (as a monoid) to \mathbb{N}^* , the strictly positive integers.

3.2. Experiments and interpretation

The previous section suggests the following definition.

Definition 3.2 (Interaction set). An *interaction set* is an infinite pointed set \mathcal{D} (the distinguished element being denoted by 0) admitting two companion bijections

$$\langle \cdot, \cdot \rangle, [\cdot, \cdot] : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{D}$$

such that $\langle 0, 0 \rangle = [0, 0] = 0$.

The examples above show that interaction sets exist; for instance, in $\Phi(\mathbb{N})$, 0 is the everywhere-zero sequence.

In the following, \mathcal{D} will be an interaction set.

Definition 3.3 (Experiment). Let μ be a net. An *experiment* on μ is a function $e : \text{Ports}(\mu) \rightarrow \mathcal{D}$ such that:

- (a) If $i, j \in \text{Ports}(\mu)$ are connected by a wire, then $e(i) = e(j)$.
- (b) If $i, j, k \in \text{Ports}(\mu)$ are the left auxiliary, right auxiliary and principal ports, respectively, of a δ cell of μ , then $e(k) = \langle e(i), e(j) \rangle$.
- (c) If $i, j, k \in \text{Ports}(\mu)$ are the left auxiliary, right auxiliary and principal ports, respectively, of a ζ cell of μ , then $e(k) = [e(i), e(j)]$.
- (d) If $i \in \text{Ports}(\mu)$ is the principal port of an ε cell, then $e(i) = 0$.

If k_1, \dots, k_n are the free ports of μ , with $n \geq 1$, the tuple $(e(k_1), \dots, e(k_n))$ is called the *result* of the experiment and is denoted by $|e|$.

In the following, we write \mathcal{D}^n for $\mathcal{D} \times \dots \times \mathcal{D}$ n times.

Definition 3.4 (Interpretation). Let μ be a net with $n \geq 1$ free ports. The *interpretation* of μ in \mathcal{D} , written $\llbracket \mu \rrbracket$, is defined to be the subset of \mathcal{D}^n containing the results of all possible experiments on μ :


$$\llbracket \mu \rrbracket = \{|e| ; e \text{ experiment on } \mu\}.$$

We can give a few examples to see some concrete applications of the above definition. Consider the package ε consisting of a single ε cell. There is only one possible experiment on it, which assigns 0 to the principal port of the ε cell and to the free port of the package, so $\llbracket \varepsilon \rrbracket = \{0\}$. This is the smallest possible interpretation a net can receive; as a matter of fact, the interpretation of a net is always a pointed set, that is, it is never empty. This is an immediate consequence of the definition of experiment and of the fact that 0 is a fixpoint of both bijections.

Proposition 3.1. For all μ with at least one free port, $(0, \dots, 0) \in \llbracket \mu \rrbracket$.


Proof. The function assigning 0 to all ports of μ is an experiment. □

If we take the package δ consisting of a single δ cell whose auxiliary ports are connected by a wire, we clearly have that all possible experiments are of the following form:




so $\llbracket \delta \rrbracket = \{ \langle d, d \rangle ; d \in \mathcal{D} \}$. Just as with an axiom in linear logic, the net ω with 2 free ports consisting of a single wire is interpreted by the diagonal relation in $\mathcal{D} \times \mathcal{D}$: $\llbracket \omega \rrbracket = \{ (d, d) ; d \in \mathcal{D} \}$.

The following is a more involved example:



In this example, a label d on a wire means that the two ports connected by the wire have both been assigned the element d by the experiment, and a is a generic element of \mathcal{D} . We therefore see that if we call the above net μ , we have $\llbracket \mu \rrbracket = \{ [0, \langle a, a \rangle] ; a \in \mathcal{D} \}$. One can check that this is also the interpretation of the net



which is the cut-free form of μ .

In fact, the interpretation is a denotational semantics for the symmetric combinators, that is, it is preserved under reduction. Additionally, it also models \simeq_η .


Lemma 3.2 (Stability under reduction). Let μ, μ' be two nets with at least one free port. Then, $\mu \rightarrow \mu'$ implies $\llbracket \mu \rrbracket = \llbracket \mu' \rrbracket$.

Proof. We need to show that for any experiment e on μ , there exists an experiment e' on μ' yielding the same result, and *vice versa*. Since the rewriting is local, it actually suffices to show that for all reduction rules, the assignment given by the experiment e

on the interface of the left member of the rule can be reproduced by e' on the interface of the right member, and *vice versa*; at this point e and e' can be assumed to be equal everywhere else, which guarantees that the results are the same.


The case of the $\varepsilon\varepsilon$ annihilation is trivial: e' is just e restricted to the ports that do not disappear after the application of the rule.

The $\delta\delta$ and $\zeta\zeta$ annihilations are structurally identical, so we shall only consider the first of them:




Here, a, b, a', b' are generic elements of \mathcal{D} . The assignment on the left-hand side must satisfy $\langle a, b \rangle = \langle a', b' \rangle$, which by the injectivity of $\langle \cdot, \cdot \rangle$ implies $a = a'$ and $b = b'$, so the assignment on the right-hand side is correct. The converse is trivial.

For the commutations, the $\delta\varepsilon$ and $\zeta\varepsilon$ commutations are again structurally identical, so we only need to consider the first of them:



Again, by injectivity of $\langle \cdot, \cdot \rangle$, the requirement on the left-hand side that $\langle a, b \rangle = 0$ implies $a = 0$ and $b = 0$, so the assignment on the right-hand side is correct. The converse holds because of the hypothesis that $\langle 0, 0 \rangle = 0$.

On the other hand, for the $\delta\zeta$ commutation, we get



On the left-hand side we must have $\langle a, b \rangle = \langle a', b' \rangle$. By surjectivity of $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$, there exist $a_1, a_2, a'_1, a'_2, b_1, b_2, b'_1, b'_2 \in \mathcal{D}$ such that $a = [a_1, a_2]$, $b = [b_1, b_2]$, $a' = \langle a'_1, a'_2 \rangle$, and $b' = \langle b'_1, b'_2 \rangle$. The above equality and the fact that $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ are companions imply $[\langle a'_1, a'_2 \rangle, \langle b'_1, b'_2 \rangle] = [[a_1, b_1], \langle a_2, b_2 \rangle]$, which by injectivity of $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ in turn implies $a'_1 = a_1$, $a'_2 = b_1$, $b'_1 = a_2$, and $b'_2 = b_2$. Therefore, the assignment defined above for the right-hand side of the rule is correct. Conversely, if we know from the right-hand side that $a'_1 = a_1 = c_1$, $a'_2 = b_1 = c_2$, $b'_1 = a_2 = c_3$, and $b'_2 = b_2 = c_4$, we have $a = [c_1, c_3]$, $b = [c_2, c_4]$, $a' = \langle c_1, c_2 \rangle$, and $b' = [c_3, c_4]$, which means that $\langle a, b \rangle = \langle [c_1, c_3], [c_2, c_4] \rangle$ and $[a', b'] = [[c_1, c_2], \langle c_3, c_4 \rangle]$. But since $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ are companions, this implies that $\langle a, b \rangle = [a', b']$, so the assignment on the left-hand side is correct. \square

In the following, if μ is a net with n free ports, we say any net C with at least $n + 1$ free ports, n of which are in bijection with the free ports of μ , is a *context* for μ ; the application of the context to μ , written $C[\mu]$, is the net obtained by plugging each free port of μ into the corresponding free port of C .

Lemma 3.3 (Congruence). Let μ, μ' be two nets with n free ports, and let C be a context for μ, μ' . Then, $[[\mu]] = [[\mu']]$ implies $[[C[\mu]]] = [[C[\mu']]]$.

Proof. An experiment on $C[\mu]$ must be the union of an experiment e on μ and an experiment f on C , such that, if i and j are two free ports of μ and C , respectively, that are connected in $C[\mu]$, then $e(i) = f(j)$. The same holds for $C[\mu']$, so, if m is the number of free ports of $C[\mu]$ and $C[\mu']$, we have


$$[[C[\mu]]] = \{d \in \mathcal{D}^m ; (d, c) \in [[C]] \text{ and } c \in [[\mu]]\}$$

$$[[C[\mu']]] = \{d' \in \mathcal{D}^m ; (d', c) \in [[C]] \text{ and } c \in [[\mu']]\},$$

from which we clearly see that if $[[\mu]] = [[\mu']]$, then $[[C[\mu]]] = [[C[\mu']]]$. \square

Lemma 3.4 (Extensionality). Let μ, μ' be two nets with at least one free port. Then, $\mu \simeq_\eta \mu'$ implies $[[\mu]] = [[\mu']]$.


Proof. The proof follows exactly the same argument used for Lemma 3.2. We start by considering the η -expansion for δ , the corresponding rule involving ζ being structurally identical:



The left-hand side requires $a'_1 = a_1$ and $a'_2 = a_2$, which implies $a = a'$. Conversely, the right-hand side requires $a'_1 = a_1$; by surjectivity of $\langle \cdot, \cdot \rangle$, a_1 and a_2 such that $\langle a_1, a_2 \rangle = a$ exist, so the assignment on the left-hand side is correct.

The cases of the $\delta\varepsilon$ and $\zeta\varepsilon$ commutations are trivial, and rest upon the fact that $\langle 0, 0 \rangle = [0, 0] = 0$.

The case of the $\delta\zeta$ commutation is also trivial:



The two experiments are the same thanks to the fact that $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ are companions. □

Lemmas 3.2 and 3.3 together prove that $\llbracket \cdot \rrbracket$ is a denotational semantics of \simeq_β : it is preserved under reduction, and it is a congruence. Lemma 3.4 proves that the semantics actually models $\simeq_{\beta\eta}$. Moreover, the examples given after Definition 3.4 show that there exist two nets μ, μ' such that $\llbracket \mu \rrbracket \neq \llbracket \mu' \rrbracket$, so the semantics is non-trivial.

The fact that two $\beta\eta$ -equivalent nets have the same interpretation holds for *any* net; if we restrict consideration to total nets, the converse holds as well.

Theorem 3.5 (Injectivity on total nets). Let μ, μ' be two total nets with at least one free port. Then, $\mu \simeq_{\beta\eta} \mu'$ if and only if $\llbracket \mu \rrbracket = \llbracket \mu' \rrbracket$.

Proof. As we have already observed, the implication $\mu \simeq_{\beta\eta} \mu' \Rightarrow \llbracket \mu \rrbracket = \llbracket \mu' \rrbracket$ is a consequence of Lemmas 3.2 and 3.4. For the converse, which is the actual injectivity property, we can restrict consideration to packages, since:

- total nets containing active pairs can be reduced and their cut-free form considered;
- cut-free nets with more than one free port can be ‘closed’ by means of any fixed tree, and the argument below can thus be easily generalised.

So, we take two packages π, π' such that $\pi \not\simeq_{\beta\eta} \pi'$. By Theorem 2.5, there exists a test θ such that $\theta[\pi] \rightarrow^* \varepsilon$ and $\theta[\pi'] \rightarrow^* \delta$ (or *vice versa*, but we do not lose generality in assuming this situation), where ε is the package consisting of a single ε cell and δ is the package consisting of a single δ cell with its auxiliary ports connected by a wire. By Lemma 3.2, and the examples given after Definition 3.4, we have

$$\begin{aligned} \llbracket \theta[\pi] \rrbracket &= \{0\} \\ \llbracket \theta[\pi'] \rrbracket &= \{ \langle d, d \rangle ; d \in \mathcal{D} \}. \end{aligned}$$

Since \mathcal{D} is infinite, it is clear that $\llbracket \mu \rrbracket \neq \llbracket \mu' \rrbracket$, which by Lemma 3.3 implies $\llbracket \pi \rrbracket \neq \llbracket \pi' \rrbracket$. □


One may wonder whether injectivity with respect to $\simeq_{\beta\eta}$ can be extended to non-total nets. The answer is in general negative, as the following example shows. Let $\Phi(\mathbb{N})$ be the space of almost-everywhere-null sequences of natural numbers. The distinguished element 0 is the everywhere-zero sequence. We know that $\Phi(\mathbb{N})$ admits the two companion bijections defined in Section 3.1; in particular, we recall the definition of $\langle \cdot, \cdot \rangle$:

$$\langle x, y \rangle_n = \begin{cases} x_k & \text{if } n = 2k \\ y_k & \text{if } n = 2k + 1 \end{cases}$$

It is not hard to verify the following result.

Lemma 3.6. $\langle x, y \rangle = x$ if and only if $y = 0$ and x is a sequence such that, for all $n \geq 1$, $x_n = 0$.

Now consider the following non-total net μ :




The labels indicate that the generic experiment on μ in $\Phi(\mathbb{N})$ assigns the sequence x to the left auxiliary and principal ports of the δ cell, and the sequence y to the right auxiliary port of the δ cell and to the free port of μ . Therefore, we have

$$\llbracket \mu \rrbracket = \{y \in \Phi(\mathbb{N}) ; \exists x \in \Phi(\mathbb{N}). \langle x, y \rangle = x\}.$$

But Lemma 3.6 proves that the only possible such y is 0, so $\llbracket \mu \rrbracket = \{0\} = \llbracket \varepsilon \rrbracket$, where ε is the package consisting of a single ε combinator. Now, both μ and ε do not contain active pairs, hence the only hope of rewriting one into the other is through η -equivalence. But a simple inspection of the η -rules in Section 2.5 reveals that the presence of ε cells is preserved by \simeq_{η} : no rule can produce a net containing no ε cell from a net containing one, and no rule can add ε cells if there are none. Therefore, $\mu \not\simeq_{\beta\eta} \varepsilon$, and again $\llbracket \mu \rrbracket = \llbracket \varepsilon \rrbracket$.

Note that identifying the net μ above with the ε combinator is observationally sound. Indeed, if we define a *blind net* as a net such that all of its free ports are auxiliary and will never become principal through reduction, then it is consistent to identify any blind net with n free ports to any net $\beta\eta$ -equivalent to the following:



In fact, they interact in the same way. In particular, this means that Proposition 2.6 does not hold for generic nets.

One may wonder whether this failure of injectivity goes beyond the particular model chosen above, and actually holds for any interaction set. Although at present we are not able to answer the question, we tend to believe that this is indeed the case, and we actually see this phenomenon as quite natural, considering the above remark about observational

equivalence. Indeed, the ε combinator and the deadlocked net used in the counter-example can be seen as the equivalents of two *unsolvable* terms in the λ -calculus: like Ω and $\lambda x.\Omega$, they are not $\beta\eta$ -equivalent, but are identified by any sensible model.

3.3. Full completeness

If \mathcal{D} is an interaction set, even a denumerable one, for obvious reasons of cardinality not every subset of \mathcal{D}^n is the interpretation of some net. In this section we characterise those that are interpretations of *total* nets.

In the following, \mathcal{D} is a generic interaction set.

Definition 3.5 (Bracket expression). Let x range over a denumerable set of variables. A *simple bracket expression* \mathbf{b} is a syntactical expression belonging to the following grammar:

$$\mathbf{b} ::= x \mid 0 \mid \langle \mathbf{b}, \mathbf{b} \rangle \mid [\mathbf{b}, \mathbf{b}].$$

A *bracket expression* is a tuple of simple bracket expressions.

We use $\text{var}(\mathbf{b})$ to denote the set of variables occurring in the simple bracket expression \mathbf{b} . We define the substitution of a variable y in place of x in \mathbf{b} , denoted $\mathbf{b}[y/x]$, in the normal way. If $x \in \text{var}(\mathbf{b})$ and $z \notin \text{var}(\mathbf{b})$, we say that \mathbf{b} and $\mathbf{b}[x/z]$ are α -equivalent.

Similarly, if $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ is a bracket expression, we define

$$\text{var}(\mathbf{B}) = \text{var}(\mathbf{b}_1) \cup \dots \cup \text{var}(\mathbf{b}_n),$$

that is, variables are shared by the simple expressions in the tuple, and substitution is performed on the whole expression; α -equivalence is trivially extended, and bracket expressions are always considered modulo α -equivalence.

If \mathbf{B} is a bracket expression containing n simple expressions such that

$$\text{var}(\mathbf{B}) \subseteq \{x_1, \dots, x_m\}$$

and if $d_1, \dots, d_m \in \mathcal{D}$, we can define an element $\mathbf{B}\{x_1 := d_1, \dots, x_m := d_m\}$ of \mathcal{D}^n in the obvious way: just assign each d_i to x_i , and compute the expression considering the symbols 0 , $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ as the distinguished element and the two bijections of \mathcal{D} , respectively. For example, if we assume that d, e, f are three elements of \mathcal{D} such that $f = [\langle d, e \rangle, d]$, then, if $\mathbf{B} = [\langle x, y \rangle, x]$, we have $\mathbf{B}\{x := d, y := e\} = f$. In this way, each bracket expression \mathbf{B} containing n simple bracket expressions and a total of m variables defines a function from \mathcal{D}^m to \mathcal{D}^n . Because of the obvious shortage of bracket expressions, the assignment cannot be surjective; it is not injective either, as the expressions 0 and $\langle 0, 0 \rangle$ show (they both represent the constant function 0).

In the following, $\text{occ}_x(\mathbf{B})$ denotes the number of occurrences of the variable x in the bracket expression \mathbf{B} .

Definition 3.6 (Balanced bracket expression). A bracket expression \mathbf{B} is *balanced* if and only if, for any variable x , either $\text{occ}_x(\mathbf{B}) = 0$ or $\text{occ}_x(\mathbf{B}) = 2$. A function from \mathcal{D}^m to \mathcal{D}^n is said to be *balanced* if it can be defined through a balanced bracket expression containing n expressions using m variables. We say a set $\mathcal{B} \subseteq \mathcal{D}^n$ is *balanced* if it is the codomain of a balanced function.

Theorem 3.7 (Full completeness). If μ is a total net with $n \geq 1$ free ports, then $\llbracket \mu \rrbracket$ is balanced. Conversely, if $n \geq 1$, given a balanced set $\mathcal{B} \subseteq \mathcal{D}^n$, there exists a cut-free net μ with n free ports such that $\llbracket \mu \rrbracket = \mathcal{B}$.


Proof. Since μ is total, we can consider its cut-free form ν . By Lemma 3.2, $\llbracket \mu \rrbracket = \llbracket \nu \rrbracket$. Now, if we remember that cut-free nets are trees of δ and ζ cells with wires and ε cells ‘on top’, it is clear that the first statement is a straight-forward consequence of the definition of experiments.

For the converse, let $\mathcal{B} \subseteq \mathcal{D}^n$ be balanced and B be the bracket expression such that \mathcal{B} is the codomain of B . Simple bracket expressions can obviously be provided with a complexity measure $\#(\cdot)$, which is the total number of binary syntactical constructs used:

- $\#(x) = \#(0) = 0$;
- $\#(\langle b_1, b_2 \rangle) = \#[b_1, b_2] = \#(b_1) + \#(b_2) + 1$.

For an expression $B = (b_1, \dots, b_n)$, we put $\#(B) = \#(b_1) + \dots + \#(b_n)$.


We can then reason by induction on $\#(B)$. If $\#(B) = 0$, knowing that B is balanced, we can assume without loss of generality that $B = (0, \dots, 0, x_1, x_1, \dots, x_k, x_k)$. It is easy to see that if we interpret the net



(where $n' + k = n$), we obtain exactly \mathcal{B} .

If $\#(B) > 0$, we can assume without loss of generality that $B = (\langle b'_1, b''_1 \rangle, b_2, \dots, b_n)$ or $B = ([b'_1, b''_1], b_2, \dots, b_n)$. In both cases, the bracket expression $B' = (b'_1, b''_1, b_2, \dots, b_n)$ has measure strictly smaller than B , so the induction hypothesis applies, giving us a net μ' with $n + 1$ free ports such that $\llbracket \mu' \rrbracket$ is the codomain of B' . Clearly, adding a δ or a ζ cell (according to the shape of B) to μ' yields a net with n free ports μ such that $\llbracket \mu \rrbracket = \mathcal{B}$. \square

The above proof actually tells us that, for $n \geq 1$, the balanced subsets of \mathcal{D}^n are in bijection with the cut-free nets with n free ports. As an example, take the balanced sets induced by the expressions $[0, \langle x, x \rangle]$ and $([0, x], \langle x, 0 \rangle)$, which correspond to the following two nets:



This is not at all surprising: a balanced expression is just a list of trees, the connections between leaves being expressed as pairs of occurrences of the same variable. Therefore, balanced expressions are just an alternative, linear syntax for cut-free nets.

It must also be mentioned that this full completeness result is very similar to that proved by Michele Pagani for multiplicative linear logic proof-nets (Pagani 2007); in fact, this strengthens the claim that the symmetric combinators are an extension of multiplicative proof-structures, and may be deeply related to them.

4. Algebraic semantics

In this section we develop an algebraic semantics for the symmetric combinators in the style of Girard’s Geometry of Interaction (GoI) (Girard 1989), which has already been sketched by Lafont (Lafont 1997). We do this in connection with the denotational semantics introduced above, and prove that there is a strong link between the two.

4.1. Interaction monoids

We have already seen that any infinite set \mathcal{D} can serve as the domain for the denotational semantics of nets of symmetric combinators. We shall now show that it may be of interest to add some algebraic structure to \mathcal{D} ; the least we can require is that \mathcal{D} is a monoid.

In the following, the composition $u \circ v$ of two monoid endomorphisms is simply denoted by uv .

Definition 4.1 (Interaction monoid). An *interaction monoid* is a commutative monoid $(M, +, 0)$ admitting eight endomorphisms $\mathbf{c}, \mathbf{c}^*, \mathbf{d}, \mathbf{d}^*, \mathbf{f}, \mathbf{f}^*, \mathbf{g}, \mathbf{g}^*$ such that the functions $\langle x, y \rangle = \mathbf{c}(x) + \mathbf{d}(y)$ and $[x, y] = \mathbf{f}(x) + \mathbf{g}(y)$ are companion isomorphisms between $M \oplus M$ and M , and $\mathbf{c}^*, \mathbf{d}^*$ and $\mathbf{f}^*, \mathbf{g}^*$ are their respective projections.

Proposition 4.1. A commutative monoid $(M, +, 0)$ is an interaction monoid if and only if there exist eight endomorphisms $\mathbf{c}, \mathbf{c}^*, \mathbf{d}, \mathbf{d}^*, \mathbf{f}, \mathbf{f}^*, \mathbf{g}, \mathbf{g}^*$ of M such that:

- 1 $\mathbf{c}^* \mathbf{c} = \mathbf{d}^* \mathbf{d} = \mathbf{f}^* \mathbf{f} = \mathbf{g}^* \mathbf{g} = \mathbf{1}$, where $\mathbf{1}$ is the identity on M .
- 2 $\mathbf{c}^* \mathbf{d} = \mathbf{d}^* \mathbf{c} = \mathbf{f}^* \mathbf{g} = \mathbf{g}^* \mathbf{f} = \mathbf{0}$, where $\mathbf{0}$ is the everywhere-zero endomorphism on M .
- 3 $\mathbf{c} \mathbf{c}^* + \mathbf{d} \mathbf{d}^* = \mathbf{f} \mathbf{f}^* + \mathbf{g} \mathbf{g}^* = \mathbf{1}$.
- 4 $\mathbf{c}, \mathbf{c}^*, \mathbf{d}, \mathbf{d}^*$ commute with $\mathbf{f}, \mathbf{f}^*, \mathbf{g}, \mathbf{g}^*$.

Proof. We first prove that Definition 4.1 implies the four statements above:

- 1 By definition, for every $x \in M$, $\mathbf{c}(x) = \langle x, 0 \rangle$, and by the hypothesis that \mathbf{c}^* is the left projection of $\langle \cdot, \cdot \rangle$, we obtain $\mathbf{c}^* \mathbf{c}(x) = \mathbf{c}^*(\langle x, 0 \rangle) = x$. The same applies to the other annihilations.
- 2 As above, for every $x \in M$ we have $\mathbf{d}(x) = \langle 0, x \rangle$, from which we obtain $\mathbf{c}^* \mathbf{d}(x) = \mathbf{c}^*(\langle 0, x \rangle) = 0$. The same applies to the other annihilations.
- 3 From the surjectivity of $\langle \cdot, \cdot \rangle$, given $x \in M$ we know that there exist $y, z \in M$ such that $x = \langle y, z \rangle$. Then we have

$$(\mathbf{c} \mathbf{c}^* + \mathbf{d} \mathbf{d}^*)(x) = \mathbf{c} \mathbf{c}^*(x) + \mathbf{d} \mathbf{d}^*(x) = \mathbf{c}(y) + \mathbf{d}(z) = \langle y, z \rangle = x.$$

The case $\mathbf{f} \mathbf{f}^* + \mathbf{g} \mathbf{g}^* = \mathbf{1}$ is identical.

- 4 To prove that \mathbf{c}, \mathbf{d} commute with \mathbf{f}, \mathbf{g} , simply consider that, by the companion hypothesis, for all $x \in M$, $\langle [x, 0], 0 \rangle = [x, 0]$, from which we get $\mathbf{c} \mathbf{f}(x) = \mathbf{f} \mathbf{c}(x)$, and $\langle [0, x], 0 \rangle = [0, x]$, from which we get $\mathbf{c} \mathbf{g}(x) = \mathbf{g} \mathbf{c}(x)$, and so on.

To prove that \mathbf{c}, \mathbf{d} commute with $\mathbf{f}^*, \mathbf{g}^*$, consider, given a generic $x \in M$, the (unique) decomposition $x = [y, z]$, so

$$\mathbf{f}^* \mathbf{c}(x) = \mathbf{f}^* \mathbf{c}\mathbf{f}(y) + \mathbf{f}^* \mathbf{c}\mathbf{g}(z) = \mathbf{f}^* \mathbf{f}\mathbf{c}(y) + \mathbf{f}^* \mathbf{g}\mathbf{c}(z) = \mathbf{c}(y) = \mathbf{c}\mathbf{f}^*(x),$$

where we have used points 1 and 2 proved above. The other cases are handled similarly, as are the commutations between \mathbf{f}, \mathbf{g} and $\mathbf{c}^*, \mathbf{d}^*$.

To prove that $\mathbf{c}^*, \mathbf{d}^*$ commute with $\mathbf{f}^*, \mathbf{g}^*$, we consider the same decomposition above for the generic $x \in M$, and get

$$\mathbf{f}^* \mathbf{c}^*(x) = \mathbf{f}^* \mathbf{c}^* \mathbf{f}(y) + \mathbf{f}^* \mathbf{c}^* \mathbf{g}(z) = \mathbf{f}^* \mathbf{f}\mathbf{c}^*(y) + \mathbf{f}^* \mathbf{g}\mathbf{c}^*(z) = \mathbf{c}^*(y) = \mathbf{c}^* \mathbf{f}^*(x),$$

and similarly for the other cases.

Assume now that the eight endomorphisms verify the four statements above. The fact that the maps $(x, y) \mapsto \mathbf{c}(x) + \mathbf{d}(y)$ and $(x, y) \mapsto \mathbf{f}(x) + \mathbf{g}(y)$ are homomorphisms from $M \oplus M$ to M is obvious, but we need to prove that they are bijective. We shall do it for the first map, the second being structurally identical.

Suppose that, given $x, x', y, y' \in M$, we have $\mathbf{c}(x) + \mathbf{d}(y) = \mathbf{c}(x') + \mathbf{d}(y')$. Then, applying \mathbf{c}^* (respectively, \mathbf{d}^*) to both sides of the equation and using points 1 and 2, we get $x = x'$ (respectively, $y = y'$), which proves injectivity. For surjectivity, by point 3, for any element $x \in M$, there exist $y, z \in M$ such that $x = \mathbf{c}(y) + \mathbf{d}(z)$: just put $y = \mathbf{c}^*(x)$ and $z = \mathbf{d}^*(x)$. The fact that $\mathbf{c}^*, \mathbf{d}^*$ are the projections associated with this isomorphism is trivial.

We are left to prove that the two isomorphisms are companions. If we put $\langle x, y \rangle = \mathbf{c}(x) + \mathbf{d}(y)$ and $[x, y] = \mathbf{f}(x) + \mathbf{g}(y)$, using point 4 we have

$$\begin{aligned} \langle [w, x], [y, z] \rangle &= \mathbf{c}\mathbf{f}(w) + \mathbf{c}\mathbf{g}(x) + \mathbf{d}\mathbf{f}(y) + \mathbf{d}\mathbf{g}(z) \\ &= \mathbf{f}\mathbf{c}(w) + \mathbf{f}\mathbf{d}(y) + \mathbf{g}\mathbf{c}(x) + \mathbf{g}\mathbf{d}(z) \\ &= \langle w, y \rangle, \langle x, z \rangle, \end{aligned}$$

which completes the proof. □

If $(A, +, 0)$ is a commutative monoid, the set $\mathcal{S}_2(A)$ of all sequences of elements of A indexed by pairs of non-negative integers is an example of an interaction monoid. Addition is defined pointwise, and the neutral element is the everywhere-zero sequence; the eight endomorphisms are defined as follows:

$$\begin{aligned} \mathbf{c}(x)_{m,n} &= \begin{cases} x_{k,n} & \text{if } m = 2k \\ 0 & \text{if } m = 2k + 1 \end{cases} & \mathbf{c}^*(x)_{m,n} &= x_{2m,n} \\ \mathbf{d}(x)_{m,n} &= \begin{cases} 0 & \text{if } m = 2k \\ x_{k,n} & \text{if } m = 2k + 1 \end{cases} & \mathbf{d}^*(x)_{m,n} &= x_{2m+1,n} \\ \mathbf{f}(x)_{m,n} &= \begin{cases} x_{m,k} & \text{if } n = 2k \\ 0 & \text{if } n = 2k + 1 \end{cases} & \mathbf{f}^*(x)_{m,n} &= x_{m,2n} \\ \mathbf{g}(x)_{m,n} &= \begin{cases} 0 & \text{if } n = 2k \\ x_{m,k} & \text{if } n = 2k + 1 \end{cases} & \mathbf{g}^*(x)_{m,n} &= x_{m,2n+1} \end{aligned}$$

From this it is not hard to check that points 1, 2 and 3 of Proposition 4.1 are satisfied. For point 4, just note that $\mathbf{c}, \mathbf{c}^*, \mathbf{d}, \mathbf{d}^*$ and $\mathbf{f}, \mathbf{f}^*, \mathbf{g}, \mathbf{g}^*$ act on separate indexes, so all operations commute.

To get a better understanding of the example, observe that the two isomorphisms defined above build a new sequence by interleaving two sequences: one of them interleaves them ‘horizontally’; the other ‘vertically’. More precisely, our sequences are bidimensional, that is, ‘sequences of sequences’: in the first case we consider them as ‘sequences of columns’, and interleave them horizontally; in the second case, we consider them as ‘sequences of rows’, and interleave them vertically.

Note that the same construction can be applied to almost-everywhere-null bidimensional sequences; when A is countable, this yields a countable interaction monoid.

Interaction monoids are clearly interaction sets (the distinguished element is the zero of the monoid), so Definitions 3.3 and 3.4 can be applied just as they are, yielding a semantics that interprets a net with $n \geq 1$ free ports as a subset of M^n , where $M^n = M \oplus \dots \oplus M$ n times.

Since we are considering monoids, it makes sense to add experiments pointwise, that is, given a net μ and two experiments e_1, e_2 on μ over an interaction monoid M , we can define the function $e_1 + e_2$ from $\text{Ports}(\mu)$ to M that associates the element $e_1(i) + e_2(i)$ with a port i . One may wonder whether this yields another experiment; the answer is indeed positive.

Lemma 4.2 (Additivity). Let μ be a net, and e_1, e_2 be two experiments on μ over an interaction monoid M . Then, $e_1 + e_2$ is an experiment.

Proof. The fact that $e_1 + e_2$ respects conditions (a) and (d) of Definition 3.3 is obvious. Conditions (b) and (c) are consequences of the fact that our companion bijections are monoid isomorphisms. For example, in the case of a δ cell whose auxiliary and principal ports are i, j and k , respectively, we have

$$\begin{aligned} (e_1 + e_2)(k) &= e_1(k) + e_2(k) \\ &= \langle e_1(i), e_1(j) \rangle + \langle e_2(i), e_2(j) \rangle \\ &= \langle e_1(i) + e_2(i), e_1(j) + e_2(j) \rangle \\ &= \langle (e_1 + e_2)(i), (e_1 + e_2)(j) \rangle. \end{aligned}$$

The case of a ζ cell is identical. □

Therefore, if μ is a net with $n \geq 1$ free ports, $\llbracket \mu \rrbracket$ is not just any subset of M^n , it is a *submonoid*.

Corollary 4.3. Let μ be a net with $n \geq 1$ free ports and M be an interaction monoid. Then, the interpretation of μ in M is a submonoid of M^n .

Proof. The result follows from the Additivity lemma (Lemma 4.2) and Proposition 3.1. □

4.2. The GoI semantics

Given an interaction monoid M , we shall now define a semantics that interprets a cut-free net with $n \geq 1$ free ports as an endomorphism of $M^n = M \oplus \dots \oplus M$. This is just a reformulation of what has already been done by Lafont (Lafont 1997), so no proofs will be given in this section.

In the following, we use W to denote the sub-semiring (with unit) of $\text{End}(M)$ generated by $\mathbf{c}, \mathbf{c}^*, \mathbf{d}, \mathbf{d}^*, \mathbf{f}, \mathbf{f}^*, \mathbf{g}, \mathbf{g}^*$.

Definition 4.2 (Weight). Let M be an interaction monoid, μ be a net and p be a straight path of μ (see Definition 2.1). We define the *weight* of p in M , which is an element of W and is denoted $w(p)$, by induction on the length of p :

- If p contains just one port, then $w(p) = \mathbf{1}$ (the identity endomorphism).
- If $p = p' \cdot i$ and the ending port of p' and i do not belong to the same cell, then $w(p) = w(p')$.
- If $p = p' \cdot i$, where p' ends with the left (respectively, right) auxiliary port of a δ cell and i is the principal port of the same δ cell, then $w(p) = \mathbf{c}w(p')$ (respectively, $w(p) = \mathbf{d}w(p')$).
- If $p = p' \cdot i$, where p' ends with the principal port of a δ cell and i is the left (respectively, right) auxiliary port of the same δ cell, then $w(p) = \mathbf{c}^*w(p')$ (respectively, $w(p) = \mathbf{d}^*w(p')$).
- If $p = p' \cdot i$, where p' ends with the left (respectively, right) auxiliary port of a ζ cell and i is the principal port of the same ζ cell, then $w(p) = \mathbf{f}w(p')$ (respectively, $w(p) = \mathbf{g}w(p')$).
- If $p = p' \cdot i$, where p' ends with the principal port of a ζ cell and i is the left (respectively, right) auxiliary port of the same ζ cell, then $w(p) = \mathbf{f}^*w(p')$ (respectively, $w(p) = \mathbf{g}^*w(p')$).

Given a graph-theoretical path, one can always consider its *reversal*, that is, the same path walked from target to source. Notice that the reversal of a straight path is still straight. The unit semiring W can be equipped with an involution $(\cdot)^*$:

- $(\mathbf{c})^* = \mathbf{c}^*$, $(\mathbf{c}^*)^* = \mathbf{c}$, and similarly for the other generators.
- $\mathbf{0}^* = \mathbf{0}$, and for all $u, v \in W$, $(u + v)^* = u^* + v^*$.
- $\mathbf{1}^* = \mathbf{1}$, and for all $u, v \in W$, $(uv)^* = v^*u^*$.

It is then straight-forward to check the following lemma.

Lemma 4.4 (Reversal). Let μ be a net, p be a straight path of μ and p' be the reversal of p . Then, $w(p') = w(p)^*$.

We are now ready to define the GoI interpretation of a cut-free net.

Definition 4.3 (GoI interpretation). Let M be an interaction monoid, ν be a cut-free net with $n \geq 1$ free ports and P_{ji} be the set of straight paths of ν starting from the free port j and ending in the free port i . The GoI interpretation of ν in M is an endomorphism of

M^n , which we represent as a formal $n \times n$ matrix v^\bullet , whose entries are defined as follows:

$$v_{ij}^\bullet = \sum_{p \in P_{ji}} w(p).$$

i and j range over the free ports of v , and the sum is intended to be equal to $\mathbf{0}$ (the everywhere-zero endomorphism) if $P_{ji} = \emptyset$.

If A is a formal matrix with coefficients in W , we can define A^* as the transpose-involute matrix of A : $(A^*)_{ij} = (A_{ji})^*$. Then, the following clearly holds from Lemma 4.4 applied to Definition 4.3.


Proposition 4.5. If v^\bullet is the GoI interpretation of a cut-free net v , then

$$v^{\bullet*} = v^\bullet.$$

We can give a few examples to clarify the definition. If ε, δ and ω are the three nets defined in Section 3.2 after Definition 3.4, we have $\varepsilon^\bullet = \mathbf{0}$, $\delta^\bullet = \mathbf{cd}^* + \mathbf{dc}^*$, while ω^\bullet is the endomorphism of $M \oplus M$ represented by the matrix

$$\omega^\bullet = \begin{bmatrix} \mathbf{0} & \mathbf{1} \\ \mathbf{1} & \mathbf{0} \end{bmatrix}.$$

A slightly more complicated example is the net




whose GoI interpretation is

$$\begin{bmatrix} \mathbf{cg}(\mathbf{cd}^* + \mathbf{dc}^*)\mathbf{g}^*\mathbf{c}^* & \mathbf{df}^* \\ \mathbf{fd}^* & \mathbf{0} \end{bmatrix}.$$

In all cases, one can check that Proposition 4.5 is verified.

One may ask why we have restricted our interpretation to cut-free nets, in sharp contrast to Definition 3.4, where the denotational semantics is defined for *any* net. The reason is quite simple: in the absence of any restriction, Definition 4.3 would not make sense in general, since P_{ji} may contain an infinite number of non-zero-weighting paths. In particular, consider the following example:




There is obviously an infinite number of straight paths from, for instance, the left free port to the right one; their weights are $\mathbf{c}^*(\mathbf{fd}^*)^n \mathbf{g}$, for every $n \in \mathbb{N}$. No element of W can be associated with the sum of all these paths, so the interpretation would be undefined. On the other hand, the following result assures us that our formulation of Definition 4.3 is sound.

Proposition 4.6. Let v be a cut-free net with at least one free port, and let i, j be two free ports (which may be the same) of v . Then, P_{ji} is finite.

Proof. We begin by recalling the general decomposition of a cut-free net with $n \geq 1$ free ports: n trees τ_1, \dots, τ_n with an ε -wiring $\tilde{\omega}$ ‘on top’. Now, a straight path p from port j to port i is necessarily of the following shape: p goes up along one branch of τ_j from the root to its leaf, which is connected through a wire of $\tilde{\omega}$ to a leaf k of τ_i ; p follows this connection, and then goes down the branch of τ_i leading from k to its root. Therefore, the number of straight paths in P_{ji} is bounded by the number of leaves of τ_j , which is of course finite. □

Any net with $n \geq 1$ free ports and k active pairs and/or vicious circles can be decomposed as follows:




where v is cut-free and has $n + 2k$ free ports. Note that, because of the possible presence of vicious circles, v is not unique in general. Nevertheless, given an interaction monoid M and a net μ with $n \geq 1$ free ports and k active pairs and/or vicious circles, we can associate with μ at least one endomorphism μ^\bullet of M^{n+2k} , which is the GoI interpretation of the net v in one of the possible decompositions; the association will be unique exactly when μ does not contain vicious circles.

Theorem 4.8. Let v be a cut-free net with $n \geq 1$ free ports. Then

$$\llbracket v \rrbracket = \text{fix}(v^\bullet),$$

where the interpretations are taken in any interaction monoid M .

Proof. We prove both inclusions by induction on the number m of binary cells of v . If $m = 0$, then v is an ε -wiring, containing n' ε cells and k wires, with $n' + 2k = n$. We can assume without loss of generality that v has the shape




hence $v^\bullet = \sigma_{n',k}$. Now, if $x \in \llbracket v \rrbracket$,

$$x = (\underbrace{0, \dots, 0}_{n'}, x_1, x_1, \dots, x_k, x_k),$$

so we obviously have $v^\bullet(x) = x$. Conversely, it is trivial to check that every fixpoint of $\sigma_{n',k}$ is of this form.

Now let $m > 0$. Since v is cut-free, at least one of its free ports is the principal port of a binary cell. Suppose it is a δ cell; we can assume without loss of generality that v has the shape



where v' is also cut-free. Now, if we order the free ports of v and v' from left to right, we have that their GoI interpretations are given by the following formal matrices:

$$v'^\bullet = \left[\begin{array}{cc|c} h & u^* & A^* \\ u & k & B^* \\ \hline A & B & C \end{array} \right]$$

$$v^\bullet = \left[\begin{array}{cc|c} chc^* + du^*c^* + cud^* + dkd^* & cA^* + dB^* \\ \hline Ac^* + Bd^* & C \end{array} \right].$$

We have used the block notation to represent an arbitrary number (which may even be zero) of entries in the bottom-right part of the matrices, corresponding to the $n - 1$ ports that are free both in v and v' . Here, h, u and k are endomorphisms of W , with $h^* = h$ and $k^* = k$, while A and B are $(n - 1) \times 1$ matrices and C is an $(n - 1) \times (n - 1)$ matrix with entries in W , and $C^* = C$.

We now take $x \in \llbracket v \rrbracket$. We know that x is an element of M^n , so $x = (y, z)$, where $y \in M$ and $z \in M^{n-1}$. Moreover, since y is associated with the free ports of a δ cell, we have $y = \mathbf{c}(y') + \mathbf{d}(y'')$ for some $y', y'' \in M$ such that $x' = (y', y'', z) \in \llbracket v' \rrbracket$. If we apply v'^\bullet to x' , we get

$$\left[\begin{array}{cc|c} h & u^* & A^* \\ u & k & B^* \\ \hline A & B & C \end{array} \right] \cdot \left[\begin{array}{c} y' \\ y'' \\ z \end{array} \right] = \left[\begin{array}{c} h(y') + u^*(y'') + A^*(z) \\ u(y') + k(y'') + B^*(z) \\ \hline A(y') + B(y'') + C(z) \end{array} \right].$$

But by the induction hypothesis, $v'^\bullet(x') = x'$, so the following equalities hold:

$$\begin{aligned} h(y') + u^*(y'') + A^*(z) &= y' \\ u(y') + k(y'') + B^*(z) &= y'' \\ A(y') + B(y'') + C(z) &= z. \end{aligned}$$

From this, if we compute $v^\bullet(x)$, we get

$$\begin{aligned} &\left[\begin{array}{cc|c} \mathbf{c}h\mathbf{c}^* + \mathbf{d}u^*\mathbf{c}^* + \mathbf{c}u\mathbf{d}^* + \mathbf{d}k\mathbf{d}^* & \mathbf{c}A^* + \mathbf{d}B^* & \\ \hline A\mathbf{c}^* + B\mathbf{d}^* & C & \end{array} \right] \cdot \left[\begin{array}{c} \mathbf{c}(y') + \mathbf{d}(y'') \\ z \end{array} \right] \\ &= \left[\begin{array}{c} \mathbf{c}(h(y') + u^*(y'') + A^*(z)) + \mathbf{d}(u(y') + k(y'') + B^*(z)) \\ \hline A(y') + B(y'') + C(z) \end{array} \right] \\ &= x. \end{aligned}$$

Now consider a fixpoint x of v^\bullet . Again, this is an element of M^n and can be decomposed into (y, z) with $y \in M$ and $z \in M^{n-1}$. Now, by surjectivity there exist $y', y'' \in M$ such that $y = \mathbf{c}(y') + \mathbf{d}(y'')$, so we can define $x' = (y', y'', z)$, for which the same computations done above show that $v'^\bullet(x') = x'$. By the induction hypothesis, $x' \in \llbracket v' \rrbracket$, which means that there is an experiment of v' with result (y', y'', z) ; the ‘same’ experiment then gives (y, z) on v , which proves that $x \in \llbracket v \rrbracket$.

The proof in the case of a ζ combinator is identical: we just need to replace \mathbf{c}, \mathbf{d} with \mathbf{f}, \mathbf{g} . □

Theorem 4.8 tells us that, for any cut-free net v with at least one free port, if we know v^\bullet , we also know $\llbracket v \rrbracket$, but is the converse true? The rest of this section is devoted to proving that it is.

Let M be an interaction monoid and v be a cut-free net with $n \geq 1$ free ports. By Theorem 3.7, we know that the elements of $\llbracket v \rrbracket$ are described by a balanced bracket expression \mathbf{B} . From this expression, it is not hard to build an endomorphism ϕ of M^n such that $\text{fix}(\phi) = \llbracket v \rrbracket$:

- Suppose that $\text{var}(\mathbf{B}) = \{x_1, \dots, x_m\}$. We make a new expression \mathbf{B}' (which is no longer balanced) that is structurally identical to \mathbf{B} , but such that, for each variable x_i of \mathbf{B} , one occurrence of x_i is replaced by x'_i and one by x''_i , with x'_i, x''_i distinct and fresh, and

each occurrence of 0 is replaced by a distinct fresh variable z_j ; this is always possible since B is balanced. Notice that B' contains every variable at most once.

- From the fact that $\langle \cdot, \cdot \rangle$ and $[\cdot, \cdot]$ are bijections, we know that, for each $d \in M^n$, there exist unique

$$d'_1, d''_1, \dots, d'_m, d''_m, e_1, \dots, e_k \in M$$

such that

$$d = B' \{ \dots x'_i := d'_i, x''_i := d''_i \dots z_j := e_j \dots \}$$


(see Section 3.3, following Definition 3.5). We then define ϕ by

$$\phi(d) = B' \{ \dots x'_i := d''_i, x''_i := d'_i \dots z_j := 0 \dots \},$$

that is, we ‘swap’ the elements assigned to x'_i and x''_i , and set each z_j to 0.

- Clearly, $\phi(0) = 0$, and because $\langle \cdot, \cdot \rangle, [\cdot, \cdot]$ are isomorphisms, we also have $\phi(x + y) = \phi(x) + \phi(y)$, so ϕ is indeed an endomorphism of M^n (but not an isomorphism, since, in general, some non-zero elements may be mapped to zero). It is not hard to check that ϕ satisfies $\phi^3 = \phi$, that is, it is a *partial symmetry*.
- By construction, the fixpoints of ϕ are those elements described by B , so $\text{fix}(\phi) = \text{Im}(B)$.

The following example will clarify this construction. The balanced expression $([0, x], \langle x, 0 \rangle)$, which generates the interpretation of



is turned into $([z_1, x'], \langle x'', z_2 \rangle)$. Now, for each element $x \in M \oplus M$, there exist unique $z_1, x', x'', z_2 \in M$ such that $x = ([z_1, x'], \langle x'', z_2 \rangle)$ (we have used the same notation for the variables in the expression and the elements of M to avoid writing the substitution explicitly). We then define ϕ so that

$$\phi(x) = \phi([z_1, x'], \langle x'', z_2 \rangle) = ([0, x''], \langle x', 0 \rangle).$$

The zero of $M \oplus M$ is $(0, 0) = ([0, 0], \langle 0, 0 \rangle)$, hence $\phi(0, 0) = (0, 0)$, and if we take $x, y \in M \oplus M$, we decompose them as $x = ([x_1, x_2], \langle x_3, x_4 \rangle)$ and $y = ([y_1, y_2], \langle y_3, y_4 \rangle)$, and we have $x + y = ([x_1 + y_1, x_2 + y_2], \langle x_3 + y_3, x_4 + y_4 \rangle)$, from which we obtain

$$\begin{aligned} \phi(x + y) &= ([0, x_3 + y_3], \langle x_2 + y_2, 0 \rangle) \\ &= ([0, x_3], \langle x_2, 0 \rangle) + ([0, y_3], \langle y_2, 0 \rangle) \\ &= \phi(x) + \phi(y), \end{aligned}$$

so ϕ is an endomorphism of $M \oplus M$. We also have

$$\text{fix}(\phi) = \{ ([0, x], \langle x, 0 \rangle) \in M \oplus M ; x \in M \},$$

which is exactly the interpretation of the above net.

All that is left to do is verify that $\phi = v^\bullet$. This is proved by induction on the number m of binary cells in v . If $m = 0$, v is an ε -wiring, and its interpretation can be assumed without loss of generality to be generated by a balanced expression of the form

$$(0, \dots, 0, x_1, x_1, \dots, x_k, x_k),$$

where the symbol 0 appears n' times, with $n' + 2k = n$ (the number of free ports of v). Then, ϕ is the endomorphism such that

$$\phi(x_1, \dots, x_{n'}, y'_1, y''_1, \dots, y'_k, y''_k) = (0, \dots, 0, y''_1, y'_1, \dots, y''_k, y'_k),$$


which is exactly the endomorphism we introduced following Proposition 4.6 under the name $\sigma_{n',k}$, and which is equal to v^\bullet . If $m > 0$, calculations virtually identical to those of the proof of Theorem 4.8 show that we have $\phi = v^\bullet$ in this case also; the details are left to the reader.

5. Conclusions

Why do we use the symmetric interaction combinators?

One may wonder why we have chosen to work with the *symmetric* combinators instead of the ‘standard’ ones, which enjoy a stronger universality property (Theorem 2.2 holds for *any* interaction net system, not just polarised ones). The answer is technical: there is a detail in the reduction rules of the interaction combinators that means it is impossible to formulate a relational semantics like the one considered here.

Recall that the ‘standard’ interaction combinators are defined in exactly the same way as the symmetric ones, except that instead of ζ there is a binary cell γ , which interacts with itself as follows:




All the other interaction rules are unchanged. Note that this rule ‘exchanges’ the auxiliary ports of the γ cell: according to our convention (at the beginning of Section 2.1), the left port of each occurrence of γ is connected to the right port of the other occurrence, unlike the $\delta\delta$ (and $\zeta\zeta$) annihilation, which connects left with left and right with right.

Now, in a relational semantics, reduction is modelled by composition of relations: if the ‘rightmost’ free port of a net μ is connected to the ‘leftmost’ free port of a net μ' , the denotational semantics of the resulting net will be (see the proof of Lemma 3.3)


$$\llbracket \mu \rrbracket \circ \llbracket \mu' \rrbracket = \{(a, c) ; (a, b) \in \llbracket \mu \rrbracket, (b, c) \in \llbracket \mu' \rrbracket\}.$$

This is ensured by our definition of an experiment. But if we try to define experiments in the presence of the $\gamma\gamma$ annihilation, we see that the only way for the interpretation to

model reduction is for both of the auxiliary ports of γ cells to receive the same value. In fact, in the rule



we clearly need $a = b$. This is an unreasonable restriction; for example, it would imply that the following two nets receive the same semantical interpretation:



These two nets are not $\beta\eta$ -equivalent; from Proposition 2.6, we infer that in such a situation, if we ever managed to model \simeq_β , we would do so by identifying all total nets with a non-empty interface.

Of course, the argument given here does not rule out the possibility of finding a denotational semantics for the interaction combinators; it simply shows that the standard definitions do not work, and justifies our shift to the symmetric combinators.

Further work

Our efforts are now concentrated on producing a *typed* semantics for the symmetric combinators. This would not only yield a typing discipline for the combinators, which would ensure good properties, such as deadlock-freeness and termination, but, more interestingly, a new logical system, which should be an extension of multiplicative linear logic.

The existence of such a system would be very intriguing, since it would combine the simplicity of the multiplicative fragment of linear logic with the high expressive power of the symmetric combinators.

References

Danos, V. and Regnier, L. (1989) The structure of multiplicatives. *Archive for Mathematical Logic* **28** 181–203.
 Ehrhard, T. (2005) Finiteness spaces. *Mathematical Structures in Computer Science* **15** (4) 615–646.
 Fernández, M. and Mackie, I. (2003) Operational equivalence for interaction nets. *Theoretical Computer Science* **297** (1–3) 157–181.

- Girard, J.-Y. (1987a) Linear logic. *Theoretical Computer Science* **50** (1) 1–102.
- Girard, J.-Y. (1987b) Multiplicatives. In: Lolli, G. (ed.) *Logic and Computer Science: New Trends and Applications*, Rendiconti del Seminario Matematico dell'Università e Politecnico di Torino 11–34.
- Girard, J.-Y. (1989) Geometry of interaction I: interpretation of System F. In: *Proceedings of the Logic Colloquium '88*, North Holland 221–260.
- Gonthier, G., Abadi, M. and Lévy, J.-J. (1992) The geometry of optimal lambda reduction. In: *Conference Record of the Nineteenth ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 92)*, ACM Press 15–26.
- Lafont, Y. (1990) Interaction nets. In: *Conference Record of POPL'90*, ACM Press 95–108.
- Lafont, Y. (1995) From proof nets to interaction nets. In: Girard, J.-Y., Lafont, Y. and Regnier, L. (eds.) *Advances in Linear Logic*, Cambridge University Press 225–247.
- Lafont, Y. (1997) Interaction combinators. *Information and Computation* **137** (1) 69–101.
- Lippi, S. (2002) Encoding left reduction in the lambda-calculus with interaction nets. *Mathematical Structures in Computer Science* **12** (6) 797–822.
- Mackie, I. and Pinto, J. S. (2002) Encoding linear logic with interaction combinators. *Information and Computation* **176** (2) 153–186.
- Mazza, D. (2006) Observational equivalence for the interaction combinators and internal separation. In: Mackie, I. (ed.) Proceedings of TERMGRAPH 2006. *Electronic Notes in Theoretical Computer Science* 7–16.
- Pagani, M. (2007) Proofs, denotational semantics and observational equivalence in Multiplicative Linear Logic. *Mathematical Structures in Computer Science* **17** (2) 341–359.