
OBSERVATIONAL EQUIVALENCE AND FULL ABSTRACTION IN THE SYMMETRIC INTERACTION COMBINATORS

DAMIANO MAZZA *

LIPN – UMR 7030 CNRS - Université Paris 13, F-93430 Villetaneuse, France
e-mail address: Damiano.Mazza@lipn.univ-paris13.fr

ABSTRACT. The symmetric interaction combinators are an equally expressive variant of Lafont’s interaction combinators. They are a graph-rewriting model of deterministic computation. We define two notions of observational equivalence for them, analogous to normal form and head normal form equivalence in the lambda-calculus. Then, we prove a full abstraction result for each of the two equivalences. This is obtained by interpreting nets as certain subsets of the Cantor space, called edifices, which play the same role as Böhm trees in the theory of the lambda-calculus.

INTRODUCTION

A foundational study of interaction nets. Lafont’s interaction nets [Laf90] are a powerful and versatile model of deterministic computation, derived from the proof nets of Girard’s linear logic [Gir87, Gir96, Laf95]. Interaction nets are characterized by the atomicity and locality of their rewriting rules. As in Turing machines, computational steps are elementary enough to be considered as constant-time operations, but, unlike Turing machines, several steps can be executed in parallel, *i.e.*, interaction nets actually model a kind of distributed computation.

Several interesting applications of interaction nets exist. The most notable ones are implementations of optimal evaluators for the λ -calculus [?, Mac04], but efficient evaluation of other functional programming languages using richer data structures is also possible with interaction nets [Mac05].

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 [Laf97] and Fernández and Mackie’s work on operational equivalence [FM03], no foundational study of interaction nets can be found in the existing literature. For example, until very recently [Maz07a], no denotational semantics had been proposed for interaction nets.

1998 ACM Subject Classification: F.3.2, F.4.1.

Key words and phrases: Interaction nets, observational equivalence, denotational semantics, full abstraction, Böhm trees.

* Partially supported by a post-doctoral fellowship of the *Fondation Sciences Mathématiques de Paris*.

This work aims precisely at studying and expanding the theory of interaction nets, in particular of the symmetric interaction combinators. These latter are especially interesting because of their *universality*: any interaction net system can be translated in the symmetric interaction combinators [Laf97]. Therefore, within the graph-rewriting paradigm given by interaction nets, the symmetric interaction combinators stand out as a prototypical language, just like the λ -calculus is the prototypical language of the functional paradigm.

More specifically, the contribution of this work is twofold, and can be seen as a methodology for addressing the following two questions:

Observational equivalence: Given two nets in the system of the symmetric interaction combinators, when can we say that they behave in the same way?

Denotational semantics and full abstraction: Any answer to the above question yields an equivalence on nets; can we denotationally characterize this equivalence? In other words, can we find an abstract interpretation of the syntax so that such equivalence becomes an equality?

Observations and contexts. The first question is a central one in all programming languages. Indeed, any programmer is aware that, given two syntactically different programs, it may as well be that they “do the same thing”, *i.e.*, that one can be replaced by the other without anyone noticing the difference. Of course, the heart of the question lies in what differences we judge worth noticing: the final result of executing a program, the time it takes to obtain such result, etc. Different choices will of course lead to different notions of “doing the same thing”.

In any case, the key notion is that of *observation*: the program interacts with an environment, and we observe the outcome, based on our choice of what we consider relevant to be observed. If our two programs yield the same observations upon interacting with all possible environments, it is fair to say that, as far as we are concerned, they “do the same thing”; more formally, we say that they are *observationally equivalent*, according to our chosen notion of observation.

For functional programming languages, and in particular for the λ -calculus, Morris [Mor68] was the first to propose the now widely accepted idea that *an environment is a context*. In this way, observations are internalized, *i.e.*, they may be made directly on programs, because, given a program P and a context C , $C[P]$ is still a program.

In interaction nets, there is a very natural notion of context for a net μ : it is simply another net C whose interface is big enough so that μ can be plugged into it, forming a new net $C[\mu]$. Morris’ idea can therefore be straightforwardly applied in our framework.

Internal separation and observable axioms. In order to choose what to observe in our nets, we draw inspiration from our previous work on internal separation [Maz07b], in which we proved a result similar to the celebrated Böhm’s theorem for the λ -calculus [Böh68]. Böhm’s theorem states the following: given two $\beta\eta$ -normal λ -terms T, U , $T \neq U$ implies that there exists a context C such that $C[T] \rightarrow_{\beta}^* x$ and $C[U] \rightarrow_{\beta}^* y$, where x and y are two different variables.

A consequence of Böhm’s result is that *it is impossible to equate two distinct $\beta\eta$ -normal forms, unless one equates all λ -terms*. This makes us better understand the importance of Böhm’s theorem, because it brings forth its negative content: as underscored for example by Giuseppe Longo [Lon05], negative results are crucial in the development of a theory,

since they witness the presence of a structure in the underlying objects. If “everything is possible”, then the objects of our theory are shapeless, we can tamper with them at will, and the theory loses any scientific interest.

In the case of the symmetric interaction combinators, internal separation (*cf.* Theorem 1.18) cannot be realized using two arbitrary nets (by contrast, in the λ -calculus, x and y may be replaced by two arbitrary distinct λ -terms). Indeed, one of the two nets used contains a special kind of connection, which we call *observable axiom*, while the other does not. Since identifying these two nets induces the identification of all nets, we are led to take the presence of observable axioms as the key phenomenon to observe.

Axiom-equivalences. The discovery of observable axioms in interaction nets is the backbone of the development of our theory of observational equivalence. Indeed, we have several results hinting to the fact that observable axioms are analogous to *head variables*; these occupy an arguably important place in the theory of the λ -calculus, so it is perhaps not surprising that we give observable axioms a central role in interaction nets too.

Furthermore, observable axioms are related to a certain kind of paths of Girard’s geometry of interaction [Gir89], as reformulated by Lafont for the (symmetric) interaction combinators [Laf97]. In particular, it is possible to show (*cf.* Sect. 2.2) that the observable axioms generated by a net in the course of its reduction correspond to its *execution paths* [DR95]; these are the paths which are preserved by reduction, and are hence present in every reduct. In some sense, each execution path describes a portion of information produced by the computation of a net; in particular, if a net is normalizable, then its execution paths describe exactly its normal form. This idea of approximation, which is also already present in Böhm trees, is another way of looking at observable axioms as meaningful objects to study the behavior of a net (*cf.* Sect. 5.2).

We thus introduce *observable nets* and *finitarily observable nets*: the first are nets which, in the course of their reduction, develop at least one observable axiom; the second are observable nets which develop only finitely many of them. It is useful to keep in mind an analogy with the λ -calculus: observable nets are similar to λ -terms having a head normal form, and finitarily observable nets are akin to normalizable λ -terms. In the first case, we may additionally introduce the notion of *solvable net*, and prove that solvable and observable nets coincide, just like solvable λ -terms coincide with λ -terms having a head normal form. In the second case, the correspondence is somewhat looser, because the symmetric combinators already have a notion of normalizable net, and it does not coincide with that of finitarily observable net. Nevertheless, there are several facts supporting this analogy.

The notions of observable and finitarily observable net can be used to define two observational equivalences on nets: *axiom-equivalence* and *finitary axiom-equivalence*. The first one is similar to head normal form equivalence (hnf-equivalence) in the λ -calculus (two λ -terms T, U are hnf-equivalent iff, for every context C , $C[T]$ is head-normalizable iff $C[U]$ is). The second one is similar to normal form equivalence (nf-equivalence) in the λ -calculus (two λ -terms T, U are nf-equivalent iff, for every context C , $C[T]$ is normalizable iff $C[U]$ is). By “similar” we mean that finitary axiom-equivalence is strictly included in axiom-equivalence, as nf-equivalence is strictly included in hnf-equivalence in the λ -calculus, and that the examples proving strict inclusion are all related to a phenomenon similar to infinite η -expansion [Wad76], as is the case for the λ -calculus. Moreover, after transporting from the λ -calculus to the symmetric interaction combinators the concepts of *theory* and *sensible*

theory (cf. Sect. 5.3), axiom-equivalence can be shown to be a maximal consistent theory, indeed the greatest consistent sensible theory, just like hnf-equivalence.

Our axiom-equivalences are not the only existing observational equivalences for the symmetric interaction combinators; in particular, Fernández and Mackie [FM03] proposed another notion of observational equivalence, based on *visible nets*. This equivalence, which seems to correspond to *weak* head normal form equivalence (whnf-equivalence) in the λ -calculus, can be proved to be strictly stronger, *i.e.*, more discriminative, than our finitary axiom-equivalence, and hence than axiom-equivalence (cf. Sect. 5.1). This is in accord with the λ -calculus analogy: whnf-equivalence is strictly included in nf- and hnf-equivalence.

Equivalence as equality. The second part of our work starts with the development of a denotational semantics of the symmetric interaction combinators. Denotational semantics originated in the late 1960's with the work of Scott and Strachey [SS71, Sco76]. Its goal is to model the syntax of a programming language by means of a more abstract mathematical structure, on which a broader range of tools and proof techniques are available. In this way, one may be able to prove results about the language which would be very difficult, or even impossible, to prove by syntactic methods only.

In a nutshell, we could say that the ultimate goal of denotational semantics is *to transform equivalences into equalities*. A typical example is precisely that of observational equivalence, as discussed above. If a denotational semantics gives the same interpretation to two programs exactly when they are observationally equivalent, then it is said to be *fully abstract* with respect to the given observational equivalence. Finding a fully abstract denotational semantics can be a very hard problem: a notable example is that of PCF, a λ -calculus-like functional language for which completely new game-semantic models had to be developed to achieve full abstraction [AJM00, HO00].

In the λ -calculus, both nf- and hnf-equivalence have been abstractly characterized in several different ways: Hyland [Hyl76] proved that two terms are nf-equivalent iff their Böhm trees are equal up to η -equivalence, and went on to prove that nf-equivalence coincides with equality in Plotkin's $P\omega$ model [Plö72]; Wadsworth [Wad76] obtained similar results for hnf-equivalence, showing that two terms are hnf-equivalent iff their Böhm trees are equal up to *infinite* η -expansion, and that this equivalence corresponds to equality in Scott's D_∞ model [Sco76]. Shortly after, Nakajima [Nak75] introduced a similar characterization of hnf-equivalence in terms of what are now known as Nakajima trees.

Edifices, the Cantor topology, and full abstraction. Besides the description of a new theory of observational equivalence for interaction nets, the other principal contribution of the present work is the introduction of *edifices*, which play the same role as Böhm or Nakajima trees, in that they provide a fully abstract model of the two axiom-equivalences mentioned above.

The starting point for defining edifices is the same as that of Böhm trees, reflecting the analogy between observable axioms and head variables: just like the Böhm tree of a λ -term T is basically the collection of the head variables appearing in the reducts of T (with the additional information concerning their hierarchical structure and the abstractions preceding them), the edifice of a net μ is built from the collection of all observable axioms appearing in the reducts of μ .

Nevertheless, the parallelism of interaction nets, unmatched in the λ -calculus, induces some fundamental differences between the two constructions: in fact, apart from collecting the information concerning the position of observable axioms within the net in which they appear (analogous to the abstractions preceding a head variable), no evident hierarchical structure emerges for observable axioms (although some kind of structure might be attached to them, as briefly discussed in Sect. 5.4). This is why edifices are not at all trees. Still, just like the Böhm tree of a λ -term, the edifice of a net is an invariant of reduction in the symmetric interaction combinators (Proposition 3.19).

To achieve full abstraction, we endow edifices with a topological structure, which turns them into subsets of the Cantor space. In the case of finitary axiom-equivalence, this is needed for technical purposes: in fact, edifices characterize this equivalence *as plain sets*, *i.e.*, two nets are finitarily axiom-equivalent iff their edifice is the same, independently of any topology attached to it; however, the only way we are able to prove this is through a topological property, namely the compactness of the edifices interpreting a certain class of nets (Proposition 4.3). In the case of axiom-equivalence, topology plays a more fundamental role: in fact, we prove that two nets are axiom-equivalent iff the *topological closure* of their edifices is the same; obviously, the notion of closure is meaningless without referring to a topology. This last result is particularly nice, because the phenomenon of infinite η -expansion (which, as mentioned above, is also present in the symmetric interaction combinators) receives a precise topological explanation.

Another nice aspect of edifices is that they are quite interesting in their own right, independently of the symmetric interaction combinators. In fact, several results in our theory hold for wider classes of edifices than those which interpret nets. In particular, there is a notion of *trace* defined on edifices (Sect. 3.1), which is completely general, and which reminds of the notion of composition of strategies in games semantics [AJM00, HO00]. When applied to the special case of edifices which interpret nets, the trace can be seen as an extension of the execution formula of the geometry of interaction, which works in all cases (*cf.* Proposition 3.16); by contrast, Girard's original execution formula, and its rephrasing developed by Lafont for the symmetric interaction combinators, is only defined under certain normalizability assumptions.

Acknowledgments. Many thanks to the anonymous referees for their useful comments and suggestions, and a special thanks to the editor Simona Ronchi della Rocca for her patience in waiting for the revised version of this paper.

1. THE SYMMETRIC INTERACTION COMBINATORS

1.1. **Nets.** The symmetric interaction combinators, or, more simply, the symmetric combinators, are an interaction net system [Laf90, Laf97]. An interaction net is built out of *cells* and *wires*. Each cell has a number of *ports*, exactly one of which is *principal*, the other being *auxiliary*. In the case of the symmetric combinators, there are three kinds of cells: cells of type δ and ζ , which have two auxiliary ports, numbered by the integers 1 and 2, and cells of type ε , which have no auxiliary ports. Cells of the first two kinds are called *binary*, while those of the latter kind are called *nullary*.




Figure 1: A net.

Each wire has two extremities; each extremity may be attached to the port of a cell, so we can use wires to connect cells together. We also allow *loops*, which are wires whose extremities are attached one to the other. Wires which are not loops are called *proper*.

A *net* is any configuration of cells and wires, such that each port of each cell is attached to the extremity of a wire. Note that a net may contain wires with one or both extremities not attached to any port of any cell; these unattached extremities will be called the *free ports* of the net.

Nets are usually presented in graphical form, as in Fig. 1. Binary cells are represented by triangles, nullary cells by circles; in both cases, the symbol denoting the kind of cell is written inside the figure representing it. For a binary cell, the principal port is depicted as one of the “tips” of the triangle representing it. The numbering of the auxiliary ports of binary cells is assigned clockwise: in particular, the auxiliary port number 1 of a binary cell is the left one if the cell is drawn with its principal port pointing towards the bottom of the picture, and it is the right one if the cell is drawn with its principal port pointing “up”. Wires and loops are represented as... wires and loops, and the free ports appear as extremities of “pending” wires. For example, the net in Fig. 1 has 11 cells, of which 4 nullary, 1 loop, 16 proper wires, and 7 free ports.

The above description is precise enough to develop the rest of the paper, and almost all of the theory of interaction nets. However, a more formal definition can be given, by considering an interaction net as the union of two structures: a labelled, directed hypergraph, and an undirected graph. The idea is that labelled and directed hyperedges correspond to cells, and undirected edges to wires. In what follows, we fix a denumerably infinite set of *ports*, which we assume contains the positive integers.

Definition 1.1 (Wire, cell, net). A *wire* is a set of ports of cardinality 1 or 2; in the first case, we speak of a *loop*, in the second case of a *proper wire*. We fix three *symbols* $\delta, \varepsilon, \zeta$; we say that δ and ζ are binary, while ε is nullary. A *cell* is a tuple $(\alpha, p_0, p_1, \dots, p_n)$ where α is a symbol, p_0, p_1, \dots, p_n are ports, and $n = 2$ if α binary, or $n = 0$ if α is nullary. In both cases, p_0 is the *principal port* of the cell, while p_1, \dots, p_n are the *auxiliary ports*.

A *net* μ is a couple $(\text{Cells}(\mu), \text{Wires}(\mu))$, where $\text{Cells}(\mu)$ is a finite set of cells and $\text{Wires}(\mu)$ is a finite set of wires, satisfying the following:

- each port appears at most twice in $\text{Cells}(\mu) \cup \text{Wires}(\mu)$;
- if a port appears in $\text{Cells}(\mu) \cup \text{Wires}(\mu)$, then it appears in exactly one wire.

The set of ports appearing in μ is denoted by $\text{Ports}(\mu)$. A port appearing only once in $\text{Cells}(\mu) \cup \text{Wires}(\mu)$ is called *free*; the set of all free ports of μ is referred to as its *interface*. We shall always assume that if a net has n free ports, then its interface is $\{1, \dots, n\}$.




Figure 2: A generic active pair (left) and an example of vicious circle (right).

Definition 1.2 (α -equivalence). A *renaming* for a net with n free ports is an injective function from ports to ports which is the identity on $\{1, \dots, n\}$. Two nets are α -equivalent iff they are equal modulo a renaming.


Nets are always considered modulo α -equivalence. In fact, observe that graphical representations equate exactly α -equivalent nets. As in most of the existing literature on interaction nets, we shall preferentially disregard Definition 1.1, in favor of more intuitive graphical notations. This is especially convenient for treating the dynamic aspects of nets, such as reduction (*cf.* Sect. 1.2); however, for static aspects, it is sometimes quite convenient to use Definition 1.1, because it gives succinct, formal descriptions of the components of a net (*e.g.* ports, wires as sets of ports, etc.).

Let us introduce some remarkable nets, which will be useful in the sequel:

Wirings: A net containing no cell and no loop is called a *wiring*. Wirings are permutations of free ports; they are ranged over by ω . We shall often use ω also to denote a single wire.

E-nets: The ε -net with n free ports, denoted by \mathbf{E}_n , is the net consisting of n ε cells;

Trees: A *tree* is a net defined by induction as follows. A single ε cell is a tree with no leaf, denoted by ε ; a proper wire is a tree with one leaf (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 α is a binary symbol, the net



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

As the reader may have noticed, in the above picture we represented trees adopting the same graphical notations as cells. We shall avoid possible ambiguities by never using $\delta, \varepsilon, \zeta$ to denote trees, and by using α, β exclusively to range over cell symbols, so that a triangle annotated with α or β will unambiguously represent a cell (or a tree consisting of a single cell, if the reader prefers).

Active pairs: An *active pair* (Fig. 2, left) is a net consisting of two cells whose principal ports are connected by a wire.

Vicious circles:— A *vicious circle* is either a loop, or a net 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 . An example is given in Fig. 2 (right).

It is also useful to identify two special sorts of wires in nets:

Definition 1.3 (Axiom, cut, cut-free net). Let μ be a net, and $\omega = \{p, q\}$ a wire of μ .

Proper axiom: We say that ω is a *proper axiom* if it is a proper wire and none of p, q is the principal port of a cell of μ .

Proper cut: We say that ω is a *proper cut* if it is a proper wire and both p and q are principal ports of cells of μ .


Axiom-cut: We say that ω is an *axiom-cut* if it is a loop, or if μ contains a tree τ such that p is the root of τ , and q one of its leaves.

An *axiom* (resp. *cut*) of μ is either a proper axiom (resp. proper cut) or an axiom-cut; in the latter case, we refer to it as an *improper axiom* (resp. *improper cut*). We say that μ is *cut-free* if it contains no cuts.

As an example, consider the net in Fig. 1, in which the reader should find 7 proper axioms, 2 proper cuts, and 2 axiom-cuts. Note that proper cuts are in one-to-one correspondence with active pairs. On the other hand, axiom-cuts are in many-to-one correspondence with vicious circles, *i.e.*, an axiom-cut implies the presence of exactly one vicious circle, but a vicious circle implies the presence of at least one axiom-cut. Although the correspondence is one-to-one for the net in Fig. 1, we have for instance that, despite having a single vicious circle, the net on the left in Fig. 2 contains 4 axiom-cuts.

The following gives us a general understanding of the structure of cut-free nets:

Lemma 1.4 (Canonical form of a cut-free net). *Let ν be a cut-free net with n free ports. Then, for each $1 \leq i \leq n$ there exist a unique tree τ_i , and there exists a unique wiring ω such that*



Proof. By induction on the number of cells of ν . □

Note the wiring drawn as a rectangle; in the sequel, this graphical notation will be used also to represent generic nets, but ω will always denote a wiring. Observe that all wires in ω are proper axioms; in fact, the above is the shape of a generic cut-free multiplicative proof net, with the axiom links in ω and the logical links in τ_1, \dots, τ_n , whence our terminology.

A fundamental notion for developing the rest of the paper is that of context:

Definition 1.5 (Context, test, feedback). Let μ be a net with n free ports. A *context* for μ is a net C with at least n free ports. We denote by $C[\mu]$ the *application* of C to μ , which is the net obtained by plugging the free port i of μ to the free port i of C , with $i \in \{1, \dots, n\}$. A *test* for μ is a particular context consisting of n trees τ_1, \dots, τ_n such that the root of each τ_i is the free port i . A *feedback context* for μ is a context σ consisting of a wiring connecting some of the free ports of μ between them.

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




Figure 3: The interaction rules: annihilation (left) and commutation (right). In the annihilation, the right member is empty in case $\alpha = \varepsilon$.

can concisely formulate a decomposition result which, combined with Lemma 1.4, uncovers the structure of a generic net:

Lemma 1.6 (Decomposition). *Let μ be a net. Then, there exists a cut-free net ν and a feedback context σ such that $\mu = \sigma[\nu]$.*

Proof. Simply let σ contain all the proper cuts of μ , plus one axiom-cut for each vicious circle of μ , and let ν be the subnet of μ obtained by removing σ . \square

Observe that the net ν of Lemma 1.6 is unique as soon as μ does not contain vicious circles (or, equivalently, axiom-cuts).

1.2. β -reduction and η -equivalence. The dynamics of interaction nets is based on the contextual rewriting of active pairs, which are thus analogous to redexes in the λ -calculus. In the case of the symmetric combinators, the active pairs are rewritten according to the *interaction rules* of Fig. 3: the annihilations, concerning the interaction of two cells of the same type, and the commutations, concerning the interaction of two cells of different type.


Definition 1.7 (β -reduction and β -equivalence). *β -reduction* is the reflexive-transitive closure of the relation defined as follows: given two nets μ, μ' , we set $\mu \rightarrow_{\beta} \mu'$ iff there exists C such that $\mu = C[\mu_0]$, $\mu' = C[\mu'_0]$, and μ_0, μ'_0 match the left and right members of one of the rules of Fig. 3, respectively. We define $\mu \simeq_{\beta} \nu$ iff there exists o such that $\mu \rightarrow_{\beta}^* o$ and $\nu \rightarrow_{\beta}^* o$.

Proposition 1.8 (Strong confluence). *If $\mu \rightarrow_{\beta} \mu'$ and $\mu \rightarrow_{\beta} \mu''$ with $\mu' \neq \mu''$, there exists ν such that $\mu' \rightarrow_{\beta} \nu$ and $\mu'' \rightarrow_{\beta} \nu$. Hence, the relation \rightarrow_{β}^* is confluent, and \simeq_{β} is an equivalence relation.*

Proof. Immediate: there are no critical pairs, because active pairs are always disjoint. \square


We now give a few basic results concerning β -reduction. The first two are generalizations of the annihilation and commutation rules:

Lemma 1.9. *Let τ be a tree. Then, we have*



Proof. By induction on the structure of τ . □


Lemma 1.10. *Let α, β range over binary symbols, with $\alpha \neq \beta$. Let A be a tree not containing β cells, and let B be a tree not containing α cells. Then, we have*



Proof. By double induction on the structures of A and B . □


The following is an easy corollary of Lemmas 1.4, 1.9 and 1.10:

Lemma 1.11 (Duplication). *Let α be a binary symbol, let ν be a cut-free net containing no α cell, and let τ be a tree containing only α cells. Then, we have*



Observe that the only cut-free net with an empty interface is the empty net. Then, the next result shows, as a special case, that cut-free nets can be freely erased:

Lemma 1.12 (Erasing). *Let ν be a cut-free net, and let μ be any net obtained from ν by plugging any number of ε cells to its free ports, as follows:*



Then, there exists a cut-free net ν' such that $\mu \rightarrow_{\beta}^ \nu'$.*

Proof. An immediate consequence of Lemmas 1.4 and 1.10. □

We now introduce η -equivalence, which is similar to the homonymous relation in the λ -calculus, with an essential difference: in the symmetric combinators η -equivalence cannot be presented as the symmetrization of a rewriting relation, like β -equivalence. In fact, one of the equations defining it (namely the η_1 equation applied to binary cells, *cf.* Fig. 4) cannot be meaningfully oriented and transformed into a rewriting step. The η_1 equation was already known to Lafont [Laf97]; the η_0 equation was introduced by Fernández and Mackie [FM03].

Definition 1.13 (η - and $\beta\eta$ -equivalence). We define the relations \simeq_{η_0} and \simeq_{η_1} as the reflexive, transitive, and contextual closure of respectively the left and right equation of Fig. 4, which we call η_0 and η_1 equation, respectively. Then, we define η - and $\beta\eta$ -equivalence respectively as $\simeq_{\eta} = (\simeq_{\eta_0} \cup \simeq_{\eta_1})^+$ and $\simeq_{\beta\eta} = (\simeq_{\beta} \cup \simeq_{\eta})^+$.

The following results are the counterparts of Lemmas 1.9 and 1.10 for η -equivalence.






Figure 4: The equations defining η -equivalence. In both equations, α, β range over cell symbols. In the left equation, α is binary; in the right equation, $\alpha \neq \beta$.

Lemma 1.14. *Let τ be a tree without ε cells. Then, we have*



Proof. By induction on the structure of τ . □


Lemma 1.15. *Let α, β range over binary symbols, with $\alpha \neq \beta$. Let A be a tree not containing β cells, and let B be a tree not containing α cells. Then, we have*



Proof. By double induction on the structure of A and B . □

An easy corollary of Lemma 1.14 is that, modulo η -equivalence, the trees rooted at the free ports of a net in the decomposition given by Lemmas 1.4 and 1.6 can “look like” almost anything we want:

Lemma 1.16. *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 1.14. □


When oriented from right to left, the η_0 equation is formally identical to the standard η -expansion rule in multiplicative proof nets (see for example the work of Pagani [Pag07]), except that it may be applied to any wire in a net, while η -expansion in proof nets concerns only axioms. Actually, it is possible to show that, when combined with β -equivalence, the η_0 equation still generates $\beta\eta$ -equivalence even if its application is limited to axioms.

Lemma 1.17. *Let $\mu \rightarrow_{\eta_0}^- \nu$ iff $\mu = C[\mu_0]$ and $\nu = C[\omega]$, where μ_0 is a net matching the left member of the η_0 equation of Fig. 4, and ω is an axiom of ν . Let $\simeq_{\eta_0}^-$ be the reflexive-transitive closure of the symmetric closure of $\rightarrow_{\eta_0}^-$, and let $\simeq_{\eta}^- = (\simeq_{\eta_0}^- \cup \simeq_{\eta_1})^+$. Then, $(\simeq_{\beta} \cup \simeq_{\eta}^-)^+ = \simeq_{\beta\eta}$.*


Proof. Let us set $\simeq_{\beta\eta}^- = (\simeq_{\beta} \cup \simeq_{\eta}^-)^+$. The inclusion $\simeq_{\beta\eta}^- \subseteq \simeq_{\beta\eta}$ is obvious, because by definition $\simeq_{\eta}^- \subseteq \simeq_{\eta}$. For what concerns the reverse inclusion, it is enough to show that if ν is obtained from μ by a single application of the η_0 equation on a wire which is not an axiom, then $\mu \simeq_{\beta\eta}^- \nu$. We can assume without loss of generality that such a wire is in ν . Then, we must have $\mu = C[\mu_0]$ and $\nu = C[\tau]$, where τ is a tree and




We need to prove that $C[\mu_0] \simeq_{\beta\eta}^- C[\tau]$. If $\tau = \varepsilon$, we leave it to the reader to check that $C[\mu_0] \rightarrow_{\beta} C[\mu'_0] \simeq_{\eta_1} C[\varepsilon]$. Otherwise, we can assume τ to be “maximal”, *i.e.*, its leaves are either free or connected to an auxiliary port. In fact, this is not possible only if one of the leaves of τ is connected to the principal port of the α cell of μ_0 shown at the bottom of the picture; but in this case the wire in ν obtained after applying the η_0 equation would be an axiom, against our hypothesis. Now, τ can always be decomposed as follows:



where B is a tree not containing α cells (we may have $B = \bullet$), and $\tau'_1, \dots, \tau'_k, \tau''_1, \dots, \tau''_k$ are trees (k may be equal to zero). By Lemma 1.10, we have



By Lemma 1.15, we have



Now, if we apply the η_0 equation to the nets on the upper right of the above picture, the resulting wires will be axioms, because of the “maximality” of τ . Hence, $C[\mu'_0] \simeq_{\eta}^- C[\tau]$, as desired. \square

As recalled in the introduction, a fundamental result due to Böhm [Böh68] implies that no non-trivial congruence on λ -terms may equate two distinct $\beta\eta$ -normal forms. In the symmetric combinators there is a similar result [Maz07b], except that one cannot speak of $\beta\eta$ -normal forms, because, as discussed above, the symmetric combinators lack a notion of η -reduction.

There is actually a deeper difference between the symmetric combinators and the λ -calculus, given by the existence of vicious circles. Observe that such configurations are stable under β -reduction, because cells can interact only through their principal port: they are sort of deadlocks. Although diverging computations certainly exist in the λ -calculus, deadlocks are something completely new. Because of this, the notion of normalizable net (which, thanks to strong confluence, is the same as that of strongly normalizable net) does not play a central role in the theory of the symmetric combinators. Instead, cut-free nets are closer to a concept of “true” normal form: a net having no cut-free form represents either a diverging or an error-bound computation, *i.e.*, one that generates deadlocks. Other interesting notions of convergence will be introduced in Sect. 2, but none of them will coincide with simple normalization.

In the following, we say that a net is *total* if it β -reduces to a cut-free net.

Theorem 1.18 (Separation [Maz07b]). *Let μ, ν be two total nets with the same interface, such that $\mu \not\approx_{\beta\eta} \nu$. Then, there exists a test θ such that*


$$\theta[\mu] \rightarrow_{\beta}^* \text{---} \quad \theta[\nu] \rightarrow_{\beta}^* \begin{array}{c} \varepsilon \\ | \\ \varepsilon \end{array} \quad \begin{array}{c} \varepsilon \\ | \\ \varepsilon \end{array}$$

or vice versa. □

The net on the right in Theorem 1.18 is the one we denoted by \mathbf{E}_2 ; we denote the other, *i.e.*, a single wire, by ω . The following makes us understand the strength of the Separation Theorem:

Proposition 1.19. *Let \sim be a congruence on nets, such that $\simeq_\beta \subseteq \sim$. Then, $\omega \sim \mathbf{E}_2$ implies that, for all μ, ν with the same interface, $\mu \sim \nu$.*

Proof. Let μ be a net with n free ports. By the Decomposition Lemma 1.6, we have $\mu = \sigma[\nu]$ for some feedback context σ and some cut-free net ν . Since \sim is a congruence, we can write



But by Lemma 1.12, the net on the right β -reduces to \mathbf{E}_n , where \mathbf{E}_n is a net with n free ports containing n ε cells. Since \sim contains β -reduction, we have proved that $\mu \sim \mathbf{E}_n$ for any μ with n free ports, and we may conclude by symmetry and transitivity of \sim . □


Therefore, the Separation Theorem implies that any non-trivial congruence containing β -equivalence cannot equate two *total* $\beta\eta$ -different nets. In particular, on total nets, such a congruence must be contained in $\simeq_{\beta\eta}$. The Separation Theorem will be fundamental in guiding us towards a definition of observational equivalence (Sect. 2).

1.3. Expressiveness. In strictly computational terms, the interest of the symmetric combinators is given by the following result:

Theorem 1.20 (Lafont [Laf97]). *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 [Laf97, Maz07a]. An example of encoding of linear logic and the λ -calculus in the symmetric combinators¹ is given by Mackie and Pinto [MP02]. We refer the reader to Lafont's paper [Laf97] for a proper formulation and proof of Theorem 1.20.

However, to give an idea of the expressive power of the symmetric combinators, we shall show how general recursion can be implemented in the system, *i.e.*, we shall see how all recursive relations of the form



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




Figure 5: The Lafont code of a net.




Figure 6: The universal decoder for the Lafont code.

may be solved. In the λ -calculus, the above relation would correspond to

$$M \rightarrow_{\beta}^* N[M/x]$$

where x appears free in N . It is well known that a general solution can be given by resorting to a fixpoint combinator, *i.e.*, a term Θ such that, for all T , $\Theta T \rightarrow_{\beta}^* T(\Theta T)$. Then, a solution to the above recursive relation would be $M = \Theta(\lambda x.N)$.

A necessary condition for having a fixpoint combinator is the ability of duplicating any term. In the symmetric combinators, we are only able to duplicate cut-free nets as in Lemma 1.11, so we do not have a fixpoint combinator at our disposal. To compensate for this, we use a fundamental construction due to Lafont [Laf97], and a generalization of it, first considered by Fernandez and Mackie [FM01].

Given a net μ with one free port and containing n cells of type δ , we build a net $!\mu$, called the *Lafont code* of μ , as in Fig. 5. The Z_k are trees of ζ cells, having k leaves. We take Z_0 to be equal to one ε cell; the actual shape of Z_k for $k > 0$ is not important, as long as one tree is fixed for each k . Observe then that, by construction, a Lafont code never contains δ cells. The net μ can be recovered from its Lafont code by means of a “universal decoder”, *i.e.*, independent of μ , as in Fig. 6.

A similar construction removes active pairs and vicious circles, and is given in Fig. 7. This construction uses the Decomposition Lemma 1.6, and the trees Δ_k are defined as the trees Z_k , but using δ cells instead of ζ cells. The net $\xi\mu$ is called the *cut-free code* of μ . Recovering μ from its cut-free code can be done as in Fig. 8. Rigorously speaking, $\xi\mu$ is not




Figure 7: The cut-free code of a net.



Figure 8: The universal decoder for the cut-free code.

well defined because, given a net μ , the cut-free net ν is not unique in general. However, the recovery process works regardless of the particular ν we chose for the cut-free code, so the abuse of notation is not problematic.

We shall take the net $!\mu$ as the *code* of μ . Decoding is done by composing the nets \mathbf{D} and \mathbf{R} of Fig. 6 and 8, respectively; we denote by \mathbf{U} the net resulting from their composition. Observe that the code of a net is cut-free and does not contain δ cells. Hence, Lemma 1.11 applies to it, and it can be freely duplicated by means of trees of δ cells.

We leave it as an instructive exercise for the reader to check that, using Lemmas 1.9, 1.10, and 1.11, the net μ of Fig. 9 is a solution to the recursive relation introduced above (we have supposed that μ has n free ports and that there are k copies of μ on the right hand side of the relation).

2. OBSERVATIONAL EQUIVALENCE

2.1. Observable axioms. We already discussed in the introduction that in the λ -calculus, and in functional programming in general, there is a standard way of defining an observational equivalence, which was first proposed by Morris [Mor68]. The key idea is to put a term (*i.e.*, a program) in a context (*i.e.*, an environment) and to observe its behavior with respect to some interesting property (for example, termination).

In a more abstract way, once we have a language with an internal notion of context, we may take any set S of objects of the language and define S -equivalence as follows: two objects a, b of the language are S -equivalent iff, for every context C , $C[a] \in S$ iff $C[b] \in S$. In other words, we partition the set of all objects into two classes, and we observe the




Figure 9: Solving recursive relations.

contextual behavior of objects with respect to these two classes. Usually, the property defining S (and thus its complementary set) is referred to as an “observable”.

As anticipated above, a typical observable is termination. In fact, Morris himself considered what we shall call *nf-equivalence*, which is obtained by taking S to be the set of normalizable λ -terms. Another fundamental observable is head-termination, *i.e.*, we may define *hnf-equivalence* by taking S to be the set of head-normalizable λ -terms, which is the same as the set of solvable λ -terms. For other examples of observables and for an account of Morris-like observational equivalences in the context of the λ -calculus, we refer the reader to Dezani-Ciancaglini and Giovannetti’s survey [DCG01].

We have shown that in the symmetric combinators, and in interaction nets in general, we have an internal notion of context, and we can therefore hope of applying the above ideas to generate interesting notions of observational equivalence: all we need is finding the right observables. We have already remarked in Sect. 1.2 that termination is not an interesting observable in interaction nets, because of vicious circles. However, the Separation Theorem may give us a hint: in fact, it 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 points out that the appropriate notion of “connection” may be the right thing to observe.

Intuitively, an axiom in a net is observable when it can be “extracted” from the net through interaction, as in the Separation Theorem 1.18. We formalize this intuition as follows.

Definition 2.1 (Observable axiom). Let μ be a net, and let $\omega = \{p, q\}$ be a proper axiom of μ . We say that ω is *observable* iff μ contains two trees τ_i, τ_j , of respective roots i, j , such that p is a leaf of τ_i , q is a leaf of τ_j , and i, j are both free ports of μ . We say that such observable axiom is *based at i, j* .

It is perhaps useful to visualize observable axioms. A net μ contains an observable axiom ω iff it is of the shape given in Fig. 10. If $i = j$, then $\tau_i = \tau_j$, and ω connects two leaves of the same tree. Note also that one or both of τ_i, τ_j may be equal to \bullet ; in particular,




Figure 10: An observable axiom.

a wire whose both extremities are free (as in the Separation Theorem 1.18) is an observable axiom.

Observable axioms may be succinctly described by assigning them an *address*. In the following, we let a, b range over the set $\mathcal{W} = \{\mathbf{p}, \mathbf{q}\}^*$ 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 is denoted by simple juxtaposition, *i.e.*, as ab . The concatenation of two pairs of finite words $a \otimes b, a' \otimes b'$ is defined as $aa' \otimes bb'$, and is also denoted by juxtaposition.

Definition 2.2 (Address). An *address* is an unordered pair of elements of $\mathcal{W} \times \mathcal{W} \times \mathbb{N}$, denoted by $a \otimes b @ i \frown c \otimes d @ j$, and ranged over by $\mathfrak{r}, \mathfrak{q}$. Let τ be a tree, and l a leaf of τ . We associate with l an element of $\mathcal{W} \times \mathcal{W}$, denoted by $\text{br}_\tau(l)$, by induction on τ :

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

Let now $\omega = \{p, q\}$ be an observable axiom of a net μ ; we define its address to be

$$\text{addr}_\mu(\omega) = \text{br}_{\tau_i}(p) @ i \frown \text{br}_{\tau_j}(q) @ j,$$

where τ_i, τ_j are the trees among whose leaves there is p, q , respectively, and i, j are the roots of τ_i, τ_j , respectively, which are free ports of μ (hence $i, j \in \mathbb{N}$ by our convention on free ports of Definition 1.1).

In the following, we denote by $\text{ax}(\mu)$ the set of all addresses of the observable axioms of a net μ , and we define

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

Proposition 2.3. *Let $\mu \rightarrow_\beta^* \mu'$. Then, $\text{ax}(\mu) \subseteq \text{ax}(\mu')$, and $\text{ax}^*(\mu) = \text{ax}^*(\mu')$.*

Proof. Simply look at Fig. 10, which represents the generic form of μ : any active pair must be inside μ_0 , and, after reducing it, since interaction rules are completely local, the “same” axiom as ω is still present in μ' , with the same address. The invariance of $\text{ax}^*(\cdot)$ is a consequence of the confluence of β -reduction. \square




Figure 11: A non-normalizable net producing a finite number of observable axioms.

Note that a net can only have a finite number of observable axioms; then, by Proposition 2.3, $\text{ax}^*(\mu)$ is finite whenever μ is normalizable. However, normalizability is not necessary: non-normalizable nets producing a finite number of observable axioms exist, as shown by the example in Fig. 11: the only observable axiom ever to be found is the one already present in the net starting the reduction sequence, *i.e.*, $\mathbf{1} \otimes \mathbf{pp} @ 1 \frown \mathbf{1} \otimes \mathbf{pq} @ 1$.

We now have two interesting notions of observable: the appearance of observable axioms during reduction, and the fact that these are only produced in finite number.

Definition 2.4 (Observability predicates). We say that μ is *immediately observable*, and we write $\mu \downarrow$, iff $\text{ax}(\mu) \neq \emptyset$. We say that μ is *observable*, and we write $\mu \Downarrow$, iff $\text{ax}^*(\mu) \neq \emptyset$, or, equivalently, $\mu \rightarrow_{\beta}^* \mu' \downarrow$. We say that μ is *finitarily observable*, and we write $\mu \Downarrow$, iff $\text{ax}^*(\mu)$ is non-empty and finite. We write $\mu \uparrow$ and $\mu \Uparrow$ for the negations of $\mu \downarrow$ and $\mu \Downarrow$, respectively. In particular, if $\mu \uparrow$ we say that μ is *blind*.

Definition 2.5 (Observational equivalences). Two nets μ, ν with the same interface are *axiom-equivalent* (resp. *finitarily axiom-equivalent*), and we write $\mu \simeq \nu$ (resp. $\mu \cong \nu$), iff for all contexts C , $C[\mu] \downarrow$ iff $C[\nu] \downarrow$ (resp. $C[\mu] \Downarrow$ iff $C[\nu] \Downarrow$).

It helps thinking of an immediately observable net as a head normal form in the λ -calculus. This analogy can be made more precise: our definition of observable net can in fact be extended to any interaction net system [Maz06], in particular to sharing graphs [?, GAL92]; then, one can see that the head variable of a head normal form T corresponds to an observable axiom in the sharing graph of T .

Observe that, once we think of observable axioms as head variables, we are naturally led to see $\text{ax}^*(\mu)$ as nothing but a sort of “Böhm tree” of μ . For instance, $\text{ax}^*(\cdot)$ is an invariant of reduction (Proposition 2.3), just like the Böhm tree of a λ -term. In Sect. 3.1, we shall develop an abstract interpretation of nets starting from this intuition.

2.2. Observable axioms and the geometry of interaction. The geometry of interaction (GoI) was introduced by Girard [Gir89] as a mathematical formulation, using functional analysis and operator algebras, of the cut-elimination process in linear logic. Later, it was reformulated using much less sophisticated tools by Danos and Regnier [DR95], and it was also transported to the interaction combinators (symmetric and not) by Lafont [Laf97].

$$\begin{array}{lll}
\mathbf{c}^* \mathbf{c} \succ \mathbf{1} & \mathbf{c}^* \mathbf{f} \succ \mathbf{f} \mathbf{c}^* & \mathbf{f}^* \mathbf{c} \succ \mathbf{c} \mathbf{f}^* \\
\mathbf{d}^* \mathbf{d} \succ \mathbf{1} & \mathbf{c}^* \mathbf{g} \succ \mathbf{g} \mathbf{c}^* & \mathbf{f}^* \mathbf{d} \succ \mathbf{d} \mathbf{f}^* \\
\mathbf{f}^* \mathbf{f} \succ \mathbf{1} & \mathbf{d}^* \mathbf{f} \succ \mathbf{f} \mathbf{d}^* & \mathbf{g}^* \mathbf{c} \succ \mathbf{c} \mathbf{g}^* \\
\mathbf{g}^* \mathbf{g} \succ \mathbf{1} & \mathbf{d}^* \mathbf{g} \succ \mathbf{g} \mathbf{d}^* & \mathbf{g}^* \mathbf{d} \succ \mathbf{d} \mathbf{g}^*
\end{array}$$

Figure 12: Word rewriting on monomials.

As a rough approximation, we can say that the GoI interprets nets as collections of paths which are stable under β -reduction; we shall see that these paths ultimately correspond to observable axioms.

Definition 2.6 (Monomial, rewriting and value of monomials). An *atom* is an element of $\{\mathbf{c}, \mathbf{d}, \mathbf{f}, \mathbf{g}, \mathbf{c}^*, \mathbf{d}^*, \mathbf{f}^*, \mathbf{g}^*\}$; and atom is *positive* if it belongs to $\{\mathbf{c}, \mathbf{d}, \mathbf{f}, \mathbf{g}\}$, otherwise it is *negative*. A *monomial* is a finite word on the set of atoms. The set of monomials is denoted by \mathcal{M} , and ranged over by A, B ; the empty monomial will be denoted by $\mathbf{1}$. A monomial is *positive* if it is empty or contains only positive atoms. We define an involution $(\cdot)^*$ on \mathcal{M} by setting $(\mathbf{c})^* = \mathbf{c}^*$, $(\mathbf{d})^* = \mathbf{d}^*$, and similarly for \mathbf{f}, \mathbf{g} ; $(\mathbf{c}^*)^* = \mathbf{c}$, $(\mathbf{d}^*)^* = \mathbf{d}$, and similarly for $\mathbf{f}^*, \mathbf{g}^*$; $\mathbf{1}^* = \mathbf{1}$, and $(AB)^* = B^*A^*$.

We define a word rewriting relation \succ on \mathcal{M} as in Fig. 12. We say that $A' \in \mathcal{M}$ is *clash-free* iff $A' \succ^* AB^*$, where A, B are positive monomials.

Let $A \in \mathcal{M}$, and let A_0^* be an occurrence of negative atom in A , *i.e.*, $A = A'A_0^*A''$ for some $A', A'' \in \mathcal{M}$. The *value* of A_0^* in A , denoted by $v_A(A_0^*)$, is the number of positive atoms in A'' . We define the *value* of A to be $v(A) = \sum_{A_0^*} v_A(A_0^*)$, where A_0^* ranges over the occurrences of negative atoms in A .

Lemma 2.7. *Let $A \succ B$. Then, $v(A) > v(B)$.*

Proof. By simple inspection of Fig. 12. □

Proposition 2.8. *Rewriting of monomials is confluent and strongly normalizing.*

Proof. For what concerns confluence, simply observe that there are no critical pairs. Strong normalization is a consequence of Lemma 2.7. □

Definition 2.9 (Port graph). The *port graph* of a net μ , denoted by $\text{PG}(\mu)$, is an undirected multigraph whose edges are weighed in the positive monomials, defined as follows: its vertices are the elements of $\text{Ports}(\mu)$, and there is an edge between two ports p, q iff one of the following (non mutually exclusive) conditions hold:

external edge: $\{p, q\} \in \text{Wires}(\mu)$; the weight in this case is $\mathbf{1}$;

internal edge: p and q are principal and auxiliary ports of a cell of μ ; the weight in this case depends on the auxiliary port q : it is \mathbf{c} (resp. \mathbf{d}) if q is port number 1 (resp. 2) of a δ cell, and it is \mathbf{f} (resp. \mathbf{g}) if q is port number 1 (resp. 2) of a ζ cell.

As an example, in Fig. 13 we give the port graph of the net of Fig. 1.

Definition 2.10 (Paths, weights [DR95, Laf97], values). A *straight path* of a net μ is a path of $\text{PG}(\mu)$ which does not contain two consecutive internal edges. We say that a straight path *crosses an active pair* iff it contains an external edge corresponding to a proper cut. A *maximal path* is a non-empty straight path connecting two free ports of μ . An *observable path* is a maximal path crossing no active pair. We denote by $\text{mp}_{i,j}(\mu)$ the set of maximal




Figure 13: The port graph of the net of Fig. 1. Internal edges are dotted, and their weight annotated beside them; the other weights are omitted, and are all equal to $\mathbf{1}$.

paths of a net μ from its free port i to its free port j , and we denote by $\text{mp}(\mu)$ the set of all maximal paths of μ .

Let $\phi \in \text{mp}(\mu)$, and let e be an internal edge of $\text{PG}(\mu)$ used by ϕ . Note that e must correspond to a binary cell c of μ ; hence, there is no ambiguity in saying that ϕ “enters” c , uses the internal edge, and then “exits” c . We say that e is *crossed downwards* (resp. *upwards*) by ϕ if ϕ enters c through an auxiliary port (resp. through its principal port) and exists c through its principal port (resp. through one of its auxiliary ports). Let A be the weight of e in $\text{PG}(\mu)$. The *weight* of e in ϕ , denoted by $w_\phi(e)$, is A if e is crossed downwards; otherwise, it is A^* . If e is an external edge, we fix $w_\phi(e) = \mathbf{1}$.

Now let $\phi = e_1, \dots, e_n$ be a maximal path of a net; we define the *weight* of ϕ to be the following monomial:

$$w(\phi) = w_\phi(e_n) \cdots w_\phi(e_1)$$

(note the reversal of the order of edges). Moreover, we define the *value* of ϕ , abusively denoted by $v(\phi)$, as $v(w(\phi))$.

In the sequel, we shall speak of paths in a net μ without explicitly referring to $\text{PG}(\mu)$. This will not be a source of confusion, because all edges of $\text{PG}(\mu)$ correspond to either wires of μ or “internal connections” represented by the cells of μ ; hence, given a straight path of μ , we can easily trace it directly on its graphical representation, and vice versa.

Note that the presence of an observable path in a net μ implies the presence of exactly one observable axiom in μ , and vice versa (modulo orientation of paths, *i.e.*, an observable axiom actually induces two observable paths, which are the reversal of each other). In fact, if we look at Fig. 10, we find an observable path by entering the net through its free port i , “going up” τ_i through the branch leading to ω , following ω itself, and “descending” τ_j through the only branch leading us to its root, which is the free port j . Conversely, it is easy to see that any observable path must be of this form (with τ_i, τ_j which may be equal to \bullet), because of the absence of active pairs.

Definition 2.11 (Residue and lift of a maximal path). Let $\mu \rightarrow_\beta \mu'$, and let $\phi \in \text{mp}(\mu)$. The *residue* ϕ' of ϕ in μ' is, if it exists, the maximal path of μ' defined as follows. If ϕ does not cross the active pair reduced, then by the locality of interaction rules, “the same” path as ϕ is found in μ' , and this is taken to be ϕ' . Otherwise, we call 1, 2 and 3, 4 the auxiliary ports of the cells (which must be binary, because ϕ is maximal) composing the active pair reduced, and we distinguish two cases:

- the two cells have the same symbol:

- if ϕ connects 1 to 3 or 2 to 4, then this connection becomes a wire in μ' , so ϕ' is equal to what is left of ϕ with the active pair replaced by a wire;
- if ϕ connects 1 to 4 or 2 to 3, then ϕ has no residue;
- the two cells have different symbols; then, whatever ports are connected by ϕ , the connection is still present in μ' , so ϕ' is equal to what is left of ϕ with the active pair replaced by this new connection.

Of course ϕ may cross the active pair more than once, but this is not problematic: its residue is still defined as above, replacing every crossing with the appropriate path described above. Conversely, if $\phi' \in \text{mp}(\mu')$, then it is the residue of exactly one maximal path of μ , which is called the *lift* of ϕ' in μ .

Remark that the notions of residue and lift can be extended to reductions of arbitrary length: if $\mu \rightarrow_{\beta}^* \mu'$, and if ϕ is a maximal path of μ , we can look for its residue (if it exists) by tracing the successive residues of ϕ along the reduction; conversely, if ϕ' is a maximal path of μ' , by successively lifting ϕ' along the reduction, we obtain the lift of ϕ' in μ . Remark also that residues connect the same free ports as their lifts: if $\phi \in \text{mp}_{i,j}(\mu)$, and if ϕ' is the residue of ϕ in a reduct μ' of μ , then $\phi' \in \text{mp}_{i,j}(\mu')$.

The following shows that monomial rewriting is related to β -reduction:

Proposition 2.12. *Let $\mu \rightarrow_{\beta} \mu'$, let $\phi' \in \text{mp}(\mu')$, and let ϕ be the lift of ϕ' in μ . Then, $w(\phi) \succ^* w(\phi')$.*

Proof. If ϕ does not cross the active pair reduced, we have $\phi' = \phi$, and the result trivially holds. So suppose that ϕ crosses the active pair reduced. We start by observing that, since ϕ is maximal, the active pair reduced must concern two binary cells c, c' . We set i_1, i_2 to be the auxiliary port number 1 and number 2 of c , respectively, and i'_1, i'_2 to be the auxiliary port number 1 and number 2 of c' , respectively. We have two cases: an annihilation, or a commutation. Suppose we are in the first case. Since ϕ has a reduct in μ' , each time ϕ crosses the active pair made by c, c' , it must do so by using the pairs of ports i_1, i'_1 or i_2, i'_2 . Hence, if ϕ crosses the active pair n times, there exist n occurrences of positive atoms A_1, \dots, A_n such that $w(\phi) = \dots A_1^* A_1 \dots A_n^* A_n \dots$, *i.e.*, the weight of ϕ contains a word $A_k^* A_k$ for each crossing of the active pair. Now, by Definition 2.11, all such crossings are replaced by wires in ϕ' , so $w(\phi') = \dots \mathbf{1} \dots \mathbf{1} \dots$, and we thus clearly have $w(\phi) \succ^* w(\phi')$. Suppose now that we are in the case of a commutation, with ϕ crossing n times the active pair made of c, c' . This time, we must have n negative atoms A_1^*, \dots, A_n^* and n positive atoms B_1, \dots, B_n such that $w(\phi) = \dots A_1^* B_1 \dots A_n^* B_n \dots$, and, for each $1 \leq k \leq n$, the atoms satisfy that, if $A_k^* \in \{\mathbf{c}^*, \mathbf{d}^*\}$, then $B_k \in \{\mathbf{f}, \mathbf{g}\}$, and if $A_k^* \in \{\mathbf{f}^*, \mathbf{g}^*\}$, then $B_k \in \{\mathbf{c}, \mathbf{d}\}$. Hence, again by Definition 2.11, we have $w(\phi) = \dots B_1 A_1^* \dots B_n A_n^* \dots$, and, by looking at Fig. 12, $w(\phi) \succ^* w(\phi')$. \square

Proposition 2.12 is the basis of the GoI. In fact, the above result basically transforms β -reduction into a word rewriting system; the idea then is to take a model of this rewriting system and build from it a model of β -reduction. By “model”, we mean a function $\llbracket \cdot \rrbracket$ interpreting monomials in some mathematical structure such that, for all $A, B \in \mathcal{M}$, $A \succ B$ implies $\llbracket A \rrbracket = \llbracket B \rrbracket$. The following construction, given by Lafont [Laf97] and inspired by previous work of Girard [Gir89], does precisely this.

In the following, an *involutive monoid* is a couple $(M, (\cdot)^*)$ where M is a multiplicative monoid, and $(\cdot)^*$ an involutive antiautomorphism of M . A homomorphism f between two involutive monoids $(M, (\cdot)^*), (M', (\cdot)^\dagger)$ is a homomorphism of monoids preserving the

involution, *i.e.*, such that, for all $a \in M$, $f(a^*) = f(a)^\dagger$. Similarly, an *involutive unit semiring* is a couple $(S, (\cdot)^*)$ where S is a unit semiring, whose additive and multiplicative units are denoted by $\mathbf{0}$ and $\mathbf{1}$, respectively, and $(\cdot)^*$ is an involutive antiautomorphism of S .

Let \mathcal{R} be the involutive unit semiring generated by the set $\{\mathbf{p}, \mathbf{q}\}$ and by the relations

$$\begin{aligned} \mathbf{p}^* \mathbf{p} &= \mathbf{1} & \mathbf{p}^* \mathbf{q} &= \mathbf{0}, \\ \mathbf{q}^* \mathbf{q} &= \mathbf{1} & \mathbf{q}^* \mathbf{p} &= \mathbf{0}. \end{aligned}$$

Consider now the semiring $\mathcal{R} \otimes \mathcal{R}$; this is still an involutive unit semiring, the multiplicative unit being $\mathbf{1} \otimes \mathbf{1}$ and the involution being defined by $(\sum_{i=1}^n a_i \otimes b_i)^* = \sum_{i=1}^n a_i^* \otimes b_i^*$. Our model will interpret monomials in $\mathcal{R} \otimes \mathcal{R}$, as follows. By definition, the set \mathcal{M} is a free involutive monoid, and $\mathcal{R} \otimes \mathcal{R}$ may also be seen as an involutive monoid, by taking its multiplicative part; hence, we can define a homomorphism $\llbracket \cdot \rrbracket$ by setting

$$\begin{aligned} \llbracket \mathbf{c} \rrbracket &= \mathbf{p} \otimes \mathbf{1} & \llbracket \mathbf{f} \rrbracket &= \mathbf{1} \otimes \mathbf{p}, \\ \llbracket \mathbf{d} \rrbracket &= \mathbf{q} \otimes \mathbf{1} & \llbracket \mathbf{g} \rrbracket &= \mathbf{1} \otimes \mathbf{q}. \end{aligned}$$

The fact that \mathcal{M} is free ensures that $\llbracket \cdot \rrbracket$ is defined everywhere once it is defined on its generators.

Proposition 2.13. *For all $A, B \in \mathcal{M}$, $A \succ B$ implies $\llbracket A \rrbracket = \llbracket B \rrbracket$.*

Proof. The rules of Fig. 12 of the form $A^* A \succ \mathbf{1}$ are modelled by the annihilation relations $a^* a = \mathbf{1}$ defining \mathcal{R} ; the rules of the form $A^* B \succ B A^*$ are modelled by the commutations $(a \otimes \mathbf{1})(\mathbf{1} \otimes b) = (a \otimes b) = (\mathbf{1} \otimes b)(a \otimes \mathbf{1})$ in $\mathcal{R} \otimes \mathcal{R}$. \square

Now, given a cut-free net ν with n free ports, the GoI assigns to it a formal $n \times n$ matrix $\llbracket \nu \rrbracket$ with coefficients belonging to $\mathcal{R} \otimes \mathcal{R}$ and defined as follows:

$$\llbracket \nu \rrbracket_{j,i} = \sum_{\phi \in \text{mp}_{i,j}(\nu)} \llbracket w(\phi) \rrbracket$$

Note that $\text{mp}_{i,j}(\mu)$ coincides with the set of observable paths from i to j , because ν is cut-free; there are obviously finitely many of these, so the above sum is finite and defines an element of $\mathcal{R} \otimes \mathcal{R}$.

If μ is a net, we know by the Decomposition Lemma 1.6 that we can always write $\mu = \sigma[\nu_0]$ with ν_0 cut-free and σ a feedback context; the key result of the GoI is that, if μ is total of cut-free form ν , we can compute $\llbracket \nu \rrbracket$ starting from $\llbracket \nu_0 \rrbracket$. For this, we use a formal matrix associated with σ , which we denote by $\llbracket \sigma \rrbracket$, defined as follows: if ν_0 has $n + 2k$ free ports, then $\llbracket \sigma \rrbracket$ is a $(n + 2k) \times (n + 2k)$ matrix, such that $\llbracket \sigma \rrbracket_{i,j} = \mathbf{1}$ if σ connects the free ports i and j of ν_0 , and $\llbracket \sigma \rrbracket_{i,j} = \mathbf{0}$ otherwise.

Theorem 2.14. *Let $\mu = \sigma[\nu_0]$ be a net with n free ports, with ν_0 having $n + 2k$ free ports. Then, μ is total iff $\llbracket \sigma \rrbracket \llbracket \nu_0 \rrbracket$ is nilpotent, and in that case, if ν is the cut-free form of μ , we have*

$$\llbracket \nu \rrbracket = \pi^t \left(\sum_{h=0}^{\infty} \llbracket \nu_0 \rrbracket (\llbracket \sigma \rrbracket \llbracket \nu_0 \rrbracket)^h \right) \pi,$$

where π is the formal matrix of the inclusion morphism of $(\mathcal{R} \otimes \mathcal{R})^n$ into $(\mathcal{R} \otimes \mathcal{R})^{n+2k}$, and π^t its transpose.

Proof. This was originally proved by Girard for linear logic² [Gir89], with nilpotency only as a necessary condition for strong normalization (*i.e.*, totality); Danos and Regnier later proved the converse [DR95]. For the symmetric combinators, the result with nilpotency as a necessary condition is due to Lafont [Laf97]; a proof that nilpotency is also sufficient for totality can be found in the author’s Ph.D. thesis [Maz06] (as the proof of Theorem 3.70). \square

The formula for computing $\llbracket \nu \rrbracket$ from $\llbracket \nu_0 \rrbracket$ and $\llbracket \sigma \rrbracket$ given in Theorem 2.14 is known as the *execution formula*. For total nets, the execution formula is an invariant of β -reduction, and therefore gives a model of \simeq_β . However, this does not work for non-total nets, because in that case the sum in the execution formula has an infinite number of terms. To handle it, one solution would be to put some topology on $\mathcal{R} \otimes \mathcal{R}$ (or rather, on its algebra of formal matrices), and study the convergence of the execution formula in this topology. Another solution, which is the one we chose in this paper, is to deal directly with possibly infinite objects, just like one deals with possibly infinite Böhm trees in the λ -calculus.

In fact, even when it “diverges”, the execution formula is not completely meaningless: it computes the interpretations of the weights of those maximal paths which are never “destroyed” by reduction. We shall see that these correspond to the observable axioms generated during reduction. This gives a further justification for our definitions of Sect. 2.1.

Definition 2.15 (Execution path [DR95]). Let μ be a net. An *execution path* of μ is a maximal path ϕ of μ such that, whenever $\mu \rightarrow_\beta^* \mu'$, ϕ has a residue in μ' .

Remark that, as an immediate consequence of the definition, any residue of an execution path is itself an execution path. As we said above, execution paths are those which are preserved by reduction, and which eventually generate observable paths (and hence observable axioms).

Lemma 2.16. *Let μ be a net, and let ϕ be an execution path of μ . Then:*

- (1) $v(\phi) = 0$ implies that ϕ is observable;
- (2) $v(\phi) > 0$ implies that ϕ crosses an active pair, reducing which we obtain $\mu \rightarrow_\beta \mu'$, and the residue ϕ' of ϕ in μ' is such that $v(\phi') < v(\phi)$.

Proof. For part (1), observe that $v(\phi) = 0$ implies $w(\phi) = AB^*$ for some positive monomials A, B , which is clearly the weight of an observable path. For part (2), note that $v(\phi) > 0$ implies $w(\phi) = A'A_0^*B_0B'$ for some $A', B' \in \mathcal{M}$ and some positive atoms A_0, B_0 . Then, ϕ crosses an active pair, so we apply Proposition 2.12 (or rather its proof) and Lemma 2.7. \square

Lemma 2.17. *Let μ be a net, and let $\phi \in \text{mp}(\mu)$. Then, the following are equivalent:*

- (1) ϕ is an execution path;
- (2) $\mu \rightarrow_\beta^* \mu'$ such that the residue of ϕ in μ' is an observable path;
- (3) $w(\phi)$ is clash-free;
- (4) $\llbracket w(\phi) \rrbracket \neq \mathbf{0}$.

Proof. (1) implies (2) is proved by induction on $v(\phi)$, using Lemma 2.16.

For (2) implies (3), let ϕ' be the residue of ϕ in μ' . Since it is an observable path, we have $w(\phi') = AB^*$ for some positive monomials A, B , and we conclude by Proposition 2.12.

²Actually, as formulated here, this result holds only for multiplicative linear logic; some technical constraints are needed in full linear logic, because the execution formula does not model cut-elimination in the general case.

For (3) implies (4), observe that normal form of $w(\phi)$ is by definition of the form AB^* , so by Proposition 2.13 $\llbracket w(\phi) \rrbracket = \llbracket AB^* \rrbracket = \llbracket A \rrbracket \llbracket B \rrbracket^*$, which is never equal to $\mathbf{0}$ by definition of $\llbracket \cdot \rrbracket$.

For (4) implies (1), we prove the contrapositive. Suppose ϕ is a maximal path of μ which is not an execution path. This means that $\mu \rightarrow_{\beta}^* \mu' \rightarrow_{\beta} \mu''$, and the residue ϕ' of ϕ in μ' has no residue in μ'' . By Definition 2.11, the only possibility is that ϕ' crosses an active pair made of two binary cells c, c' of the same kind, using auxiliary port 1 of c and auxiliary port 2 of c' . Hence, we have, for some $A', B' \in \mathcal{M}$, $w(\phi') = A'A_0^*B_0B'$, and either $A_0, B_0 \in \{\mathbf{c}, \mathbf{d}\}$, or $A_0, B_0 \in \{\mathbf{f}, \mathbf{g}\}$, but in both cases $A_0 \neq B_0$. Hence, by Proposition 2.13 and by definition of \mathcal{R} , we have $\llbracket w(\phi) \rrbracket = \llbracket w(\phi') \rrbracket = \llbracket A' \rrbracket \llbracket A_0 \rrbracket^* \llbracket B_0 \rrbracket \llbracket B' \rrbracket = \llbracket A' \rrbracket \mathbf{0} \llbracket B' \rrbracket = \mathbf{0}$. \square

Now, we invite the reader to check that, if a net ν has an observable axiom ω whose address is $s @ i \frown t @ j$, and if ϕ is the observable path from i to j induced by ω , we have $\llbracket w(\phi) \rrbracket = ts^*$. Vice versa, if ν has an observable path ϕ from free port i to free port j such that $w(\phi) = BA^*$, then the corresponding observable axiom will have address $\llbracket A \rrbracket @ i \frown \llbracket B \rrbracket @ j$. This is the reason behind our choice of notations for the addresses of observable axioms; moreover, it allows us to state the following:

Proposition 2.18. *For every net μ , we have*

$$\text{ax}^*(\mu) = \{s @ i \frown t @ j \mid \phi \in \text{mp}_{i,j}(\mu), \llbracket w(\phi) \rrbracket = ts^*\}.$$

Proof. For what concerns the inclusion from left to right, we have $\mu \rightarrow_{\beta}^* \mu'$ such that μ' contains an observable axiom of address $s @ i \frown t @ j$. By the remark made above, we know that this observable axiom induces an observable path ϕ' of μ' from free port i to free port j such that $\llbracket w(\phi') \rrbracket = ts^*$; then, we take its lift ϕ in μ , and conclude by Propositions 2.12 and 2.13.

For the other inclusion, by hypothesis $\llbracket w(\phi) \rrbracket \neq \mathbf{0}$, so by Lemma 2.17 we have $\mu \rightarrow_{\beta}^* \mu'$ such that the residue ϕ' of ϕ in μ is observable. By Propositions 2.12 and 2.13, $\llbracket w(\phi') \rrbracket = \llbracket w(\phi) \rrbracket$, and, as remarked above, we know that ϕ' induces an observable axiom of address $s @ i \frown t @ j$, which is in $\text{ax}^*(\mu)$ by definition. \square

We encourage the reader to compare Proposition 2.18 with the definition of $\llbracket \nu \rrbracket$ given above: basically, we can see $\text{ax}^*(\cdot)$ as the extension of $\llbracket \cdot \rrbracket$ to arbitrary nets. In fact, if we tried to define $\llbracket \mu \rrbracket$ in the general case, the sum ranging over $\text{mp}_{i,j}(\mu)$ might be infinite; as mentioned above, instead of introducing a topology to handle series, we opted for a ‘‘Böhm-tree’’ approach: we reduce μ , and collect the observable axioms showing up along the way, accepting that there may be an infinity of them. Using formal matrices does not make much sense at this point, but we still need to retain the information concerning the free ports connected by the observable paths: this is the reason behind the presence of integers in addresses. The fact that addresses are unordered pairs reflects the fact that an observable path carries the same information as its reversal; indeed, in the GoI interpretation of a cut-free net ν one can show that $\llbracket \nu \rrbracket_{i,j} = \llbracket \nu \rrbracket_{j,i}^*$ (all operators are ‘‘Hermitian’’ [Gir89]).

The question of what can be done by taking the topological approach instead of the ‘‘Böhm tree’’ approach is left open, and is out of the scope of this work. We shall see that topology will eventually play a fundamental role in our approach too, but for quite different purposes.



Figure 14: A quasi-wire; μ is an arbitrary net with an empty interface.

2.3. Solvability and ε -reduction. The analogy “an immediately observable net is like a λ -term in head normal form” can be given a justification within the theory of the symmetric combinators. In the following, we call a net of the form given in Fig. 14 a *quasi-wire*.


Definition 2.19 (Solvable net). A net μ is *solvable* iff there exist a test θ such that $\theta[\mu] \rightarrow_{\beta}^* W$, where W is a quasi-wire.

We shall see that, just as λ -terms having a head normal form coincide with solvable λ -terms, so in the symmetric combinators observable nets and solvable nets coincide. Furthermore, in Sect. 5.2 we shall give evidence supporting the fact that the above notion of solvable net is indeed analogous to that of a solvable λ -term.

Let C be a net, and let I be a subset of its interface. We say that C is *relatively blind* on I iff, whenever $s@i \frown t@j \in \text{ax}^*(C)$, $i \in I$ implies $j \notin I$, i.e., C generates no observable axiom all based within I . A context C for nets with n free ports will be said to be relatively blind if its interface is $\{1, \dots, n\} \uplus I$ and C is relatively blind on I .

Lemma 2.20. *For every net μ and relatively blind context C , $\mu \uparrow$ implies $C[\mu] \uparrow$.*

Proof. Let $\mu \uparrow$, and suppose for the sake of absurdity that $C[\mu] \rightarrow_{\beta}^* \mu' \downarrow$. The observable axiom in μ' induces an observable path ϕ' in μ' , which has a lift ϕ in $C[\mu]$, connecting two free ports of the relatively blind interface of C . Suppose that $\phi \in \text{mp}(C)$. By looking at the proof of Proposition 2.12, we see that ϕ can be transformed into an observable path by reducing only the active pairs it crosses; hence, we would have $C \rightarrow_{\beta}^* C'$, with ϕ' an observable path of C' , which would contradict the fact that C is relatively blind. Then, we must have the following situation



where we drew ϕ as a dashed line. As a consequence, we can write $\phi = \phi' \phi_0 \phi''$, where $\phi_0 \in \text{mp}(\mu)$. We thus have $w(\phi) = w(\phi'')w(\phi_0)w(\phi')$, which means that $w(\phi_0)$ is a subword of $w(\phi)$. But, by Lemma 2.17, $w(\phi)$ is clash-free, so $w(\phi_0)$ is clash-free too; if we apply Lemma 2.17 again, we obtain that ϕ_0 has an observable residue, which, by the above remark, can be obtained by reducing only the active pairs it crosses, which are all within μ . Therefore, we obtain $\mu \downarrow$, a contradiction. \square


Lemma 2.21. *Let μ be a net. Then, $\mu \uparrow$ implies $\theta[\mu] \uparrow$ for any test θ .*

Proof. Simply observe that a test is a relatively blind context. \square

Now, the forward implication of the equivalence “observable iff solvable” holds because observable axioms are defined precisely with the intuition that we can “extract” a wire from μ by means of a test; the converse is a consequence of Lemma 2.21:

Proposition 2.22. *A net μ is observable iff it is solvable.*

Proof. Let $\mu \Downarrow$; by Definition 2.4, we know that μ reduces to a net of the shape given in Fig. 10. Then, consider the test



where we leave free exactly the two leaves of τ_i, τ_j corresponding to those connected by the observable axiom ω of Fig. 10. By Lemma 1.9, we have that $\theta[\mu]$ β -reduces to a quasi-wire, as desired. Suppose now $\mu \Uparrow$. Remark that quasi-wires are immediately observable, so if μ were solvable, we would have a test θ such that $\theta[\mu] \Downarrow$, contradicting Lemma 2.21. \square

Another interesting consequence of Lemma 2.21 is that finitary axiom-equivalence is stronger than axiom-equivalence:

Proposition 2.23. *For all μ, ν , $\mu \cong \nu$ implies $\mu \simeq \nu$.*

Proof. Assume $\mu \cong \nu$, and let C be a context such that $C[\mu] \Downarrow$. We need to show that $C[\nu] \Downarrow$; by symmetry of \cong , this will be enough to prove the result. By Proposition 2.22, $C[\mu]$ is solvable, so we have a test θ such that $\theta[C[\mu]]$ β -reduces to a quasi-wire. But quasi-wires generate exactly one observable axiom, so $\theta[C[\mu]] \Downarrow$, which by $\mu \cong \nu$ implies $\theta[C[\nu]] \Downarrow$, which by Definition 2.4 implies $\theta[C[\nu]] \Downarrow$, which implies $C[\nu] \Downarrow$ thanks to Lemma 2.21. \square

The λ -calculus has an interesting notion of Ω -reduction, which is defined by the reduction rule $M \rightarrow_{\Omega} \Omega$ iff M is unsolvable and different from Ω , where Ω is itself some fixed unsolvable term (usually one takes $\Omega = (\lambda x.xx)(\lambda x.xx)$). Added to other reductions, Ω -reduction has interesting properties: $\beta\Omega$ -reduction and $\beta\eta\Omega$ -reduction characterize provability in the theory \mathcal{H} (the smallest sensible theory, cf. Sect. 5.3) and $\mathcal{H}\eta$, respectively [Bar84]. This latter coincides with nf-equivalence, defined in Sect. 2.1.

Since we have our own notion of unsolvable net, it may be interesting to study the behavior of the following rewriting rule, directly inspired by Ω -reduction:

$$\boxed{\begin{array}{c} \mu \Uparrow \\ \vdots \end{array}} \rightarrow_{\varepsilon} \begin{array}{c} \epsilon \\ \vdots \end{array} \dots \begin{array}{c} \epsilon \\ \vdots \end{array} = \mathbf{E}_n$$

where μ has n free ports and $\mu \neq \mathbf{E}_n$.

Definition 2.24 (ε - and $\beta\varepsilon$ -reduction). We write $\rightarrow_{\varepsilon}$ for the contextual closure of the above rule, and we define $\rightarrow_{\beta\varepsilon} = (\rightarrow_{\beta} \cup \rightarrow_{\varepsilon})$.

Of course the rewriting rule defining ε -reduction is not recursive, because it is undecidable whether a net is blind (intuitively, this is a consequence of the Turing-completeness of the symmetric combinators, and of Rice's theorem—any non-trivial class of recursive functions, hence of nets, is undecidable). This is exactly the same situation as Ω -reduction, where it is undecidable whether a λ -term is unsolvable. We shall see that the interest of $\beta\varepsilon$ -reduction is in its relationship with finitary axiom-equivalence (Corollary 2.31).

We say that a binary relation on nets \rightsquigarrow has the *quasi-diamond property* iff $\mu \rightsquigarrow \mu_1$ and $\mu \rightsquigarrow \mu_2$ implies that there exists ν such that $\mu_1 \rightsquigarrow \nu$ or $\mu_1 = \nu$, and $\mu_2 \rightsquigarrow \nu$ or $\mu_2 = \nu$.

Lemma 2.25. *Let \rightsquigarrow be a binary relation on nets satisfying the quasi-diamond property. Then, its reflexive transitive closure \rightsquigarrow^* satisfies the diamond property, i.e., it is confluent.*


Proof. A standard diagram-chasing argument. \square

Lemma 2.26. *The relation $\rightarrow_{\beta\varepsilon}$ satisfies the quasi-diamond property.*


Proof. Let $\mu \rightarrow_{\beta\varepsilon} \mu_1$ and $\mu \rightarrow_{\beta\varepsilon} \mu_2$. We may suppose $\mu_1 \neq \mu_2$, otherwise there is nothing to prove. If the two reductions come from two active pairs, we conclude by applying Proposition 1.8. Otherwise, suppose without loss of generality that $\mu \rightarrow_{\beta} \mu_1$ and $\mu \rightarrow_{\varepsilon} \mu_2$, *i.e.*, we have $\mu = C[\nu]$ and $\mu_2 = C[\mathbf{E}_n]$, where ν is a blind net with n free ports. We have three cases:

- $\nu \rightarrow_{\beta} \nu'$ and $\mu_1 = C[\nu']$, *i.e.*, the active pair reduced to obtain μ_1 is contained in ν . In that case, we still have $\nu' \uparrow$, so $\mu_1 \rightarrow_{\varepsilon} \mu_2$.
- The active pair reduced to obtain μ_1 is “between” ν and C , *i.e.*, one of its cells is in ν and the other, call it c , is in C . We suppose c to be binary; the nullary case is easier, and left to the reader. Then, the cell c together with a suitable identity wiring (which may be empty in case $n = 1$) forms a test θ , and we can write $\mu = C'[\theta[\nu]]$ for a suitable context C' . Now $\theta[\nu] \rightarrow_{\beta} \nu'$ and $\mu_1 = C'[\nu']$, while $\mu_2 = C'[\theta[\mathbf{E}_n]]$. By Lemma 2.21, we have both $\theta[\nu] \uparrow$ and $\theta[\mathbf{E}_n] \uparrow$, so both μ_1 and μ_2 ε -reduce in one step to $C'[\mathbf{E}_{n+1}]$.
- The active pair reduced is completely disjoint from ν . This case is trivial.

We are left with the situation in which both μ_1 and μ_2 are obtained by means of ε -steps. Let ν_1, ν_2 be the blind subnets of μ reduced to obtain μ_1 and μ_2 , respectively. If ν_1 and ν_2 are disjoint, then the diamond property holds trivially. Otherwise, we have $\mu = C[\nu]$, where ν is the net



and ν_1 is equal to ν'_1 plus ν_0 , while ν_2 is equal to ν'_2 plus ν_0 . Now, if we put



we have $\mu_1 = C[o_1]$ and $\mu_2 = C[o_2]$. But ν'_1 and ν'_2 must be relatively blind on I_1 and I_2 , respectively, because ν_1 and ν_2 are blind. Hence, by Lemma 2.20, the subnets marked by the dashed rectangles in the above picture are both blind, so μ_1 and μ_2 both reduce in at most one ε -step to $C[\mathbf{E}_n]$, where n is the number of free ports of ν . \square

Proposition 2.27 (Confluence of $\beta\varepsilon$ -reduction). *The relation $\rightarrow_{\beta\varepsilon}^*$ is confluent.*

Proof. Apply Lemmas 2.25 and 2.26. \square

The confluence of $\beta\varepsilon$ -reduction allows us to introduce the following congruences:

Definition 2.28 ($\beta\varepsilon$ - and $\beta\varepsilon\eta$ -equivalence). $\beta\varepsilon$ -equivalence is defined by $\mu \simeq_{\beta\varepsilon} \nu$ iff there exists o such that $\mu \rightarrow_{\beta\varepsilon}^* o$ and $\nu \rightarrow_{\beta\varepsilon}^* o$; $\beta\varepsilon\eta$ -equivalence is defined by $\simeq_{\beta\varepsilon\eta} = (\simeq_{\beta\varepsilon} \cup \simeq_{\eta})^+$.

Note that $\beta\varepsilon$ -normal forms are always cut-free. In particular, we have the following characterization, whose proof is left to the reader. In the following, an ε -tree is a tree with no leaves; the ε -tree ε is called *trivial*.

Proposition 2.29 ($\beta\varepsilon$ -normal forms). *A net μ is $\beta\varepsilon$ -normal iff it is cut-free, and each ε -tree contained in μ is trivial.* \square

The following result shows that $\beta\varepsilon$ -reduction is related to finitary axiom-equivalence. We first need to extend Proposition 2.3 to $\beta\varepsilon$ -reduction, which is unproblematic:

Proposition 2.30. $\mu \rightarrow_{\beta\varepsilon}^* \nu$ implies $\text{ax}(\mu) \subseteq \text{ax}(\nu)$ and $\text{ax}^*(\mu) = \text{ax}^*(\nu)$. \square

Proof. It is enough to check one-step reductions. If $\mu \rightarrow_{\beta} \nu$, we conclude by Proposition 2.3. If $\mu \rightarrow_{\varepsilon} \nu$, we have that $\mu = C[\mu_0]$ and $\nu = C[\mathbf{E}_n]$ for some context C and blind net μ_0 with n free ports; precisely because μ_0 is blind, the axioms disappearing from μ are not observable; moreover, ν does not contain new observable axioms, so $\text{ax}(\mu) = \text{ax}(\nu)$, and we conclude $\text{ax}^*(\mu) = \text{ax}^*(\nu)$ by confluence of $\beta\varepsilon$ -reduction (Proposition 2.27). \square

Corollary 2.31. μ is $\beta\varepsilon$ -normalizable iff $\text{ax}^*(\mu)$ is finite.

Proof. The forward implication is an immediate consequence of Proposition 2.30. For the converse, $\text{ax}^*(\mu)$ finite and Proposition 2.30 imply that any reduction starting from μ stumbles upon a net μ' such that $\text{ax}(\mu') = \text{ax}^*(\mu)$. This means that all subnets of μ' containing active pairs are blind, *i.e.*, they do not produce further observable axioms. There is of course at most a finite number of such subnets, so μ' reduces in finitely many ε -steps to a $\beta\varepsilon$ -normal net. \square

Up to now we have five congruences strictly extending β -equivalence: $\simeq_{\beta\eta}$, $\simeq_{\beta\varepsilon}$, $\simeq_{\beta\eta\varepsilon}$, \cong , and \simeq (the fact that these last two strictly extend \simeq_{β} is an immediate consequence of Definition 2.5). We shall see that the last four congruences form a sequence of strictly stronger equivalences:

$$\simeq_{\beta} \subsetneq \simeq_{\beta\varepsilon} \subsetneq \simeq_{\beta\eta\varepsilon} \subsetneq \cong \subsetneq \simeq .$$

The first two strict inclusions are obvious. We shall give a semantic proof of the third inclusion (Corollary 4.20 of the full abstraction Theorem 4.14); Fig. 19 shows that it is strict. We already established the fourth inclusion in Proposition 2.23; Fig. 15 shows that the inclusion is strict, as an application of the full abstraction Theorem 4.18.

3. DENOTATIONAL SEMANTICS

3.1. Edifices. In what follows, $\mathcal{C} = \{\mathbf{p}, \mathbf{q}\}^{\mathbb{N}}$ is the set of infinite binary words, ranged over by x, y . As in the case of finite words, the elements of $\mathcal{C} \times \mathcal{C}$ will be denoted by $x \otimes y$, and ranged over by u, v, w . Given two words or pairs of words s, u , where s is finite and u may be infinite, we say that s is *prefix* of u iff there exists u' such that $u = su'$.

Definition 3.1 (Pillar, arch, edifice, vault). Let $I \subseteq \mathbb{N}$, and set $\mathcal{P}_I = \mathcal{C} \times \mathcal{C} \times I$. 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 .

Set $\vec{\mathcal{A}}_I = \mathcal{P}_I \times \mathcal{P}_I$, and, given $(\xi, v), (\xi', v') \in \vec{\mathcal{A}}_I$, define $(\xi, v) \rightleftharpoons (\xi', v')$ iff $\xi' = v$ and $v' = \xi$, or $\xi' = \xi$ and $v' = v$. We then set $\mathcal{A}_I = \vec{\mathcal{A}}_I / \rightleftharpoons$. 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}, \mathbf{b} . An arch is said to be *based* at the unordered pair where its two pillars are based.

An *edifice* is a set of arches; edifices are ranged over by $\mathfrak{E}, \mathfrak{F}$. A *vault* is an edifice \mathfrak{A} such that there exists an address $\mathfrak{r} = s @ i \frown t @ j$ such that

$$\mathfrak{A} = \{sw @ i \frown tw @ j \mid w \in \mathcal{C} \times \mathcal{C}\} .$$

The address \mathfrak{r} is said to *generate* \mathfrak{A} . We denote by $\mathfrak{A}(\mathfrak{r})$ the vault generated by \mathfrak{r} .

We now introduce two special kinds of edifices, which will be useful in the sequel.

Definition 3.2 (Uniform edifice). If \mathfrak{E} is an edifice and $\mathfrak{a} \in \mathfrak{E}$, \mathfrak{a} is said to be *uniform in \mathfrak{E}* iff there exists a vault $\mathfrak{A} \subseteq \mathfrak{E}$ such that $\mathfrak{a} \in \mathfrak{A}$. An edifice is *uniform* if all of its arches are uniform.

Proposition 3.3. *An edifice is uniform iff it is a union of vaults.*

Proof. That a union of vaults is uniform is obvious. For the converse, take a uniform edifice \mathfrak{E} , and let $\mathfrak{a} \in \mathfrak{E}$. By definition, there exists a vault contained in \mathfrak{E} which contains \mathfrak{a} ; we call such vault $\mathfrak{A}_{\mathfrak{a}}$. Then, it is easy to check that $\mathfrak{E} = \bigcup_{\mathfrak{a} \in \mathfrak{E}} \mathfrak{A}_{\mathfrak{a}}$. \square

Definition 3.4 (Coherence, simple edifice). If $\mathfrak{a} = \xi \frown v$ is an arch, we define its *support* to be $|\mathfrak{a}| = \{\xi, v\}$. Then, given $\mathfrak{a}, \mathfrak{a}' \in \mathcal{A}$, we say that \mathfrak{a} and \mathfrak{a}' are *coherent*, and we write $\mathfrak{a} \Xi \mathfrak{a}'$, iff $|\mathfrak{a}| \cap |\mathfrak{a}'|$ is either empty or of cardinality 2. An edifice \mathfrak{E} is *simple* iff it is a clique with respect to coherence, *i.e.*, for all $\mathfrak{a}, \mathfrak{a}' \in \mathfrak{E}$, $\mathfrak{a} \Xi \mathfrak{a}'$.

Note that, although obviously symmetric, coherence is not reflexive: all arches of the form $\xi \frown \xi$, which we may refer to as degenerated, are not coherent with themselves.

Edifices are naturally endowed with a *trace* operation. We shall see that this operation closely corresponds to the execution formula of the GoI (*cf.* Sect. 2.2); it is also reminiscent of the notion of composition of strategies in games semantics.

Definition 3.5 (Feedback function, trace sequence). A *feedback function* σ is a fixpoint-free partial involution on \mathbb{N} of finite domain. In other words, $\sigma(i)$ is defined for finitely many $i \in \mathbb{N}$, and in that case $\sigma(i) \neq i$, and $\sigma^2(i) = i$. We denote by $\text{dom } \sigma$ the domain of σ .

Let σ be a feedback function, let \mathfrak{E} be an edifice, and let $\mathfrak{s} = \mathfrak{s}_1, \dots, \mathfrak{s}_n$ be a non-empty finite sequence of arches of \mathfrak{E} , for which we set, for $1 \leq k \leq n$, $\mathfrak{s}_k = u_k @ i_k \frown v_k @ j_k$. We say that \mathfrak{s} is a *trace sequence of \mathfrak{E} along σ* iff:

chain: for all $2 \leq k \leq n$, $j_{k-1} \in \text{dom } \sigma$ and $i_k = \sigma(j_{k-1})$;

match: for all $2 \leq k \leq n$, $u_k = v_{k-1}$.

A trace sequence is *visible* if it further satisfies $i_1, j_n \notin \text{dom } \sigma$.

The *length* of a trace sequence \mathfrak{s} is denoted by $|\mathfrak{s}|$. We denote by $\text{seq}_{\sigma}(\mathfrak{E})$ the set of trace sequences of \mathfrak{E} along σ . If $\mathfrak{s} \in \text{seq}_{\sigma}(\mathfrak{E})$ such that $|\mathfrak{s}| = n$, we define the *arch generated by \mathfrak{s}* as $\mathfrak{a}(\mathfrak{s}) = u_1 @ i_1 \frown v_n @ j_n$.

Observe that, if \mathfrak{E} is an edifice and σ a feedback function, then $\mathfrak{s} = \mathfrak{s}_1, \dots, \mathfrak{s}_n \in \text{seq}_{\sigma}(\mathfrak{E})$ implies $\mathfrak{s}' = \mathfrak{s}_n, \dots, \mathfrak{s}_1 \in \text{seq}_{\sigma}(\mathfrak{E})$, $\mathfrak{a}(\mathfrak{s}') = \mathfrak{a}(\mathfrak{s})$, and \mathfrak{s}' is visible iff \mathfrak{s} is. Intuitively, this reflects the fact that a maximal path in a net can always be “reversed”. In Sect. 3.2 we shall formalize the relation between visible trace sequences and maximal paths. We also remark that the role of non-visible trace sequences will be purely technical: their purpose is to allow proofs by induction on the length of sequences. In fact, visible trace sequences are not suitable for such proof technique, because an initial or final segment of a visible trace sequence is never visible.

Definition 3.6 (Trace). Let \mathfrak{E} be a set of arches, and σ a feedback function. We define the *trace* of \mathfrak{E} along σ as

$$\text{Tr}_{\sigma}(\mathfrak{E}) = \{\mathfrak{a}(\mathfrak{s}) \mid \mathfrak{s} \in \text{seq}_{\sigma}(\mathfrak{E}), \mathfrak{s} \text{ visible}\}.$$

The trace is obviously monotonic:

Lemma 3.7 (Monotonicity of the trace). *Let $\mathfrak{E}, \mathfrak{F}$ be edifices, and let σ be a feedback function. Then, $\mathfrak{E} \subseteq \mathfrak{F}$ implies $\text{Tr}_\sigma(\mathfrak{E}) \subseteq \text{Tr}_\sigma(\mathfrak{F})$.*

Proof. Obvious. \square

In the following, if σ, σ' are feedback functions of disjoint domain, we denote by $\sigma \uplus \sigma'$ the function defined by

$$(\sigma \uplus \sigma')(i) = \begin{cases} \sigma(i) & \text{if } i \in \text{dom } \sigma \\ \sigma'(i) & \text{if } i \in \text{dom } \sigma' \\ \text{undefined} & \text{otherwise} \end{cases},$$

which is obviously a feedback function.

Lemma 3.8 (Associativity of the trace). *Let \mathfrak{E} be an edifice, and let σ, σ' be feedback functions of disjoint domain. Then,*

$$\text{Tr}_\sigma(\text{Tr}_{\sigma'}(\mathfrak{E})) = \text{Tr}_{\sigma \uplus \sigma'}(\mathfrak{E}).$$

Proof. We start by establishing the inclusion from left to right. Let $\mathfrak{a} \in \text{Tr}_\sigma(\text{Tr}_{\sigma'}(\mathfrak{E}))$, and let $\mathfrak{s} \in \text{seq}_\sigma(\text{Tr}_{\sigma'}(\mathfrak{E}))$ be the trace sequence generating \mathfrak{a} , with $|\mathfrak{s}| = m$. By definition, if we let $1 \leq k \leq m$, each \mathfrak{s}_k is generated by a trace sequence $\mathfrak{s}^k \in \text{seq}_{\sigma'}(\mathfrak{E})$, with $|\mathfrak{s}^k| = n_k$. Then, it is not hard to see that $\mathfrak{s}' = \mathfrak{s}_1^1, \dots, \mathfrak{s}_{n_1}^1, \mathfrak{s}_1^2, \dots, \mathfrak{s}_{n_{m-1}}^{m-1}, \mathfrak{s}_1^m, \dots, \mathfrak{s}_{n_m}^m$ is a visible trace sequence of \mathfrak{E} along $\sigma \uplus \sigma'$, such that $\mathfrak{a}(\mathfrak{s}') = \mathfrak{a}(\mathfrak{s})$.

For the reverse inclusion, let $\mathfrak{a} \in \text{Tr}_{\sigma \uplus \sigma'}(\mathfrak{E})$ and let $\mathfrak{s} \in \text{seq}_{\sigma \uplus \sigma'}(\mathfrak{E})$ be the trace sequence generating \mathfrak{a} , with $|\mathfrak{s}| = n$, and $\mathfrak{s}_k = u_k @ i_k \frown v_k @ j_k$ for $1 \leq k \leq n$. We say that k is a *breaking point* of \mathfrak{s} if $j_k \notin \text{dom } \sigma'$. Let now $B = \{k_1 < \dots < k_m\}$ be the set of breaking points of \mathfrak{s} , ordered from the smallest to the greatest. We partition \mathfrak{s} into m sequences, as follows: $\mathfrak{s}^1 = \mathfrak{s}_1, \dots, \mathfrak{s}_{k_1}$, and, for $2 \leq h \leq m$, $\mathfrak{s}^h = \mathfrak{s}_{k_{h-1}+1}, \dots, \mathfrak{s}_{k_h}$. Once again, it is not hard to check that, for all $1 \leq h \leq m$, $\mathfrak{s}^h \in \text{seq}_{\sigma'}(\mathfrak{E})$, and that $\mathfrak{s}' = \mathfrak{a}(\mathfrak{s}^1), \dots, \mathfrak{a}(\mathfrak{s}^m) \in \text{seq}_\sigma(\text{Tr}_{\sigma'}(\mathfrak{E}))$, with $\mathfrak{a}(\mathfrak{s}') = \mathfrak{a}(\mathfrak{s})$. \square

The trace of a simple edifice has a very nice property, namely that each arch in it is generated by a unique trace sequence:

Lemma 3.9. *Let \mathfrak{E} be a simple edifice, let σ be a feedback function, and let $\mathfrak{s}, \mathfrak{s}' \in \text{seq}_\sigma(\mathfrak{E})$, such that $\mathfrak{a}(\mathfrak{s}) = \zeta_1 \frown \zeta_2$ and $\mathfrak{a}(\mathfrak{s}') = \zeta'_1 \frown \zeta'_2$. Then, $\zeta_1 = \zeta'_1$ implies $\mathfrak{s} = \mathfrak{s}'$; in particular, $\mathfrak{a}(\mathfrak{s}) = \mathfrak{a}(\mathfrak{s}')$ implies $\mathfrak{s} = \mathfrak{s}'$.*

Proof. Let k be the smallest integer such that $\mathfrak{s}_k \neq \mathfrak{s}'_k$, and let $\mathfrak{s}_k = \xi \frown v$, $\mathfrak{s}'_k = \xi' \frown v'$. If $k = 0$, $\xi = \zeta_1$ and $\xi' = \zeta'_1$, so $\xi = \xi'$ by hypothesis; if $k > 0$, the chain and match conditions also imply $\xi = \xi'$, because we supposed $\mathfrak{s}_{k-1} = \mathfrak{s}'_{k-1}$. But \mathfrak{E} is simple, so $\mathfrak{s}_k \Xi \mathfrak{s}'_k$, which further implies $v = v'$. Then, $\mathfrak{s}_k = \mathfrak{s}'_k$, a contradiction. \square

On the other hand, uniform edifices are preserved by the trace:

Proposition 3.10. *Let \mathfrak{E} be a uniform edifice, and let σ be a feedback function. Then, $\text{Tr}_\sigma(\mathfrak{E})$ is uniform.*

Proof. Let $\mathfrak{F} = \{\mathfrak{a}(\mathfrak{s}) \mid \mathfrak{s} \in \text{seq}_\sigma(\mathfrak{E})\}$. Let \mathfrak{s} be a trace sequence; we shall prove, by induction on $|\mathfrak{s}|$, that $\mathfrak{a}(\mathfrak{s})$ is uniform in \mathfrak{F} , i.e., that there exists $\mathfrak{r} = s @ i \frown t @ j$ such that $\mathfrak{a} \in \mathfrak{A}(\mathfrak{r}) \subseteq \mathfrak{F}$. This will be enough to prove the result; in fact, the trace sequences generating the arches of $\mathfrak{A}(\mathfrak{r})$ are all visible if \mathfrak{s} is visible, i.e., $\mathfrak{A}(\mathfrak{r}) \subseteq \text{Tr}_\sigma(\mathfrak{E})$, because this depends only on i and j , which are the same in all such sequences.

For the base case, $|\mathfrak{s}| = 1$, so \mathfrak{s} consists of a single arch, and the result is a trivial consequence of the uniformity of \mathfrak{E} . The inductive case is $|\mathfrak{s}| = n > 1$. Let $\mathfrak{s}' = \mathfrak{s}_1, \dots, \mathfrak{s}_{n-1}$. By induction hypothesis, $\mathfrak{a}(\mathfrak{s}')$ is uniform in \mathfrak{F} : we have an address \mathfrak{r}' such that $\mathfrak{a}(\mathfrak{s}') \in \mathfrak{A}(\mathfrak{r}') \subseteq \mathfrak{F}$. In particular, there exist $s, r_0 \in \mathcal{W} \times \mathcal{W}$ and $w_0 \in \mathcal{C} \times \mathcal{C}$ such that the first and last arches of \mathfrak{s}' are respectively of the form

$$\begin{aligned} sw_0 @ i \frown v, \\ \xi \frown r_0 w_0 @ h, \end{aligned}$$

and, for all $w \in \mathcal{C} \times \mathcal{C}$, there exists a trace sequence \mathfrak{s}^w whose first and last arches are respectively of the form

$$\begin{aligned} sw @ i \frown v, \\ \xi \frown r_0 w @ h. \end{aligned}$$

Similarly, \mathfrak{E} is uniform, so there exist $r_1, t \in \mathcal{W} \times \mathcal{W}$ and $w_1 \in \mathcal{C} \times \mathcal{C}$ such that

$$\mathfrak{s}_n = r_1 w_1 @ l \frown t w_1 @ j,$$

with $\sigma(h) = l$, and, for all $w \in \mathcal{C} \times \mathcal{C}$, the arch

$$\mathfrak{a}^w = r_1 w @ l \frown t w @ j$$

is also in \mathfrak{E} . Now, by the match condition, we have $r_0 w_0 = r_1 w_1$, which means that r_0, r_1 are one prefix of the other. Suppose $r_0 = r_1 r'$, and consider the trace sequences obtained by appending $\mathfrak{a}^{r'w}$ to \mathfrak{s}^w ; these generate all the arches of the form $r_0 w @ i \frown t r' w @ j$, among which there is $\mathfrak{a}(\mathfrak{s})$, which is therefore uniform in \mathfrak{F} . The other case is $r_1 = r_0 r'$; then, consider the sequences obtained by appending \mathfrak{a}^w to $\mathfrak{s}^{r'w}$. As above, these prove that $\mathfrak{a}(\mathfrak{s})$ is uniform in \mathfrak{F} . \square

3.2. Nets as edifices. The basic idea to assign an edifice to a net is that arches model observable axioms/paths.³ In fact, we have already seen that an observable axiom may be conveniently represented by an unordered pair of couples of the form $s @ i$, where s is the address of a leaf and i a free port. A pillar contains the same information; the need for infinite words arises from η -expansion (the η_0 equation of Fig. 4), which can be applied indefinitely, as in the pure λ -calculus.

Definition 3.11 (Edifice of a net). Let μ be a net. We associate an edifice with μ , denoted by $\mathfrak{E}(\mu)$, as follows:

$$\mathfrak{E}(\mu) = \bigcup_{\mathfrak{r} \in \text{ax}^*(\mu)} \mathfrak{A}(\mathfrak{r}).$$

The union of Definition 3.11 is actually disjoint, and the resulting edifice is simple:

Lemma 3.12. *Let μ be a net, and let $\mathfrak{r}, \mathfrak{r}' \in \text{ax}^*(\mu)$. Then*

- (1) $\mathfrak{A}(\mathfrak{r}) \cup \mathfrak{A}(\mathfrak{r}')$ is simple, and hence $\mathfrak{E}(\mu)$ is simple;
- (2) $\mathfrak{r} \neq \mathfrak{r}'$ implies $\mathfrak{A}(\mathfrak{r}) \cap \mathfrak{A}(\mathfrak{r}') = \emptyset$.

³Graphically (Fig. 10), observable axioms/paths look like arches, hence the terminology.

Proof. We start by observing that two arches in the same vault are coherent, because in an address $s @ i \frown t @ j$ we always have $s \neq t$. This proves point (1) in case $\mathfrak{x} = \mathfrak{x}'$, so we may suppose $\mathfrak{x} \neq \mathfrak{x}'$. By confluence of β -reduction and by Proposition 2.3, $\mathfrak{x}, \mathfrak{x}' \in \text{ax}^*(\mu)$ implies that there exists a net μ' such that $\mu \rightarrow_{\beta}^* \mu'$ and μ' contains two observable axioms of address $\mathfrak{x}, \mathfrak{x}'$, respectively, which are distinct because $\mathfrak{x} \neq \mathfrak{x}'$. Now, put $\mathfrak{x} = s @ i \frown t @ j$ and $\mathfrak{x}' = s' @ i' \frown t' @ j'$, and take $\mathfrak{a} \in \mathfrak{A}(\mathfrak{x})$, $\mathfrak{a}' \in \mathfrak{A}(\mathfrak{x}')$. By definition of vault, we have $\mathfrak{a} = sw @ i \frown tw @ j$, $\mathfrak{a}' = s'w' @ i' \frown t'w' @ j'$, for some $w, w' \in \mathcal{C} \times \mathcal{C}$. Suppose $sw @ i = s'w' @ i'$; we would obtain that s is a prefix of s' , or vice versa. But this is absurd, because s, s' are addresses of distinct leaves of μ' . This proves point (2); for point (1), simply note that the same holds for t, t' . To see that $\mathfrak{E}(\mu)$ is itself simple, note that a union of pairwise coherent simple edifices is obviously simple. \square

Note that if μ has n free ports and σ is a feedback function whose domain is included in $\{1, \dots, n\}$, then σ defines a feedback context for μ : it is the one connecting the free port i to the free port $\sigma(i)$, or leaving it free if $\sigma(i)$ is undefined. Conversely, each feedback context for a net with n free ports defines a feedback function of domain included in $\{1, \dots, n\}$. Hence, we shall use σ to range over both feedback functions and feedback contexts, and make no distinction between the two, speaking more generally of a “feedback” σ for a net μ .

The next result shows how the trace construction is related to the execution formula of the GoI. In fact, in Sect. 2.2 we mentioned that this latter is invariant under reduction: if $\mu = \sigma[\nu] \rightarrow_{\beta} \sigma'[\nu'] = \mu'$ with ν, ν' cut-free and μ, μ' total, then the formula of Theorem 2.14 applied to $\llbracket \sigma \rrbracket, \llbracket \nu \rrbracket$ or $\llbracket \sigma' \rrbracket, \llbracket \nu' \rrbracket$ yields the same result. Once again, our work generalizes this to non-total nets.

Proposition 3.13 (Invariance of the trace). *Let $\mu \rightarrow_{\beta} \mu'$, and let $\mu = \sigma[\nu]$ and $\mu' = \sigma'[\nu']$ according to the Decomposition Lemma 1.6. Then, $\text{Tr}_{\sigma}(\mathfrak{E}(\nu)) = \text{Tr}_{\sigma'}(\mathfrak{E}(\nu'))$.*

Proof. The proof is a bit too long and not interesting enough to be included here. We prefer to give it in Appendix A. \square

Let ν be a cut-free net, let σ be a feedback for ν , and let $\mathfrak{a} \in \text{Tr}_{\sigma}(\mathfrak{E}(\nu))$. Part (1) of Lemma 3.12 and Lemma 3.9 guarantee us that \mathfrak{a} induces a unique $\mathfrak{s} \in \text{seq}_{\sigma}(\mathfrak{E}(\nu))$ such that $\mathfrak{a}(\mathfrak{s}) = \mathfrak{a}$. If we let $\mathfrak{s} = \mathfrak{s}_1, \dots, \mathfrak{s}_n$, we see that each \mathfrak{s}_k determines an axiom of ν , which is unique by part (2) of Lemma 3.12. Hence, \mathfrak{s} induces a sequence of observable paths; thanks to the chain condition, these are all composable through σ , and, by the visibility condition, the first and last paths start and end at a free port of $\sigma[\nu]$. Therefore, their composition forms a maximal path of $\sigma[\nu]$.

To sum up, we found out that each arch \mathfrak{a} of the trace of $\mathfrak{E}(\nu)$ along σ determines a unique maximal path of $\sigma[\nu]$; in what follows, we shall denote this path by $\phi(\mathfrak{a})$.

Lemma 3.14. *Let ν be a cut-free net, let σ be a feedback for ν , let \mathfrak{x} be an address, and let $\mathfrak{a}, \mathfrak{a}' \in \mathfrak{A}(\mathfrak{x}) \subseteq \text{Tr}_{\sigma}(\mathfrak{E}(\nu))$. Then, $\phi(\mathfrak{a}) = \phi(\mathfrak{a}')$.*

Proof. Let $\mathfrak{s}_1, \dots, \mathfrak{s}_n$ and $\mathfrak{s}'_1, \dots, \mathfrak{s}'_{n'}$ be the trace sequences generating \mathfrak{a} and \mathfrak{a}' , respectively. We decompose $\phi(\mathfrak{a})$ and $\phi(\mathfrak{a}')$ into observable paths of ν , following the $\mathfrak{s}_k, \mathfrak{s}'_k$, as described in the remarks above, and we obtain $\phi_1, \dots, \phi_n, \phi'_1, \dots, \phi'_{n'}$. Let k be the smallest integer such that $\phi_k \neq \phi'_k$, and let $\mathfrak{s}_k = u @ i \frown v @ j$, $\mathfrak{s}'_k = u' @ i' \frown v' @ j'$. Suppose $k = 0$; if we put $\mathfrak{x} = s @ i \frown t @ j$, we have, for some $w, w' \in \mathcal{C} \times \mathcal{C}$, $u = sw$, $u' = sw'$, $i = i'$, and moreover for all $w_0 \in \mathcal{C} \times \mathcal{C}$ we have an arch of the form $sw_0 @ i \frown tw_0 @ j$ in $\text{Tr}_{\sigma}(\mathfrak{E}(\nu))$. By Definition 3.5, this latter implies that, for all $w_0 \in \mathcal{C} \times \mathcal{C}$, there is some arch of the


form $sw_0 @ i \frown v$ in $\mathfrak{E}(\nu)$. But, by looking at Definition 3.11, we see that this is possible only if there is an axiom of ν of address $s_0 @ i \frown t' @ j'$, with s_0 a prefix of s . Then, the arches \mathfrak{s}_0 and \mathfrak{s}'_0 correspond to the same observable axiom, and $\phi_0 = \phi'_0$. Therefore, we must have $k > 0$. In this case, since $\phi_{k-1} = \phi'_{k-1}$, we have $\mathfrak{s}_{k-1} = \xi \frown v_{k-1} @ p$ and $\mathfrak{s}'_{k-1} = \xi' \frown v'_{k-1} @ p$, and the chain condition implies $i = i' = \sigma(p)$. But v_{k-1}, v'_{k-1} come from the same observable axiom, so they have a common prefix. By the match condition, u, u' have a common prefix too; but this would be impossible if $\phi_k \neq \phi'_k$, because two distinct observable path generate pillars which have no common prefix (*cf.* the proof of Lemma 3.12). So we must conclude $\phi_k = \phi'_k$, a contradiction. \square

Lemma 3.15. *Let ν be a cut-free net, let σ be a feedback for ν , and let i, j be free ports of $\sigma[\nu]$. Then, the following are equivalent:*

- (1) $\phi \in \text{mp}_{i,j}(\sigma[\nu])$, $\llbracket w(\phi) \rrbracket = ts^*$;
- (2) $\mathfrak{A}(s @ i \frown t @ j) \subseteq \text{Tr}_\sigma(\mathfrak{E}(\nu))$.

Proof. We start with (1) implies (2), noting first that $\llbracket w(\phi) \rrbracket = ts^*$ implies, by Lemma 2.17, that ϕ is an execution path; therefore, we reason by induction on $v(\phi)$ (the value of ϕ , Definition 2.10), using Lemma 2.16.

- $v(\phi) = 0$. We know that ϕ is an observable path. The situation can be schematically depicted as follows:



In the above picture, ϕ goes “from left to right”, and $\omega_1, \dots, \omega_n$ are n proper axioms of ν . If, for $1 \leq k \leq n$, we put $\text{addr}_\nu(\omega_k) = \mathfrak{r}_k = s_k @ i_k \frown t_k @ j_k$, from the above picture we deduce that $i_1 = i$, $j_n = j$, $i, j \notin \text{dom } \sigma$, $i_{k+1} = \sigma(j_k)$ for all $1 \leq k < n$, $s_k = \mathbf{1} \otimes \mathbf{1}$ for all $2 \leq k \leq n$, $s_1 = s$, and $t_n \cdots t_1 = t$. Now, given any $w \in \mathcal{C} \times \mathcal{C}$ and $1 < k < n$, define $w_k = t_{k-1} \cdots t_1 w$, and put

$$\begin{aligned} \mathfrak{s}_1 &= sw @ i \frown t_1 w @ j_1, \\ \mathfrak{s}_k &= w_k @ i_k \frown t_k w_k @ j_k, \text{ for } 2 \leq k \leq n. \end{aligned}$$

It is easy to check that $\mathfrak{s} = \mathfrak{s}_1, \dots, \mathfrak{s}_n$ is a trace sequence of $\mathfrak{E}(\nu)$ along σ such that $\mathfrak{a}(\mathfrak{s}) = sw @ i \frown tw @ j$, as desired.

- $v(\phi) > 0$. We β -reduce $\sigma[\nu]$ along ϕ and obtain $\sigma'[\nu]$ in which the residue ϕ' of ϕ is such that $v(\phi') < v(\phi)$, so that the induction hypothesis applies to ϕ' . By Propositions 2.12 and 2.13, we have $\llbracket w(\phi') \rrbracket = ts^*$, so the induction hypothesis and Proposition 3.13 give us $\mathfrak{A}(s @ i \frown t @ j) \subseteq \text{Tr}_{\sigma'}(\nu') = \text{Tr}_\sigma(\nu)$.

We now consider (2) implies (1). By Lemma 3.14, any arch of $\mathfrak{A}(s @ i \frown t @ j)$ yields the same $\phi \in \text{mp}_{i,j}(\sigma[\nu])$. Hence, we need only check that $\llbracket w(\phi) \rrbracket = ts^*$; we do this again by induction on $v(\phi)$.

- $v(\phi) = 0$. We know that ϕ is an observable path of $\sigma[\nu]$, which we may decompose in several observable paths ϕ_1, \dots, ϕ_n of ν , as above. Now, as we remarked in the proof of Lemma 3.14, $\mathfrak{A}(s @ i \frown t @ j) \subseteq \text{Tr}_\sigma(\nu)$ is only possible if $s @ i$ and $t @ j$ are the addresses of the leaves of the trees resulting in the composition of ϕ_1, \dots, ϕ_n , similarly to the drawing above used in the previous part of the proof. So $\llbracket w(\phi) \rrbracket = ts^*$, as desired.

- $v(\phi) > 0$. We know that ϕ crosses an active pair; however, this time we must prove that, by reducing it, we obtain a net in which ϕ has a residue. If this were not the case, by Lemma 2.17 and Proposition 2.12 the only possibility is that the active pair corresponds to a clash in $w(\phi)$, *i.e.*, we have $w(\phi) = A'A_0^*B_0B'$ for some monomials A', B' and some positive atoms A_0, B_0 , such that, for example, $A_0 = \mathbf{c}$ and $B_0 = \mathbf{d}$. This would yield the presence of two observable axioms of ν of addresses $\mathfrak{r} = s@p \frown \mathbf{q}b \otimes b'@q$ and $\mathfrak{r}' = \mathbf{p}a \otimes a'@q' \frown t@p'$, such that $\sigma(q) = q'$, *i.e.*, the free ports q, q' of ν are connected by σ in $\sigma[\nu]$. Then, it is easy to see that no arch of $\mathfrak{E}(\nu)$ generated by \mathfrak{r} and \mathfrak{r}' could match: indeed, there are no $x, y, x', y' \in \mathcal{C}$ such that $\mathbf{q}bx \otimes b'y = \mathbf{p}ax' \otimes a'y'$. But this is absurd, because by hypothesis ϕ comes from an arch of $\text{Tr}_\sigma(\nu)$, which in turn comes from a unique trace sequence, and trace sequences satisfy the match condition.

Therefore, we now know that we may reduce the active pair crossed by ϕ , obtaining $\sigma[\nu] \rightarrow_\beta \sigma'[\nu']$ such that ϕ has a residue ϕ' in $\sigma'[\nu']$. Now, by Proposition 3.13, we have $\text{Tr}_{\sigma'}(\mathfrak{E}(\nu')) = \text{Tr}_\sigma(\mathfrak{E}(\nu))$, so $\mathfrak{A}(s@i \frown t@j) \subseteq \text{Tr}_{\sigma'}(\mathfrak{E}(\nu'))$. Thanks to Proposition 2.12 and Lemma 2.7, we can apply the induction hypothesis to ϕ' , and obtain $\llbracket w(\phi') \rrbracket = ts^*$. But, by Proposition 2.13, $\llbracket w(\phi) \rrbracket = \llbracket w(\phi') \rrbracket$, and we are done. \square

Proposition 3.16. *Let ν be a cut-free net, and σ a feedback for ν . Then*

$$\mathfrak{E}(\sigma[\nu]) = \text{Tr}_\sigma(\mathfrak{E}(\nu)).$$

Proof. We start with the inclusion from left to right. Let $\mathfrak{a} \in \mathfrak{E}(\sigma[\nu])$. By Definition 3.11, there is an address $\mathfrak{r} = s@i \frown t@j \in \text{ax}^*(\sigma[\nu])$ (which is also unique by Lemma 3.12) such that $\mathfrak{a} \in \mathfrak{A}(\mathfrak{r})$; by Proposition 2.18, there exists $\phi \in \text{mp}_{i,j}(\sigma[\nu])$ such that $\llbracket w(\phi) \rrbracket = ts^*$, and by Lemma 3.15 $\mathfrak{A}(\mathfrak{r}) \subseteq \text{Tr}_\sigma(\mathfrak{E}(\nu))$.

For the inclusion from right to left, let $\mathfrak{a} \in \text{Tr}_\sigma(\mathfrak{E}(\nu))$. By Lemma 3.3, $\mathfrak{E}(\nu)$ is uniform, so by Proposition 3.10 $\text{Tr}_\sigma(\mathfrak{E}(\nu))$ is also uniform. Then, there must be an address $\mathfrak{r} = s@i \frown t@j$ such that $\mathfrak{a} \in \mathfrak{A}(\mathfrak{r}) \subseteq \text{Tr}_\sigma(\mathfrak{E}(\nu))$. We can thus apply Lemma 3.15 and Proposition 2.18, which, by Definition 3.11, give us $\mathfrak{A}(\mathfrak{r}) \subseteq \mathfrak{E}(\sigma[\nu])$. \square

Compare Proposition 3.16 with Theorem 2.14: basically, the trace construction can be seen as an extension of the execution formula, which works in all cases, even when $\sigma[\nu]$ is not total. Something similar happens when one formulates the GoI in categorical terms, using certain traced monoidal categories, as shown by Haghverdi and Scott [HS04]. Thanks to the associativity of the trace, we can straightforwardly extend Proposition 3.16 to arbitrary nets:

Corollary 3.17. *For any net μ and feedback σ for μ , $\mathfrak{E}(\sigma[\mu]) = \text{Tr}_\sigma(\mathfrak{E}(\mu))$.*

Proof. By the Decomposition Lemma 1.6, we know that $\mu = \sigma'[\nu]$ for some cut-free net ν and feedback context σ' . Then, using Proposition 3.16 and Lemma 3.8, we have

$$\begin{aligned} \mathfrak{E}(\sigma[\mu]) &= \mathfrak{E}(\sigma[\sigma'[\nu]]) = \mathfrak{E}((\sigma \uplus \sigma')[\nu]) = \text{Tr}_{\sigma \uplus \sigma'}(\mathfrak{E}(\nu)) = \\ &= \text{Tr}_\sigma(\text{Tr}_{\sigma'}(\mathfrak{E}(\nu))) = \text{Tr}_\sigma(\mathfrak{E}(\sigma'[\nu])) = \text{Tr}_\sigma(\mathfrak{E}(\mu)), \end{aligned}$$

where we used the notation $\sigma \uplus \sigma'$ also to denote the “union” of the nets σ, σ' . \square

3.3. A denotational semantics. As explained in the introduction, a denotational semantics is an interpretation of the syntax transforming certain given syntactic equivalences into denotational equalities. In a syntax such as the λ -calculus, or the symmetric combinators, the typical equivalence to be modelled is that induced by β -reduction. As proposed for example by Girard [Gir91], we may describe a denotational semantics of a syntax with a reduction relation \rightsquigarrow and an internal notion of context as an interpretation satisfying at least the following:

- invariance:** for any two syntactic objects a, b , $a \rightsquigarrow b$ implies that a and b are denotationally equal;
- contextuality:** the semantics induces a congruence on the syntax, *i.e.*, if two syntactic objects a, b are denotationally equal, then for every context C , $C[a]$ and $C[b]$ are also denotationally equal.

In this section, we shall prove that our interpretation of nets as edifices satisfies these two requirements. Actually, we shall see that, instead of just β -reduction, edifices model $\simeq_{\beta\eta\varepsilon}$.

Lemma 3.18. *Let \rightarrow_{η_0} be the relation of Lemma 1.17, *i.e.*, a single application of the η_0 equation (Fig. 4), oriented from left to right, in which the wire on the right member is an axiom. Moreover, let $\simeq_{\eta_1}^1$ be the contextual closure of the η_1 equation (Fig. 4), *i.e.*, the restriction of \simeq_{η_1} to just one application of the equation. For all nets μ, ν , we have:*

- (1) $\mu \rightarrow_{\beta\varepsilon} \nu$ implies $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$;
- (2) $\mu \rightarrow_{\eta_0} \nu$ implies $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$;
- (3) $\mu \simeq_{\eta_1}^1 \nu$ implies $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$.

Proof. Point (1) is an immediate consequence of Proposition 2.30 and of Definition 3.11.

For what concerns point (2), considering the Decomposition Lemma 1.6 we obtain $\mu = \sigma[\mu_0]$ and $\nu = \sigma[\nu_0]$ such that μ_0, ν_0 are cut-free and $\mu_0 = C[o]$, $\nu_0 = C[\omega]$, where o is a net matching the left member of the η_0 equation of Fig. 4, and ω is an axiom of ν_0 . We suppose that o consists of two δ cells, the case of two ζ cells being perfectly analogous. Now, remark that, if $\mathfrak{r}_0 = \text{addr}_{\nu_0}(\omega) = a \otimes b @ i \frown c \otimes d @ j$, we have

$$\text{ax}(\mu_0) = \{\mathfrak{r} \in \text{ax}(\nu_0) \mid \mathfrak{r} \neq \mathfrak{r}_0\} \cup \{\mathfrak{a}\mathbf{p} \otimes b @ i \frown \mathfrak{c}\mathbf{p} \otimes d @ j, \mathfrak{a}\mathbf{q} \otimes b @ i \frown \mathfrak{c}\mathbf{q} \otimes d @ j\}.$$

Moreover, note that, since μ, ν_0 are cut-free, $\text{ax}^*(\mu_0) = \text{ax}(\mu_0)$ and $\text{ax}^*(\nu_0) = \text{ax}(\nu_0)$. So take $\mathfrak{a} \in \mathfrak{E}(\mu_0)$. If \mathfrak{a} is generated by an address in the left term of the above union, then we clearly have $\mathfrak{a} \in \mathfrak{E}(\nu_0)$. Otherwise, we have for example $\mathfrak{a} = \mathfrak{a}\mathbf{p}x \otimes by @ i \frown \mathfrak{c}\mathbf{p}x \otimes dy @ j$ for some $x, y \in \mathcal{C}$. But in this case too $\mathfrak{a} \in \mathfrak{E}(\nu_0)$, because $\mathfrak{r}_0 \in \text{ax}(\nu_0)$. Of course a similar reasoning applies if we had chosen \mathbf{q} instead of \mathbf{p} , so $\mathfrak{E}(\mu_0) \subseteq \mathfrak{E}(\nu_0)$. Conversely, let $\mathfrak{a} \in \mathfrak{E}(\nu_0)$. Again, if \mathfrak{a} is generated by $\mathfrak{r} \neq \mathfrak{r}_0$, then clearly $\mathfrak{a} \in \mathfrak{E}(\mu_0)$. Otherwise, we have $\mathfrak{a} = ax \otimes by @ i \frown cx \otimes dy @ j$ for some $x, y \in \mathcal{C}$. But x must be of the form $\mathbf{p}x'$ or $\mathbf{q}x'$; in either case, we see that $\text{ax}(\mu_0)$ contains an address generating \mathfrak{a} , so $\mathfrak{E}(\nu_0) \subseteq \mathfrak{E}(\mu_0)$. Point (2) can now be obtained by applying Proposition 3.16: $\mathfrak{E}(\mu) = \mathfrak{E}(\sigma[\mu_0]) = \text{Tr}_\sigma(\mathfrak{E}(\mu_0)) = \text{Tr}_\sigma(\mathfrak{E}(\nu_0)) = \mathfrak{E}(\sigma[\nu_0]) = \mathfrak{E}(\nu)$.

For point (3), we apply again the Decomposition Lemma 1.6 and write $\mu = \sigma[\mu_0]$, $\nu = \sigma[\nu_0]$ with μ_0, ν_0 cut-free and such that $\mu_0 \simeq_{\eta_1}^1 \nu_0$. Observe now that the addresses of leaves are invariant under the η_1 equation of Fig. 4; hence, $\text{ax}(\mu_0) = \text{ax}(\nu_0)$. Again, $\text{ax}^*(\mu_0) = \text{ax}(\mu_0)$ and $\text{ax}^*(\nu_0) = \text{ax}(\nu_0)$, because μ_0, ν_0 are cut-free. So we have $\mathfrak{E}(\mu_0) = \mathfrak{E}(\nu_0)$, and we can conclude once more by applying Proposition 3.16. \square

Proposition 3.19. *For all nets μ, ν , $\mu \simeq_{\beta\eta\epsilon} \nu$ implies $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$.*

Proof. A straightforward consequence of Lemma 3.18, using Lemma 1.17. \square

Proposition 3.20. *Let μ, ν be two nets such that $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$. Then, for every context C , $\mathfrak{E}(C[\mu]) = \mathfrak{E}(C[\nu])$.*

Proof. Observe that applying a context C to net μ with n free ports can be done in two steps: first, we juxtapose C and μ , forming the net which we denote by $C \bullet \mu$. We stipulate that, in $C \bullet \mu$, the free ports of μ are labelled by $1, \dots, n$, whereas the free ports of C are “shifted” by n , *i.e.*, they are labelled starting from $n + 1$. Then, we consider the feedback σ such that $\sigma(i) = i + n$ for $i \in \{1, \dots, n\}$, $\sigma(i) = i - n$ for $i \in \{n + 1, \dots, 2n\}$, and σ is undefined everywhere else. We clearly obtain $\sigma[C \bullet \mu] = C[\mu]$. Note furthermore that $\mathfrak{E}(C \bullet \mu) = \mathfrak{E}(C) \cup \mathfrak{E}(\mu)$, since the two nets are disjoint and do not share free ports by our assumption. The result is then an easy consequence of Corollary 3.17:

$$\begin{aligned} \mathfrak{E}(C[\mu]) &= \mathfrak{E}(\sigma[C \bullet \mu]) = \text{Tr}_\sigma(\mathfrak{E}(C \bullet \mu)) = \text{Tr}_\sigma(\mathfrak{E}(C) \cup \mathfrak{E}(\mu)) = \\ &= \text{Tr}_\sigma(\mathfrak{E}(C) \cup \mathfrak{E}(\nu)) = \text{Tr}_\sigma(\mathfrak{E}(C \bullet \nu)) = \mathfrak{E}(\sigma[C \bullet \nu]) = \mathfrak{E}(C[\nu]). \end{aligned}$$

\square

4. FULL ABSTRACTION

4.1. Edifices and the Cantor topology. Our aim now is to show that edifices are able to fully characterize the observational equivalences introduced in Sect. 2.1. For this, we shall take the sets introduced in Definition 3.1 and equip them with topological structures based on the Cantor topology. This will be needed for two reasons: first, to characterize the edifices which are interpretations of $\beta\epsilon$ -normalizable nets, a result which will be fundamental in characterizing finitary axiom-equivalence; second, to obtain a characterization of axiom-equivalence itself.

The idea of using topology for semantic purposes is of course far from being new: it is enough to think that the very basis of the denotational semantics of the λ -calculus is Scott’s intuition that computability should be interpreted by topological continuity [Sco76]. Moreover, also non-Scott topologies have been attached to λ -terms to obtain various kinds of results (Visser’s topology is an example [Vis80]). Closer to our work, we can mention the work of Kennaway et al. [KKSdV97], who also used a Cantor-like topology, very similar to our own, to define the infinitary λ -calculus.

Definition 4.1 (Arch topology). The set $\mathcal{C} = \{\mathbf{p}, \mathbf{q}\}^{\mathbb{N}}$ may be equipped with the Cantor topology. This is well known to be 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 $\mathfrak{B}(x, r)$ the open ball of center x and radius r .

As well known, $\mathcal{C} \times \mathcal{C}$ is also a Cantor space; if we equip \mathbb{N} with the discrete topology, we can endow the set \mathcal{P} of pillars with the product topology. This is also metrizable: if $\xi = x \otimes y @ i$ and $\nu = x' \otimes y' @ i'$, we shall consider the distance

$$d(\xi, \nu) = \max\{d_{\mathcal{C}}(x, x'), d_{\mathcal{C}}(y, y'), d_{\text{disc}}(i, i')\},$$

where d_{disc} is the discrete metric, defined as $d_{\text{disc}}(i, i') = 0$ if $i = i'$, and $d_{\text{disc}}(i, i') = 2$ if $i \neq i'$. Therefore, to be “close”, two pillars must be based at the same integer.

Similarly, we equip $\overrightarrow{\mathcal{A}}$ with the product topology; the *arch topology*, applied on the set \mathcal{A} of arches, is the quotient topology with respect to the relation \rightleftharpoons of Definition 3.1.

The following helps understanding the arch topology:

Proposition 4.2. *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. \square*

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 a Cantor space, because it is not compact. In fact, we can give a characterization of its compact subsets. Recall from Definition 3.1 that, if $I \subseteq \mathbb{N}$, \mathcal{A}_I is the set of arches based within I . Then, we have

Proposition 4.3. *$\mathfrak{E} \subseteq \mathcal{A}$ 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 all 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 $U_{i,j} = \mathfrak{E} \cap \mathfrak{B}(\mathbf{a}_{i,j}, 2)$. These are all open sets in the relative topology, and since, for all i, j , $D(\mathbf{a}_{i,j}, \mathbf{a}) < 2$ iff \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 4.2): when I is finite, there is identity between closed, compact, and complete subsets of \mathcal{A}_I .

Vaults are examples of compact edifices:

Lemma 4.4. *Vaults are compact.*

Proof. Let $\mathfrak{A} = \mathfrak{A}(\mathfrak{r})$ be a vault, with $\mathfrak{r} = s @ i \frown t @ j$. Clearly $\mathfrak{A}(\mathfrak{r}) \subseteq \mathcal{A}_{\{i,j\}}$. Now take an arch $\mathbf{a} = u @ i' \frown v @ j'$ not belonging to $\mathfrak{A}(\mathfrak{r})$. If $i' \neq i$ or $j' \neq j$, then obviously $\mathfrak{B}(\mathbf{a}, 1)$ is all outside of $\mathfrak{A}(\mathfrak{r})$. Otherwise, either s is not a prefix of u , or t is not a prefix of v ; suppose we are in the first situation, and let k be the length of the longest common prefix between u and s . Then, it is easy to see that $\mathfrak{B}(\mathbf{a}, 2^{-(k+1)})$ is all outside of $\mathfrak{A}(\mathfrak{r})$. So $\mathfrak{A}(\mathfrak{r})$ is a closed subset of $\mathcal{A}_{\{i,j\}}$, and we conclude by Proposition 4.3. \square

Observe that vaults are not open: given an arch $\mathbf{a} \in \mathfrak{A}(s @ i \frown t @ j)$, any non-empty open ball centered at \mathbf{a} contains arches of the form $su @ i \frown tv @ j$, with $u \neq v$.

It turns out that the edifice of a net is compact exactly when the net is $\beta\varepsilon$ -normalizable. This result, which we shall now prove, shows why we are interested in the compact sets of the arch topology. If a is a finite binary word, we denote by $|a|$ its length. If $k \geq |a|$, we denote by $C_k(a)$ the set of words b of length k such that $b = ab'$ for some word b' , *i.e.*, all

possible “extensions” of a to length k . Let now $\mathfrak{r} = a_0 \otimes b_0 @ i \frown c_0 \otimes d_0 @ j$ be an address, and let $k_{\mathfrak{r}} = \max\{|a_0|, |b_0|, |c_0|, |d_0|\}$. We define the set of *centers* of \mathfrak{r} as

$$\begin{aligned} \text{Ctr}(\phi) &= \{ax_0 \otimes bx_0 @ i \frown cx_0 \otimes dx_0 @ j \mid \\ &\quad a \in \mathbf{C}_{k_{\mathfrak{r}}}(a_0), b \in \mathbf{C}_{k_{\mathfrak{r}}}(b_0), c \in \mathbf{C}_{k_{\mathfrak{r}}}(c_0), d \in \mathbf{C}_{k_{\mathfrak{r}}}(d_0)\}, \end{aligned}$$

where x_0 is some fixed infinite word, whose value is irrelevant. Then we set

$$\mathcal{O}(\mathfrak{r}) = \bigcup_{\mathfrak{a} \in \text{Ctr}(\mathfrak{r})} \mathfrak{B}(\mathfrak{a}, 2^{-k_{\mathfrak{r}}+1}).$$

The set $\mathcal{O}(\mathfrak{r})$ is clearly open; additionally, we have

Lemma 4.5. *For every address \mathfrak{r} , $\mathfrak{A}(\mathfrak{r}) \subseteq \mathcal{O}(\mathfrak{r})$.* □

Proof. Let $\mathfrak{r} = a \otimes b @ i \frown c \otimes d @ j$. We assume without loss of generality that a is the longest of a, b, c, d . Then, given $\mathfrak{a} = ax \otimes by @ i \frown cx \otimes dy @ j \in \mathfrak{A}(\mathfrak{r})$, we can always write, for some $a_1, a_2, a_3 \in \mathcal{W}$, $y_1, x_2, y_3 \in \mathcal{C}$, $by = ba_1y_1$, $cx = ca_2x_2$, $dy = da_3y_3$, such that $|ba_1| = |ca_2| = |da_3| = |a| = k_{\mathfrak{r}}$. By definition, $\mathfrak{a}_0 = ax_0 \otimes ba_1x_0 @ i \frown ca_2x_0 \otimes da_3x_0 @ j$ is a center of \mathfrak{r} , and we have $D(\mathfrak{a}_0, \mathfrak{a}) = 2^{-k_{\mathfrak{r}}} < 2^{-k_{\mathfrak{r}}+1}$, so $\mathfrak{a} \in \mathfrak{B}(\mathfrak{a}_0, 2^{-k_{\mathfrak{r}}+1})$. □

A version of part (2) of Lemma 3.12 can be given for the sets $\mathcal{O}(\mathfrak{r})$:

Lemma 4.6. *Let μ be a net, and let $\mathfrak{r}, \mathfrak{r}' \in \text{ax}(\mu)$, with $\mathfrak{r} \neq \mathfrak{r}'$. Then, $\mathcal{O}(\mathfrak{r}) \cap \mathcal{O}(\mathfrak{r}') = \emptyset$.*

Proof. Let $\mathfrak{r} = s @ i \frown t @ j$, $\mathfrak{r}' = s' @ i' \frown t' @ j'$, and let $\mathfrak{r}, \mathfrak{r}' \in \text{ax}(\mu)$. If $i \neq i'$ or $j \neq j'$, then the result is obvious. Otherwise, by the same arguments given in the proof of Lemma 3.12, s and t cannot be prefixes of s' or t' , and vice versa. Now, the sets $\mathcal{O}(\mathfrak{r}), \mathcal{O}(\mathfrak{r}')$ are built precisely so that, whenever $u @ i \frown v @ j \in \mathfrak{A}(\mathfrak{r})$ and $u' @ i' \frown v' @ j' \in \mathfrak{A}(\mathfrak{r}')$, s, t are prefixes of resp. u, v , and s', t' are prefixes of resp. u', v' ; hence, the two sets cannot have any arch in common. □

Proposition 4.7. *For all μ , $\mathfrak{E}(\mu)$ is compact iff μ is $\beta\varepsilon$ -normalizable.*

Proof. The backward implication is a straightforward consequence of Corollary 2.31 and Lemma 4.4 (a finite union of compact sets is compact). Suppose now that μ is not $\beta\varepsilon$ -normalizable. Again thanks to Corollary 2.31, we know that $\text{ax}^*(\mu)$ is infinite. Consider now the family of sets $\mathcal{O}(\mathfrak{r}) \cap \mathfrak{E}(\mu)$ as \mathfrak{r} varies over $\text{ax}^*(\mu)$; by Lemma 4.5, this forms an infinite open cover of $\mathfrak{E}(\mu)$. By Lemma 4.6, removing any of these sets causes the family not to cover $\mathfrak{E}(\mu)$ anymore; hence, $\mathfrak{E}(\mu)$ is not compact. □

4.2. Closed edifices. Since the edifice of a net is always a subset of \mathcal{A}_I with I finite, by Proposition 4.3 we have a standard way to “compactify” it: we simply take its topological closure, denoted by $\overline{(\cdot)}$.

Proposition 4.8. *Let I be a finite subset of \mathbb{N} . Then, for every $\mathfrak{E} \subseteq \mathcal{A}_I$, $\overline{\mathfrak{E}}$ is compact.*

Proof. The arches based outside of I are “too far” to be adherent to \mathfrak{E} , therefore its closure is still in \mathcal{A}_I . By Proposition 4.3, this is enough to ensure compactness. □

Definition 4.9 (Closed edifice of a net, closed trace). Let μ be a net. The *closed edifice* of μ is defined as $\mathfrak{E}^\bullet(\mu) = \overline{\mathfrak{E}(\mu)}$. Similarly, if \mathfrak{E} is an edifice and σ is a feedback, the *closed trace* of \mathfrak{E} along σ is defined as $\text{Tr}_\sigma^\bullet(\mathfrak{E}) = \overline{\text{Tr}_\sigma(\mathfrak{E})}$.

Closed edifices also define a denotational semantics of the symmetric combinators. The fact that they model $\simeq_{\beta\eta\epsilon}$ is an immediate consequence of Proposition 3.19, because $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$ quite obviously implies $\mathfrak{E}^\bullet(\mu) = \mathfrak{E}^\bullet(\nu)$. What is left to prove, is that they yield a congruence, which we do next.

The following result, which is proved by a slightly tricky induction, tells us that if the closure of a uniform edifice \mathfrak{F} contains a uniform edifice \mathfrak{E} , then the trace sequences of \mathfrak{E} along any feedback may be arbitrarily approximated by trace sequences of \mathfrak{F} along the same feedback. The hypothesis that $\mathfrak{E}, \mathfrak{F} \subseteq \mathcal{A}_I$ for I finite is needed so that Proposition 4.8 can be tacitly applied.

Lemma 4.10. *Let $\mathfrak{E}, \mathfrak{F} \subseteq \mathcal{A}_I$ be uniform edifices, with I finite, such that $\mathfrak{E} \subseteq \overline{\mathfrak{F}}$, and let σ be a feedback. Then, for all $\mathfrak{a} \in \text{Tr}_\sigma(\mathfrak{E})$ and for all $\epsilon > 0$, there exists $\mathfrak{b} \in \text{Tr}_\sigma(\mathfrak{F})$ such that $D(\mathfrak{a}, \mathfrak{b}) < \epsilon$.*

Proof. Let $\tilde{\mathfrak{E}} = \{\mathfrak{a}(\mathfrak{s}) \mid \mathfrak{s} \in \text{seq}_\sigma(\mathfrak{E})\}$, and similarly $\tilde{\mathfrak{F}} = \{\mathfrak{a}(\mathfrak{s}) \mid \mathfrak{s} \in \text{seq}_\sigma(\mathfrak{F})\}$. Given $\epsilon > 0$ and a trace sequence $\mathfrak{s} = \mathfrak{s}_1, \dots, \mathfrak{s}_n$ of \mathfrak{E} along σ , we shall prove by induction on n that there exists a vault $\mathfrak{A} \subseteq \tilde{\mathfrak{F}}$ such that $\mathfrak{A} \subseteq \mathfrak{B}(\mathfrak{a}(\mathfrak{s}), \epsilon)$. This will be enough to conclude, because whenever \mathfrak{s} is visible, we have $\mathfrak{A} \subseteq \text{Tr}_\sigma(\mathfrak{F})$, and any arch in \mathfrak{A} satisfies the thesis.

In the base case, $n = 1$, so \mathfrak{s} consists of a single arch $\mathfrak{a} \in \mathfrak{E} \subseteq \tilde{\mathfrak{F}}$. By definition, \mathfrak{a} can be arbitrarily approximated in $\tilde{\mathfrak{F}}$, *i.e.*, given arbitrarily long s, t such that $\mathfrak{a} = su @ i \frown tv @ j$, there exists $\mathfrak{b} \in \tilde{\mathfrak{F}}$ such that $\mathfrak{b} = su' @ i \frown tv' @ j$. It will then be enough to show that $\tilde{\mathfrak{F}}$ contains a vault contained in $\mathfrak{A}(s @ i \frown t @ j)$, because s, t are arbitrarily long, and $\tilde{\mathfrak{F}} \subseteq \tilde{\mathfrak{F}}$. Now, by uniformity of $\tilde{\mathfrak{F}}$, we must have $\mathfrak{b} \in \mathfrak{A}(s_0 @ t \frown t_0 @ j) \subseteq \tilde{\mathfrak{F}}$ for some s_0, t_0 , and $\mathfrak{b} = s_0 w @ i \frown t_0 w @ j$ for some $w \in \mathcal{C} \times \mathcal{C}$. Observe that s_0, s and t_0, t must then be prefixes of each other. Therefore, we have four cases, depending on the possible combinations of which is prefix of which:

- $s = s_0 s'$ and $t = t_0 t'$. Then both s' and t' are prefixes of w , which means that they are prefixes of each other. Suppose that $t' = s' t''$; then, $\mathfrak{b} = s_0 s' w' @ i \frown t_0 s' t'' w' @ j$ for some w' , and $\mathfrak{A}(s @ i \frown t @ j) \subseteq \tilde{\mathfrak{F}}$. The case $s' = t' s''$ is symmetric.
- $s = s_0 s'$ and $t_0 = t t'$. We have $w = s' w'$, and $\mathfrak{b} = s_0 s' w' @ i \frown t t' s' w' @ j$, which means that $\mathfrak{A}(s @ i \frown t t' s' @ j) \subseteq \tilde{\mathfrak{F}}$.
- $s_0 = s s'$ and $t = t_0 t'$. This case is symmetric to the one above.
- $s_0 = s s'$ and $t_0 = t t'$. Then we may conclude, because $\mathfrak{A}(s_0 w @ i \frown t_0 w @ j) \subseteq \mathfrak{A}(s @ i \frown t @ j)$.

Let now $n \geq 1$, and let $\mathfrak{s} = \mathfrak{s}_1, \dots, \mathfrak{s}_n, \mathfrak{s}_{n+1}$. If we put $\mathfrak{s}' = \mathfrak{s}_1, \dots, \mathfrak{s}_n$, by Proposition 3.10, we have that $\mathfrak{a}(\mathfrak{s}')$ is uniform in $\tilde{\mathfrak{E}}$; furthermore, by hypothesis, \mathfrak{s}_{n+1} is uniform in $\tilde{\mathfrak{E}}$. Hence, there exist r_0, s_0, s_1, t_0 and h, i, j, l such that

$$\begin{aligned} \mathfrak{a}(\mathfrak{s}') &\in \mathfrak{A}(r_0 @ h \frown s_0 @ i) \subseteq \tilde{\mathfrak{E}}, \\ \mathfrak{s}_{n+1} &\in \mathfrak{A}(s_1 @ j \frown t_0 @ l) \subseteq \tilde{\mathfrak{E}}, \end{aligned}$$

with $\sigma(i) = j$. Actually, we may always suppose $s_1 = s_0$. In fact, by the match condition, s_0, s_1 are prefixes of each other; suppose for instance that $s_0 = s_1 s'$; then, we have $\mathfrak{a}(\mathfrak{s}') \in \mathfrak{A}(r_0 s' @ h \frown s_1 @ i)$, which is still contained in $\tilde{\mathfrak{E}}$. A symmetric argument applies in case $s_1 = s_0 s'$.

Let us now apply the induction hypothesis, obtaining $\mathfrak{A}(r @ h \frown s @ i) \subseteq \widetilde{\mathfrak{F}}$, with $r = r_0 r'$ and $s = s_0 s'$. The length of r', s' depends on how much we want to approximate $\mathfrak{a}(s')$; in fact, we know that the induction hypothesis allows us to use *any* approximation, as precise as we want. Now, observe that, thanks to the match condition, $s_0 s'$ and $t_0 s'$ are prefixes of the words in \mathfrak{s}_{n+1} ; therefore, because our ultimate goal is to approach $\mathfrak{a}(s)$ within distance ϵ , we apply the induction hypothesis with an ϵ' small enough so that r and $t_0 s'$ are long enough prefixes of the words contained in $\mathfrak{a}(s)$ to satisfy the requirement, *i.e.*, so that, for all $u, v \in \mathcal{C} \times \mathcal{C}$, we have $D(\mathfrak{a}(s), ru @ h \frown t_0 s' v @) < \epsilon$.

Now, observe that $\mathfrak{A}(s_0 @ j \frown t_0 @ l) \subseteq \mathfrak{E} \subseteq \overline{\mathfrak{F}}$ means that all of the arches of this vault can be arbitrarily approximated in \mathfrak{F} ; this applies in particular to the arches of the vault $\mathfrak{A}(s_0 s' @ j \frown t_0 s' @ l)$, which implies, by uniformity of \mathfrak{F} , that there exist s'', t'' such that $\mathfrak{A}_1 = \mathfrak{A}(s_0 s' s'' @ j \frown t_0 s' t'' @ l) \subseteq \mathfrak{F}$. But remark now that we have $\mathfrak{A}_0 = \mathfrak{A}(rs'' @ h \frown ss'' @ i) = \mathfrak{A}(rs'' @ h \frown s_0 s' s'' @ i) \subseteq \widetilde{\mathfrak{F}}$, so, given any $w \in \mathcal{C} \times \mathcal{C}$, a trace sequence \mathfrak{t}' generating the arch $rs'' w @ h \frown s_0 s' s'' w @ i$ of \mathfrak{A}_0 may be extended with the arch $s_0 s' s'' w @ j \frown t_0 s' t'' w @ l$ of \mathfrak{A}_1 , yielding a trace sequence \mathfrak{t} of \mathfrak{F} along σ , such that

$$\mathfrak{a}(\mathfrak{t}) = rs'' w @ h \frown t_0 s' t'' w @ l,$$

which satisfies $D(\mathfrak{a}(s), \mathfrak{a}(\mathfrak{t})) < \epsilon$, because we chose r, s' appropriately when we applied the induction hypothesis. \square

Thanks to Lemma 4.10, we can prove that if two uniform edifices have the same closure, then their closed traces coincide, with respect to any feedback. The congruence property of closed edifices is obtained as an easy corollary, with the help of Corollary 3.17.

Proposition 4.11. *Let $\mathfrak{E}, \mathfrak{F} \subseteq \mathcal{A}_I$ be uniform edifices, with I finite, and let σ be a feedback. Then, $\overline{\mathfrak{E}} = \overline{\mathfrak{F}}$ implies $\text{Tr}_\sigma^\bullet(\mathfrak{E}) = \text{Tr}_\sigma^\bullet(\mathfrak{F})$.*

Proof. By symmetry, and since $\mathfrak{E} \subseteq \overline{\mathfrak{E}}$, it is enough to show that $\mathfrak{E} \subseteq \overline{\mathfrak{F}}$ implies $\text{Tr}_\sigma^\bullet(\mathfrak{E}) \subseteq \text{Tr}_\sigma^\bullet(\mathfrak{F})$. So suppose that \mathfrak{E} is a subset of the closure of \mathfrak{F} , and let $\mathfrak{a} \in \text{Tr}_\sigma^\bullet(\mathfrak{E})$. By definition, there exists a sequence $(\mathfrak{a}_n)_{n \in \mathbb{N}} \in \text{Tr}_\sigma(\mathfrak{E})$ such that $\mathfrak{a}_n \rightarrow \mathfrak{a}$. Let now $\epsilon_m = 2^{-m}$, for $m \in \mathbb{N}$. If we apply Lemma 4.10 to each \mathfrak{a}_n and for each ϵ_m , we obtain a sequence $(\mathfrak{b}_m^n)_{m, n \in \mathbb{N}} \in \text{Tr}_\sigma(\mathfrak{F})$ such that, for all $m, n \in \mathbb{N}$, $D(\mathfrak{a}_n, \mathfrak{b}_m^n) < \epsilon_m$. Consider now the diagonalization of such sequence, *i.e.*, the sequence $(\mathfrak{a}'_n)_{n \in \mathbb{N}} \in \text{Tr}_\sigma(\mathfrak{F})$ defined by setting $\mathfrak{a}'_n = \mathfrak{b}_n^n$, for all $n \in \mathbb{N}$. We contend that $\mathfrak{a}'_n \rightarrow \mathfrak{a}$, which is enough to conclude. So let $\epsilon > 0$. Since $\mathfrak{a}_n \rightarrow \mathfrak{a}$, there exists $N \in \mathbb{N}$ such that, for all $n \geq N$, $D(\mathfrak{a}, \mathfrak{a}_n) < \epsilon/2$. Similarly, let M be smallest integer such that $\epsilon_M < \epsilon/2$, and let $K = \max(M, N)$. We then have, for all $n \geq K$,

$$D(\mathfrak{a}, \mathfrak{a}'_n) \leq D(\mathfrak{a}, \mathfrak{a}_n) + D(\mathfrak{a}_n, \mathfrak{a}'_n) = D(\mathfrak{a}, \mathfrak{a}_n) + D(\mathfrak{a}_n, \mathfrak{b}_n^n) < \epsilon,$$

which proves that \mathfrak{a}'_n tends to \mathfrak{a} as n grows to infinity, as desired. \square

Corollary 4.12. *Let μ, ν be two nets such that $\mathfrak{E}^\bullet(\mu) = \mathfrak{E}^\bullet(\nu)$. Then, for every context C , $\mathfrak{E}^\bullet(C[\mu]) = \mathfrak{E}^\bullet(C[\nu])$.*

Proof. We use the notation of the proof of Proposition 3.20 for juxtaposing nets, *i.e.*, we denote by $C \bullet \mu$ the juxtaposition of C and μ , so that $C[\mu]$ may be written as $\sigma[C \bullet \mu]$ for a suitable feedback σ . Now, observe that the arches of $\mathfrak{E}(C)$ and $\mathfrak{E}(\mu)$ are based at different integers, so that the closure of these two edifices is completely disjoint. The same remark applies to $\mathfrak{E}(C)$ and $\mathfrak{E}(\nu)$; hence, we have $\mathfrak{E}^\bullet(C \bullet \mu) = \mathfrak{E}^\bullet(C) \cup \mathfrak{E}^\bullet(\mu) = \mathfrak{E}^\bullet(C) \cup \mathfrak{E}^\bullet(\nu) = \mathfrak{E}^\bullet(C \bullet \nu)$. Then, using Corollary 3.17 and Proposition 4.11, we may write

$$\mathfrak{E}^\bullet(C[\mu]) = \mathfrak{E}^\bullet(\sigma[C \bullet \mu]) = \text{Tr}_\sigma^\bullet(\mathfrak{E}(C \bullet \mu)) = \text{Tr}_\sigma^\bullet(\mathfrak{E}(C \bullet \nu)) = \mathfrak{E}^\bullet(\sigma[C \bullet \nu]) = \mathfrak{E}^\bullet(C[\nu]).$$

□

4.3. Characterizing observational equivalence. Obtaining a semantic characterization of an observational equivalence involves two results: the proof of a first statement, usually referred to as the *adequacy* of the semantics, establishing that denotational equality implies observational equivalence; and the proof of the converse, *i.e.*, that observational equivalence implies denotational equality, which is usually known as *full abstraction*.

The presence of both results is often simply referred to as a “full abstraction result”, because the second property is in most cases harder to obtain, and is thus the fundamental one. In fact, adequacy is an immediate consequence of the following two properties:

contextuality: the semantics induces a congruence on the syntax, *i.e.*, if two syntactic objects a, b are denotationally equal, then for every context C , $C[a]$ and $C[b]$ are also denotationally equal;

discrimination: the semantics is able to discriminate between the two classes of syntactic objects used to define the observational equivalence, *i.e.*, if the set S is the basis of the equivalence, as described in Sect. 2.1, then for any $a \in S$ and $b \notin S$, one must have that a and b are denotationally different.

To see that adequacy follows from the above two properties, consider the contrapositive statement: let a, b be observationally different, *i.e.*, suppose there exists C such that $C[a] \in S$ and $C[b] \notin S$; by discrimination, we obtain that $C[a]$ and $C[b]$ are denotationally different, so we conclude by contextuality. Note that this latter property is usually taken as a basic property of denotational semantics, *i.e.*, all semantics are assumed to verify it, as described in Sect. 3.3. Hence, all that is left to verify is the discrimination property, which is often not so hard to obtain. This is why, in most cases, full abstraction receives all the attention. Nevertheless, this does not mean that adequacy itself is banal: for instance, our proofs of the contextuality property for edifices and closed edifices (Proposition 3.20 and Corollary 4.12, respectively) are far from being trivial.

So, we start by ensuring that edifices enjoy the discrimination property with respect to finitary axiom-equivalence. For this, we use the topological characterization based on compactness (Proposition 4.7):


Lemma 4.13. *For all nets μ, ν , $\mu \Downarrow$ and $\nu \Uparrow$ implies $\mathfrak{E}(\mu) \neq \mathfrak{E}(\nu)$.*

Proof. Simply observe that, by Proposition 4.7, $\mathfrak{E}(\mu)$ is non-empty and compact, while $\mathfrak{E}(\nu)$ is either empty, or not compact. □


We now have our first full abstraction result:

Theorem 4.14 (Full abstraction for \cong). *For all nets μ, ν , $\mu \cong \nu$ iff $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$.*


Proof. As discussed above, the implication from right to left, or the adequacy property, is a consequence of Proposition 3.20 and Lemma 4.13; so let us examine directly the converse, or rather its contrapositive. Suppose that $\mathfrak{E}(\mu) \neq \mathfrak{E}(\nu)$, and let $\mathfrak{a} \in \mathfrak{E}(\mu) \setminus \mathfrak{E}(\nu)$ (we are supposing without loss of generality that $\mathfrak{E}(\mu)$ is not contained in $\mathfrak{E}(\nu)$). We then have



where the observable axiom shown generates \mathbf{a} , whereas, by Lemma 1.16,



and no reduct of ν' develops a connection between the ports i, j generating \mathbf{a} . Now consider the test



By Lemma 1.9, $\theta[\mu]$ ε -reduces to a quasi-wire (cf. Fig. 14), so $\theta[\mu] \Downarrow$. On the contrary, $\theta[\nu]$ reduces to a net with 2 free ports which cannot be $\beta\eta$ -equivalent to a wire, otherwise we would have $\mathbf{a} \in \mathfrak{E}(\nu)$. We have two possibilities: either $\theta[\nu]$ is $\beta\varepsilon$ -normalizable, or it is not. In the latter case, by Corollary 2.31, we have $\theta[\nu] \Uparrow$, so we are done. In the former case, we take the $\beta\varepsilon$ -normal forms of $\theta[\mu]$ and $\theta[\nu]$, which are cut-free by Proposition 2.29, and conclude by applying the Separation Theorem 1.18. \square

The fact that closed edifices enjoy the discrimination property with respect to axiom-equivalence is trivial, and no special topological property is needed to prove it:

Lemma 4.15. *For all nets μ, ν , $\mu \Downarrow$ and $\nu \Uparrow$ implies $\mathfrak{E}^\bullet(\mu) \neq \mathfrak{E}^\bullet(\nu)$.*

Proof. $\mathfrak{E}(\mu)$ is non-empty (it contains at least one vault), and the closure of a non-empty set is non-empty; on the contrary, $\mathfrak{E}(\nu)$ is empty, and so is its closure. \square

By contrast, in the case of axiom-equivalence, compactness (and hence completeness) becomes essential for yielding a fully-abstract denotational semantics. It is crucial in the proof of the following result:

Lemma 4.16. *Let μ, ν be such that $\mathfrak{E}^\bullet(\mu) \neq \mathfrak{E}^\bullet(\nu)$. Then, one of the following holds:*

- *there exists an observable axiom $\mathfrak{x} \in \text{ax}^*(\mu)$ such that $\mathfrak{A}(\mathfrak{x}) \subseteq \mathfrak{E}(\mu) \setminus \mathfrak{E}^\bullet(\nu)$;*
- *there exists an observable axiom $\mathfrak{y} \in \text{ax}^*(\nu)$ such that $\mathfrak{A}(\mathfrak{y}) \subseteq \mathfrak{E}(\nu) \setminus \mathfrak{E}^\bullet(\mu)$.*


Proof. Let μ have n free ports, and suppose, without loss of generality, that there exists $\mathbf{a} \in \mathfrak{E}^\bullet(\mu) \setminus \mathfrak{E}^\bullet(\nu)$, based at $i, j \in \{1, \dots, n\}$. Remember that $\mathfrak{E}^\bullet(\mu)$ and $\mathfrak{E}^\bullet(\nu)$ are defined as the closures of resp. $\mathfrak{E}(\mu)$ and $\mathfrak{E}(\nu)$, and that by Proposition 4.8 they are both compact, hence complete. Then, if $\mathbf{a} \in \mathfrak{E}^\bullet(\mu) \setminus \mathfrak{E}(\mu)$, \mathbf{a} must be a “missing limit” of a Cauchy sequence $(\mathbf{a}_n)_{n \in \mathbb{N}}$ of $\mathfrak{E}(\mu)$. Since a subsequence of a Cauchy sequence is still a Cauchy sequence, there must exist an integer m such that, for all $n \geq m$, $\mathbf{a}_n \in \mathfrak{E}(\mu) \setminus \mathfrak{E}^\bullet(\nu)$, otherwise \mathbf{a} would belong to $\mathfrak{E}^\bullet(\nu)$ because of its completeness. Therefore, modulo replacing it by one of these \mathbf{a}_n , we can always assume that $\mathbf{a} \in \mathfrak{E}(\mu) \setminus \mathfrak{E}^\bullet(\nu)$. If it is so, then by Definition 3.11 there exists $\mathfrak{x} = s @ i \frown t @ j \in \text{ax}^*(\mu)$ such that $\mathbf{a} \in \mathfrak{A}(\mathfrak{x}) \subseteq \mathfrak{E}(\mu)$, which means that $\mathbf{a} = sw_0 @ i \frown tw_0 @ j$ and, for every $w \in \mathcal{C} \times \mathcal{C}$, $sw @ i \frown tw @ j \in \mathfrak{E}(\mu)$. 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 $\mathbf{a}_n = s_n u_n @ i \frown t_n v_n @ j \in \mathfrak{E}^\bullet(\nu)$; 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 $\mathfrak{E}^\bullet(\nu)$, we would obtain

$\mathfrak{a} \in \mathfrak{E}^\bullet(\nu)$, a contradiction. Therefore, there must exist an integer n such that, for all w , $s_n w @ i \frown t_n w @ j \in \mathfrak{E}(\mu) \setminus \mathfrak{E}^\bullet(\nu)$. \square


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

Lemma 4.17. *Let W be a quasi-wire (Fig. 14), and let μ be a net with two free ports, such that $s @ i \frown t @ j \in \text{ax}^*(\mu)$ implies $i = j$. Then, there exists a test θ such that $\theta[W] \Downarrow$ and $\theta[\mu] \Uparrow$.*

Proof. If $\mu \Uparrow$, the identity test suffices, 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 $\mu \rightarrow_\beta^* \mu'$, and



In the above picture, we have supposed that the observable path connects the free port 1 to itself, and that the leaves connected by the observable axiom 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 1.9, $\theta[W] \rightarrow_\beta^* W$, 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 $\theta'[W] \rightarrow_\beta^* W \Downarrow$, while $\theta'[\mu] \Uparrow$, as desired. \square


We are now ready to prove our second full abstraction theorem:

Theorem 4.18 (Full abstraction for \simeq). *For all nets μ, ν , $\mu \simeq \nu$ iff $\mathfrak{E}^\bullet(\mu) = \mathfrak{E}^\bullet(\nu)$.*

Proof. Once again, the adequacy property, *i.e.*, the backward implication, is a consequence of Corollary 4.12 and Lemma 4.15, so let us turn to the actual full abstraction property. For this, we consider the contrapositive statement, and assume $\mathfrak{E}^\bullet(\mu) \neq \mathfrak{E}^\bullet(\nu)$. Let I be the interface of μ and ν . By Lemma 4.16, we know that there exist $i, j \in I$ and $\mathfrak{x} = s @ i \frown t @ j \in \text{ax}^*(\mu)$ such that, for all w , $sw @ i \frown tw @ j \in \mathfrak{A}(\mathfrak{x}) \setminus \mathfrak{E}^\bullet(\nu)$ (it could actually be that these arches belong to $\mathfrak{A}(\eta) \setminus \mathfrak{E}^\bullet(\mu)$, where $\eta \in \text{ax}^*(\nu)$, 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 $i = j$. Since $\mathfrak{x} \in \text{ax}^*(\mu)$, we have




where we have explicitly drawn the observable axiom of address \mathfrak{r} . On the other hand, by Lemma 1.16, we have




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 the fact that closed edifices model $\simeq_{\beta\eta\varepsilon}$, the edifice of the net on the right is still $\mathfrak{E}^\bullet(\nu)$. 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 , $sw @ i \frown tw @ j \notin \mathfrak{E}^\bullet(\nu)$. 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 1.9, $\theta[\mu]$ β -reduces to a quasi-wire; on the other hand, we have



But ν' never develops observable paths between k and l , so Lemma 4.17 applies, and we obtain $\mu \not\cong \nu$. \square

By inspecting the proofs of Theorems 4.14 and 4.18, we see that only tests are used to discriminate nets. Since those two results say precisely that equality of edifices and closed edifices coincides with finitary axiom-equivalence and axiom-equivalence, respectively, we get the following Context Lemma for free:

Lemma 4.19 (Context). $\mu \cong \nu$ (resp. $\mu \simeq \nu$) iff, for every test θ , $\theta[\mu] \Downarrow$ iff $\theta[\nu] \Downarrow$ (resp. $\theta[\mu] \Downarrow$ iff $\theta[\nu] \Downarrow$). \square




Figure 15: A non- $\beta\varepsilon$ -normalizable net observationally equivalent to a wire.

Furthermore, combined with Proposition 3.19, Theorem 4.14 gives us that $\beta\eta\varepsilon$ -equivalence is included in finitary axiom-equivalence (and hence in axiom-equivalence, by Proposition 2.23); in Sect. 5.3 we shall see that this inclusion is strict (cf. Fig. 19).

Corollary 4.20. *For every nets μ, ν , $\mu \simeq_{\beta\eta\varepsilon} \nu$ implies $\mu \cong \nu$.* \square

On the other hand, as an application of Theorem 4.18 we give an example showing that the inclusion of Proposition 2.23 is strict, *i.e.*, that there exist axiom-equivalent nets which are not finitarily axiom-equivalent. Such example is based on a net which is not $\beta\varepsilon$ -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 reducing as in Fig. 15. Such a net exists by what we have shown in Sect. 1.3; furthermore, after constructing it, one can see that ι is not immediately observable, and that $\mathfrak{r} \in \text{ax}^*(\iota)$ iff $\mathfrak{r} = \mathbf{q}^n \mathbf{p} \otimes \mathbf{1} @ 1 \frown \mathbf{q}^n \mathbf{p} \otimes \mathbf{1} @ 2$ for some non-negative integer n . On the other hand, if ω denotes a wire, we have

$$\mathfrak{E}^\bullet(\omega) = \mathfrak{E}(\omega) = \{u @ 1 \frown u @ 2 ; \forall u \in \mathcal{C} \times \mathcal{C}\}.$$

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 $\mathfrak{E}(\iota)$, hence $\mathfrak{E}(\iota) \subsetneq \mathfrak{E}(\omega)$. But these arches are all adherent to $\mathfrak{E}(\iota)$: in fact, it is very easy to construct a Cauchy sequence in $\mathfrak{E}(\iota)$ of limit \mathbf{a}_y , for any y . Therefore, $\mathfrak{E}^\bullet(\iota) = \mathfrak{E}^\bullet(\omega)$, and $\iota \simeq W$. On the other hand, $\iota \not\cong \omega$, and we do not need Theorem 4.14 to prove that: in fact, the identity is a context discriminating between the two nets.

Note 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.

5. CONCLUDING REMARKS

5.1. Comparison with previous work. The first notion of observational equivalence for interaction nets introduced in the literature is due to Bechet [Bec92]. In his work, the author mentions a notion of behavioral equivalence based on Girard’s coherence spaces [Gir87]. However, we have not been able to reformulate this equivalence so as to compare it to the ones studied in the present paper.

The situation is different with Fernández and Mackie’s work [FM03], the only other existing work on observational equivalence for interaction nets, for which we have precise




Figure 16: Visible ports; α is any symbol, and τ any tree.

results. First of all, we may note that Fernández and Mackie’s approach is more general, *i.e.*, it applies to all systems of interaction nets, not just to the symmetric interaction combinators. However, we have already mentioned that our notion of observable and finitarily observable net can also be generalized to arbitrary systems of interaction nets, as shown in the author’s Ph.D. thesis [Maz06]; the details of this generalization are out of the scope of this paper though.

What is more interesting is to compare our notions of observational equivalence with the specialization of Fernández and Mackie’s observational equivalence to the symmetric interaction combinators, which we shall call here *visible equivalence*. It can be formulated as follows:

Definition 5.1 (Visible port, visible equivalence [FM03]). Let μ be a net with n free ports, and let $1 \leq i \leq n$. We say that i is *immediately visible* in μ iff μ has one of the shapes given in Fig. 16. We say that i is *visible* in μ , and we write $\mu \Downarrow_i$, iff $\mu \rightarrow_{\beta}^* \mu'$ such that i is immediately visible in μ' . We write $\mu \Uparrow_i$ for the negation of $\mu \Downarrow_i$.


Given two nets μ, ν with n free ports, we say that μ and ν are *visibly equivalent*, and we write $\mu \approx \nu$, iff, for every $1 \leq i \leq n$ and for every context C , $C[\mu] \Downarrow_i$ iff $C[\nu] \Downarrow_i$.

Fernández and Mackie [FM03] also give, in case a port i is visible, a notion of *visible agent* at i , and require furthermore that either the visible agents at i of $C[\mu]$ and $C[\nu]$ are the same, or that one of such visible agents is not a constructor. In interaction net systems, a constructor is simply a symbol declared to be such, *i.e.*, it is not an intrinsic notion. Declaring symbols to be constructors may be useful from an “intentional” point view, when one has in mind a particular semantics for the given interaction net system. In the symmetric combinators, because cells with the same symbol may interact, there are arguably no constructors, hence the simplified definition we give here.

Note that, as defined above, visible equivalence does not quite fit in the general pattern of Morris-like observational equivalences discussed in Sect. 2.1, because it is defined “port-wise”, *i.e.*, it takes free ports into account. In other words, visibility is a property of free ports, not of nets. However, Fernández and Mackie’s definition can easily be adjusted so as to conform to Morris’ pattern.

Definition 5.2 (Visible net). Let μ be a net with n free ports. We say that μ is *visible*, and we write $\mu \Downarrow$, iff $\mu \Downarrow_i$ for some $1 \leq i \leq n$.

Lemma 5.3. *Let μ be a net with n free ports, let $1 \leq i \leq n$, and let*



Then, $\nu \Downarrow$ iff $\mu \Downarrow_i$.

Proof. For the backward implication, by definition μ reduces to a net of one of the shapes given in Fig. 16. In the case on the left, the appearance of the α cell is not modified by the ε cells in ν ; in the case on the right, simply observe that the ε cell plugged at the root of τ will “eat” the tree until arriving at free port i . For the forward implication, we have that only the left case of Fig. 16 is possible, *i.e.*, there is a reduct of ν in which the principal port of a cell c appears at its only free port. Note that ε cells only produce ε cells through interaction. Hence, if c is not an ε cell, it already appears in μ ; if c is an ε cell, either it already appears in μ , or c “comes from” one of the ε cells plugged to μ in ν . Then, it is not hard to see that μ must reduce to a net of the shape at the right of Fig. 16, with the ε cell “producing” c being the one plugged to the root of τ . \square


Proposition 5.4. *For all nets μ, ν , $\mu \approx \nu$ iff, for every context C , $C[\mu] \Downarrow$ iff $C[\nu] \Downarrow$.*

Proof. The forward implication is trivial; for what concerns the converse, consider the contrapositive statement: there exist C and $1 \leq i \leq n$ such that, for instance, $C[\mu] \Downarrow_i$ and $C[\nu] \Uparrow_i$. Then, let E be the context plugging ε cells to all free ports of $C[\mu]$ and $C[\nu]$ except i ; by Lemma 5.3, we have $E[C[\mu]] \Downarrow$ and $E[C[\nu]] \Uparrow$, as desired. \square


In their paper [FM03], the authors prove that $\simeq_{\beta\eta} \subseteq \approx$, so, by Proposition 1.19, we have that \approx coincides with $\simeq_{\beta\eta}$ on total nets, just like all the other equivalences introduced in this paper (except of course \simeq_{β}). However, the situation is quite different if we consider non-total nets; indeed, we can show that visible equivalence is strictly stronger than finitary axiom-equivalence (and, by Proposition 2.23, than axiom-equivalence).

Lemma 5.5. *Let μ, ν be two nets with the same interface, and let C be a context such that $C[\mu]$ reduces to a quasi-wire, while $C[\nu]$ reduces to a net whose one of the two free ports is connected to a principal port. Then, $\mu \not\approx \nu$.*

Proof. Simply consider the context



where α is any binary symbol, and we have supposed, without loss of generality, that the free port of the reduct of $C[\nu]$ which is connected to the principal port is the one on the left in the above picture (while the free ports drawn at the top of the picture are those that are connected to μ and ν in $C[\mu]$ and $C[\nu]$, respectively). Then, we have, for some net without interface o , some symbol β , and some net ν_0 ,



so that $C'[\mu] \uparrow$, whereas $C'[\nu] \downarrow$. \square


Lemma 5.6. *Let $\mu \downarrow$ and $\nu \uparrow$. Then, $\mu \not\approx \nu$.*

Proof. By Proposition 2.22, there exists a test θ such that $\theta[\mu]$ reduced to a quasi-wire. On the other hand, by Lemma 2.21 we still have $\theta[\nu] \uparrow$. Observe that, if i is any of the two free ports of $\theta[\mu]$, we have $\theta[\mu] \downarrow_i$. Now, if one of the two free ports of $\theta[\nu]$ is not visible, we are done. We may then assume $\theta[\nu] \rightarrow^* \nu'$, where both free ports of ν' are immediately visible; then, they must both be connected to a principal port, otherwise ν' would be immediately observable (cf. Fig. 16, right), whereas we know it to be blind. We may therefore conclude by applying Lemma 5.5. \square


The fact that visible equivalence is stronger than finitary axiom-equivalence is a trivial corollary of the following:

Proposition 5.7. *Let $\mu \downarrow$ and $\nu \uparrow$. Then, $\mu \not\approx \nu$.*

Proof. Let μ and ν have n free ports. We have two possibilities: either $\nu \uparrow$, or $\text{ax}^*(\nu)$ is infinite. In the first case, observe that $\mu \downarrow$ implies $\mu \downarrow$, so we conclude by Lemma 5.6. In the second case, since $\text{ax}^*(\mu)$ is finite, there must exist $1 \leq i, j \leq n$ and pairs of words s, t such that $s @ i \frown t @ j \in \text{ax}^*(\nu) \setminus \text{ax}^*(\mu)$. Now, s and t describe two trees τ_s, τ_t and a leaf in each of them, such that



where the wire shown connects the two leaves of τ_s, τ_t described by s, t , respectively. Consider then the test



where the only leaves of τ_s, τ_t left free are again those described by s, t , respectively. We obviously have that $\theta[\nu]$ β -reduces to a quasi-wire; on the contrary, because of the way we have chosen s and t , $\theta[\mu]$ is not $\beta\eta$ -equivalent to a wire. Observe however that $\text{ax}^*(\theta[\mu])$ is still finite (although it may now be empty). Then, we may consider the $\beta\varepsilon$ -normal forms of $\theta[\nu]$ and $\theta[\mu]$, which exist by Corollary 2.31, and apply the Separation Theorem 1.18 to them. We thus obtain a further test θ' such that $\theta'[\theta[\mu]]$ β -reduces to a quasi-wire and $\theta'[\theta[\nu]] \uparrow$, or vice versa. In any case, we reason as in the proof of Lemma 5.6: if one of the ports of the blind net is not visible, we conclude; otherwise, we apply Lemma 5.5. \square



Figure 17: Two nets showing that visible equivalence is more discriminative than finitary axiom-equivalence; α is any binary symbol.




Figure 18: The parallelizing context for nets with n free ports; ν is an arbitrary net.

Corollary 5.8. *For all nets μ, ν , $\mu \approx \nu$ implies $\mu \cong \nu$.*

Proof. Consider the contrapositive statement: $\mu \not\cong \nu$ implies that there exists C such that, for example, $C[\mu] \Downarrow$ and $C[\nu] \Uparrow$; by Proposition 5.7, we have $C[\mu] \not\approx C[\nu]$, so we conclude $\mu \not\approx \nu$ by using the fact that \approx is a congruence. \square

To see that visible equivalence is *strictly* stronger than finitary equivalence, consider the nets of Fig. 17: the only free port of $\mu = \mathbf{E}_1$ is visible, while the only free port of ν is not visible, so $\mu \not\approx \nu$; on the contrary, the edifice of both nets is empty, so by Theorem 4.14 we have $\mu \cong \nu$.

An intuitive justification to Corollary 5.8 and to the example of Fig. 17 is that the difference between a visible and an observable net is seemingly akin to the difference between a head-normalizable and a *weak*-head-normalizable λ -term. In fact, the two cases of Fig. 16 are strikingly similar to the cases $\lambda x.M$ (left) and $xM_1 \dots M_n$ (right) defining weak head normal forms: the latter case is a special case of observable net, just like $xM_1 \dots M_n$ is a special case of head-normal-form; the former case guarantees that a net visible on port i is “reactive” when we plug the principal port of a cell to i itself, *i.e.*, an active pair is created, just like $\lambda x.M$ is “reactive” to application (a redex is created). In the λ -calculus, whnf-equivalence is strictly stronger than nf- and hnf-equivalence [DCG01]; this is in accord with our intuition about visible equivalence and (finitary) axiom-equivalence.

5.2. Approximations and the Genericity Lemma. In the λ -calculus, unsolvable terms are important because they represent *meaningless data*. One of the main formal arguments in favor of this intuition is the so-called Genericity Lemma [Bar84]: let M be an unsolvable λ -term, and let C be such that $C[M]$ is normalizable; then, $C[X] \simeq_\beta C[M]$ for every λ -term X . In other words, if we see $C[\cdot]$ as a function, the only functions which are able to produce something meaningful (a normal form) out of unsolvable terms are the constant functions, confirming the fact that unsolvable terms are meaningless.

In the symmetric interaction combinators, a word-by-word rephrasing of the Genericity Lemma fails; this is because of two interesting differences with respect to the λ -calculus:

- the intrinsic parallelism of interaction nets, which has no equivalent in the λ -calculus;
- the fact that, for every $n \in \mathbb{N}$, there is an unsolvable net with n free ports which is cut-free, hence normal (namely, the net we called \mathbf{E}_n); by contrast, no normal λ -term can be unsolvable.

Indeed, consider the context P_ν given in Fig. 18. We may refer to this context as the “parallelizing” context: in fact, for every net μ with n free ports, $P_\nu[\mu] = \mu \bullet \nu$, *i.e.*, the juxtaposition of μ and ν . Now, let ν be a net which is normalizable, or total, or finitarily observable, and take the unsolvable net \mathbf{E}_n . Clearly $P_\nu[\mathbf{E}_n]$ is normalizable, or total, or finitarily observable, but there always exists a net ξ such that $P_\nu[\xi]$ need not be normalizable, or total, or finitarily observable, let alone β -equivalent to $P_\nu[\mathbf{E}_n]$. This is because ξ and ν do not interact, so the properties of $P_\nu[\xi]$ basically depend solely on ξ .

However, there is a reformulation of the Genericity Lemma which holds for the symmetric interaction combinators, and which supports the fact that our notion of unsolvable net coincides indeed with that of meaningless data. Take an unsolvable net μ , and take a context C . Then, we can prove that, whenever

$$C[\mu] = \begin{array}{c} \boxed{\mu} \\ \cdots \\ \boxed{C} \\ \cdots \end{array} \xrightarrow{*}_\beta \begin{array}{c} \boxed{\mu_0} \\ \cdots \\ \boxed{C_0} \\ \cdots \end{array}$$

such that C_0 is cut-free and none of the wires connecting C_0 to μ_0 is a cut, we have that, for every net ξ with the same number of free ports as μ , there exists ξ_0 such that

$$C[\xi] = \begin{array}{c} \boxed{\xi} \\ \cdots \\ \boxed{C} \\ \cdots \end{array} \xrightarrow{*}_\beta \begin{array}{c} \boxed{\xi_0} \\ \cdots \\ \boxed{C_0} \\ \cdots \end{array}$$

In other words, every bit of information in the result of the computation represented by $C[\mu]$ is also present in $C[\xi]$, for all ξ , which means that μ actually does not produce any information, and is thus meaningless.

The above concept of “bit of information” may be formalized by the notion of *approximation*:

Definition 5.9 (Approximation). Let μ be a net. An *approximation* of μ is a cut-free net ν such that:

- $\nu = C[\mathbf{E}_n]$ for some context C and $n \in \mathbb{N}$;
- $\mu \xrightarrow{*}_\beta C[\mu_0]$ for some μ_0 with n free ports.

If ν is an approximation of μ , we write $\nu \sqsubseteq \mu$.

Intuitively, an approximation of μ is a “piece” of the hypothetical cut-free form of μ , *i.e.*, it gives a partial information on the result of the computation represented by μ . The least information, or the lack thereof, is \mathbf{E}_n , which is an approximation of every net with n free ports. If a net is total, then its cut-free form is also an approximation of it, the most complete one indeed.

Approximations and edifices are related by the following, whose proof is left to the reader:

Proposition 5.10. *Let μ be a net, and let ν be a cut-free net with the same interface as μ . Then, $\nu \sqsubseteq \mu$ iff $\mathfrak{E}(\nu) \subseteq \mathfrak{E}(\mu)$. \square*

The above result confirms in particular the idea that the edifice of a net may be seen as its “infinite cut-free form”. Indeed, the relation \sqsubseteq can be made a partial order, and the set of approximations of a net can be shown to be a directed set. However, the order given by \sqsubseteq is not complete, so this set has no least upper bound in general; to make it complete, one should introduce infinite cut-free nets, which is more or less what edifices are.

We may then state the Genericity Lemma as follows:

Lemma 5.11 (Genericity). *Let μ be an unsolvable net with n free ports. Then, for every context C and for every net ξ with n free ports, $\nu \sqsubseteq C[\mu]$ implies $\nu \sqsubseteq C[\xi]$.*

Proof. By Proposition 2.22, we have $\mu \uparrow$, so $\mathfrak{E}(\mu) = \emptyset$. Hence, we can write, using Corollary 3.17, Lemma 3.7, and the decomposition $C[\mu] = \sigma[C \bullet \mu]$ for a suitable feedback σ ,

$$\mathfrak{E}(C[\mu]) = \text{Tr}_\sigma(\mathfrak{E}(C) \cup \mathfrak{E}(\mu)) = \text{Tr}_\sigma(\mathfrak{E}(C)) \subseteq \text{Tr}_\sigma(\mathfrak{E}(C) \cup \mathfrak{E}(\xi)) = \mathfrak{E}(C[\xi]).$$

Now, by Proposition 5.10, $\nu \sqsubseteq C[\mu]$ implies $\mathfrak{E}(\nu) \subseteq \mathfrak{E}(C[\mu])$, so $\mathfrak{E}(\nu) \subseteq \mathfrak{E}(C[\xi])$, and we conclude $\nu \sqsubseteq C[\xi]$ again by Proposition 5.10. \square

5.3. Theories for the symmetric interaction combinators. In the foundational studies concerning the λ -calculus, an important role is played by λ -theories [Bar84, LS04]. These can be straight-forwardly be reformulated in the context of the symmetric interaction combinators:

Definition 5.12 (Theory). A *theory* is a binary relation \sim on nets such that:

- (1) \sim relates nets with the same interface;
- (2) \sim is a congruence;
- (3) $\simeq_\beta \subseteq \sim$.

The set of theories \mathfrak{T} is a complete bounded lattice with respect to inclusion: given any family of theories $(\sim_i)_{i \in I}$, the least upper bound (lub) is defined by $(\bigcup_{i \in I} \sim_i)^+$, and the greatest lower bound (glb) by $\bigcap_{i \in I} \sim_i$; the least element is \simeq_β , and the greatest element is the inconsistent theory \top , which equates all nets with the same interface.

Much effort has been put forth in order to understand the structure of the lattice of λ -theories; quite a few things are known about it [Vis80, Bar84, LS04], and many more are the subject of ongoing research [BMS07, BS08, BMS09, CS09]. In the case of the symmetric combinators, we suspect the structure of \mathfrak{T} to be at least as intricate as in the case of the λ -calculus. In this section, we gather everything we presently know about it (which is arguably not much!), leaving several questions open for further work.

As in the case of the λ -calculus, we may define *sensible* and *semi-sensible* theories, based on the fact that unsolvable nets are meaningless, and it is therefore sensible to identify all of them:

Definition 5.13 (Sensible and semi-sensible theory). A theory \sim is *sensible* iff, for all μ, ν unsolvable, $\mu \sim \nu$. A theory \sim is *semi-sensible* iff $\mu \sim \nu$ implies μ solvable iff ν solvable.

Note that any theory containing a sensible theory is sensible, while any theory contained in a semi-sensible theory is semi-sensible (these are both immediate consequences of the definition). The two notions are related as follows.

Lemma 5.14. *Let \sim be a sensible theory. Then:*

- (1) *for every blind net μ with n free ports, $\mu \sim \mathbf{E}_n$;*

(2) for every quasi-wire W , $W \sim \omega$, where ω is a wire.

Proof. Point (1) is obvious (modulo Proposition 2.22). For point (2), observe that $W = \omega[o]$, where o is some net without interface, necessarily blind; moreover, note that $\omega = \omega[\mathbf{E}_0]$, where \mathbf{E}_0 is the empty net. Now, by point (1), $o \sim \mathbf{E}_0$; but then we can conclude, because \sim is a congruence. \square

Proposition 5.15. *A consistent sensible theory is semi-sensible.*

Proof. Let \sim be a sensible theory, and let $\mu \Downarrow$ and $\nu \Uparrow$. We shall prove that $\mu \sim \nu$ implies $\sim = \top$. First of all, by Proposition 2.22, by the fact that \sim includes β -equivalence, and by point (2) of Lemma 5.14, there exists a test θ such that $\theta[\mu] \sim \omega$. On the other hand, by Lemma 2.21, and by point (1) of Lemma 5.14, we have $\theta[\nu] \sim \mathbf{E}_2$. But \sim is a congruence, so $\omega \sim \mathbf{E}_2$, and we may conclude by Proposition 1.19. \square

Apart from \simeq_β , in the course of this paper we introduced several theories: $\simeq_{\beta\eta}$, $\simeq_{\beta\varepsilon}$, $\simeq_{\beta\eta\varepsilon}$, \cong , \simeq , and \approx (the first and the last were actually introduced by Fernández and Mackie, cf. Sect. 5.1). All of them are semi-sensible, because they are all included in \simeq , which is semi-sensible by definition. Furthermore, since $\simeq_{\beta\varepsilon}$ is sensible by definition (ε -reduction equates precisely all unsolvable nets), all the theories including it are also sensible, namely $\simeq_{\beta\eta\varepsilon}$, \cong , and \simeq . On the contrary, \simeq_β , $\simeq_{\beta\eta}$, and \approx are not sensible: Fig. 17 gives an example of two unsolvable nets which are distinguished by all of these theories.

Indeed, Fernández and Mackie's equivalence is an example of non-sensible theory which strictly extends $\beta\eta$ -equivalence, and stands quite on its own with respect to the other theories discussed in this paper. For instance, it is completely orthogonal to $\beta\varepsilon$ -equivalence: this latter is not included in \approx , as shown again by the example of Fig. 17; and \approx is not included in $\simeq_{\beta\varepsilon}$, because the former includes η -equivalence. On the other hand, although the example of Fig. 17 tells us that $\simeq_{\beta\eta\varepsilon}$ is not included in \approx , we know nothing about the converse. All we know is that \approx is strictly contained in \cong (Corollary 5.8 and Fig. 17).

What about consistent sensible theories in general? First of all, observe that the lub and glb of a family of sensible theories is sensible, so the set of sensible theories is a complete sub-lattice of \mathfrak{T} , which is actually bounded. The least element is obviously $\simeq_{\beta\varepsilon}$, because it is defined so as to validate exactly β -equivalence plus equality of every unsolvable net. The greatest element turns out to be \simeq ; in fact, this can be shown to be a coatom of \mathfrak{T} , i.e., a maximal consistent theory (so \simeq is also the greatest semi-sensible theory).

Proposition 5.16. *Let \sim be a theory such that $\simeq \subsetneq \sim$. Then, $\sim = \top$.*

Proof. We start by observing that \sim is sensible, because it includes \simeq . Now, let $\mu \sim \nu$, with $\mu \not\sim \nu$. We then have a context C such that, for example, $C[\mu] \Downarrow$ and $C[\nu] \Uparrow$. But \sim is a congruence, so $C[\mu] \sim C[\nu]$, which proves that \sim is not semi-sensible. By Proposition 5.15, the only sensible theory which is not semi-sensible is \top . \square

So far, the situation is identical to the case of the λ -calculus, in which the lattice of consistent sensible λ -theories has least element $\simeq_{\beta\Omega}$ (also known as \mathcal{H}) and greatest element hnf-equivalence (also known as \mathcal{H}^*).

In between the two, there is $\simeq_{\beta\eta\Omega}$, which coincides with nf-equivalence (also denoted by $\mathcal{H}\eta$). Here we find the first sharp difference with respect to the λ -calculus: $\simeq_{\beta\eta\varepsilon}$, which is analogous to $\simeq_{\beta\eta\Omega}$, does not coincide with \cong , which, morally, is the counterpart of nf-equivalence. In fact, the converse of Corollary 4.20 fails: the two nets of Fig. 19 give an interesting example of this. They can be built by slightly twisting the constructions given




Figure 19: Nets showing that $\mu \cong \nu$ does not imply $\mu \simeq_{\beta\eta\epsilon} \nu$.

in Sect. 1.3. It is not hard to show that $\mu_1 \not\approx_{\beta\eta\epsilon} \mu_2$; in some sense, the two nets endlessly “chase” each other in their reduction, never managing to meet. And yet, it is evident that they generate exactly the same observable axioms, *i.e.*, $\text{ax}^*(\mu_1) = \text{ax}^*(\mu_2)$. Therefore, $\mathfrak{E}(\mu_1) = \mathfrak{E}(\mu_2)$, and $\mu_1 \cong \mu_2$ by Theorem 4.14. Note how the parallelism of interaction nets, absent in the λ -calculus, plays once again a crucial role in this example. If we use the analogy that observable axioms are head variables, here we are clearly exploiting the fact that interaction nets allow several head variables in parallel: although μ_1 and μ_2 have the same “head variables”, they “produce” them in a different order.

In the λ -calculus, it is possible to show that between $\mathcal{H}\eta$ and \mathcal{H}^* there is a continuum of sensible theories [Bar84]; we ignore whether this is the case for the symmetric interaction combinators. Indeed, a related open question is the existence of *easy* nets, *i.e.*, nets which can be consistently equated with any other net with the same interface. A first difference with the λ -calculus is that \mathbf{E}_2 , which is the prototypical unsolvable net with 2 free ports, is not easy (Proposition 1.19 shows that there is no consistent theory equating it with a wire); on the contrary, the λ -term Ω , which is the prototypical unsolvable term, can be shown to be easy.

Finally, we give an example of a consistent non-semi-sensible theory. Define *total equivalence* as $\mu \sim_t \nu$ iff, for every context C , $C[\mu]$ is total iff $C[\nu]$ is total. It can be shown that η -equivalence does not alter totality, so $\simeq_{\beta\eta} \subseteq \sim_t$. Moreover, observe that total equivalence is the only theory we introduced which distinguishes between the empty net \mathbf{E}_0 (which is total) from all other nets with no interface not reducing to \mathbf{E}_0 (which are not total). This proves in particular that the theory is consistent.

However, consider the net ξ obtained by juxtaposing two copies of the net ν given in Fig. 17, and let ι be the net of Fig. 15. Both nets have 2 free ports, and are thus comparable; furthermore, none of the two nets is total (ι is not normalizable, ξ contains vicious circles), and we clearly have $\iota \Downarrow$ and $\xi \Uparrow$. Now, non-totally has a quite singular behavior if compared to non-normalizability in the λ -calculus, in that it is “resistant” to contexts: in fact, by the locality of interaction rules, neither active pairs nor vicious circles can be eliminated through interaction; thus, if μ is not total, so is $C[\mu]$, for any C . Therefore, we have $\iota \sim_t \xi$, which proves that total equivalence is not semi-sensible. So the only possible relationship with the other known theories is $\approx \not\subseteq \sim_t$; we ignore whether this is the case.

5.4. More open questions and further work. In Fig. 20 we graphically resume what we know about theories in the symmetric interaction combinators. A solid line represents inclusion, from bottom to top; a thick solid line represents atomic inclusion, *i.e.*, there is no theory in between. The main open questions concerning Fig. 20 discussed up to now may be resumed as follows:

The situation may be different for $\simeq_{\beta\varepsilon}$ (or $\simeq_{\beta\eta\varepsilon}$). Fig. 19 shows that there exists nets such that $\mu_1 \not\sim_{\beta\varepsilon} \mu_2$ and yet $\text{ax}^*(\mu_1) = \text{ax}^*(\mu_2)$, so any semantics based on simply collecting observable axioms will not work. What would be needed is an additional structure to $\text{ax}^*(\mu)$, which takes into account the causal relationship between observable axioms. For instance, we may think of endowing $\text{ax}^*(\mu)$ with a poset structure, in the style of Winskel’s event structures [Win82]: given $\mathfrak{r}, \mathfrak{h} \in \text{ax}^*(\mu)$, $\mathfrak{r} \leq \mathfrak{h}$ iff $\mu \xrightarrow{\beta}^* \mu'$ and $\mathfrak{h} \in \text{ax}(\mu')$ imply $\mathfrak{r} \in \text{ax}(\mu')$. For example, consider the nets μ_1, μ_2 of Fig. 19. The observable axioms generated by these two nets fall within one of two categories: those based at free port 1, whose addresses we denote by $\mathfrak{r}_1^1, \mathfrak{r}_2^1, \dots$, and those based at free port 2, whose addresses we denote by $\mathfrak{r}_1^2, \mathfrak{r}_2^2, \dots$. Then, the structure of $\text{ax}^*(\mu_1)$ as a poset would be $\mathfrak{r}_1^1 < \mathfrak{r}_1^2 < \mathfrak{r}_2^1 < \mathfrak{r}_2^2 < \dots$, while the structure of $\text{ax}^*(\mu_2)$ would be $\mathfrak{r}_1^2 < \mathfrak{r}_1^1 < \mathfrak{r}_2^2 < \mathfrak{r}_2^1 < \dots$, which is enough to tell the two nets apart. We have not yet attempted to formalize these ideas, but we believe them to be a promising direction of research to obtain a full abstraction result for $\simeq_{\beta\varepsilon}$ (or $\simeq_{\beta\eta\varepsilon}$, which is perhaps more feasible—again, in the λ -calculus, no non-syntactic model is known whose theory is \mathcal{H} , which corresponds to our $\simeq_{\beta\varepsilon}$).

Of course, there is also the question of semantically characterizing Fernández and Mackie’s \approx . We currently have no clue about this question, but if, as we discussed in the end of Sect. 5.1, this equivalence is akin to whnf-equivalence in the λ -calculus, then we may be facing a difficult problem: in the λ -calculus, no full abstraction result exists at present for such equivalence.

Concerning edifices, an aspect which should further be explored is their connection with games semantics and traced monoidal categories. Indeed, the trace operation on edifices is strikingly reminiscent of the notion of “composition and hiding” for composing strategies in games semantics, trace sequences representing plays. One may wonder whether the set of edifices presented in this paper can be seen as some sort of “reflexive object” in a traced monoidal category of edifices. This would be quite interesting, because it would open the way for introducing a typed version of the symmetric interaction combinators: types would be modeled by the objects of this category, and the set of edifices introduced here would appear as a special type capable of modeling untyped nets (in the context of the λ -calculus, this would be like an object D in a Cartesian closed category such that $D \Rightarrow D$ is a retract of D). Such considerations also bring forth the question of what is a categorical model of the symmetric interaction combinators, a question for which we have some clues, but which is still unsettled.

Finally, there is the intriguing possibility of using edifices as the basis for defining new non-deterministic extensions of the symmetric interaction combinators, or modeling existing non-deterministic systems. What we have in mind is something in the vein of Ehrhard and Regnier’s differential interaction nets [ER06]. In fact, as soon as one considers non-simple edifices (*cf.* Definition 3.4), several non-deterministic phenomena emerge: arches may superpose, *i.e.*, they may share a pillar, which is reminiscent of additive slices in proof nets, or in differential interaction nets; and trace sequences, which represent computational paths in nets, are no longer uniquely determined by the arch they generate (Lemma 3.9 fails).

REFERENCES

- [AJM00] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Information and Computation*, 163(2):409–470, 2000.

- [Bar84] Henk P. Barendregt. *The Lambda Calculus*. North Holland, revised edition, 1984.
- [Bec92] Denis Bechet. Partial evaluation of interaction nets. In *Proceedings of WSA '92*, pages 331–338, 1992.
- [BMS07] Chantal Berline, Giulio Manzonetto, and Antonino Salibra. Lambda theories of effective lambda models. In Jacques Duparc and Thomas A. Henzinger, editors, *Proceedings of CSL 2007*, volume 4646 of *Lecture Notes in Computer Science*, pages 268–282. Springer, 2007.
- [BMS09] Chantal Berline, Giulio Manzonetto, and Antonino Salibra. Effective lambda-models vs recursively enumerable lambda-theories. *Mathematical Structures in Computer Science* (to appear), 2009.
- [Böh68] Corrado Böhm. Alcune proprietà delle forme $\beta\eta$ -normali nel λ -K-calcolo. *Pubblazioni dell'IAC*, 696:1–19, 1968.
- [BS08] Antonio Bucciarelli and Antonino Salibra. Graph lambda theories. *Mathematical Structures in Computer Science*, 18(5):975–1004, 2008.
- [CS09] Alberto Carraro and Antonino Salibra. Reflexive domains are not complete for the extensional lambda calculus. In *Proceedings of LICS 2009* (to appear), 2009.
- [DCG01] Mariangiola Dezani-Ciancaglini and Elio Giovannetti. From Bohm's theorem to observational equivalences: an informal account. *Electronic Notes in Theoretical Computer Science*, 50(2):85–118, 2001.
- [DR95] Vincent Danos and Laurent Regnier. Proof nets and the Hilbert space. In J.-Y. Girard, Y. Lafont, and L. Regnier, editors, *Advances in Linear Logic*, pages 307–328. Cambridge University Press, 1995.
- [ER06] Thomas Ehrhard and Laurent Regnier. Differential interaction nets. *Theoretical Computer Science*, 364(2):166–195, 2006.
- [FM01] Maribel Fernández and Ian Mackie. Packing interaction nets: Applications to linear logic and the lambda calculus. In *Argentinian Workshop of Theoretical Computer Science (WAIT2001)*, 2001.
- [FM03] Maribel Fernández and Ian Mackie. Operational equivalence for interaction nets. *Theoretical Computer Science*, 297(1–3):157–181, 2003.
- [GAL92] Georges Gonthier, Martín Abadi, and Jean-Jacques Lévy. The geometry of optimal lambda reduction. In *Conference Record of POPL 92*, pages 15–26. ACM Press, 1992.
- [Gir87] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–102, 1987.
- [Gir89] Jean-Yves Girard. Geometry of Interaction I: interpretation of System F. In *Proceedings of the Logic Colloquium '88*, pages 221–260. North Holland, 1989.
- [Gir91] Jean-Yves Girard. A new constructive logic: Classical logic. *Mathematical Structures in Computer Science*, 1(3):255–296, 1991.
- [Gir96] Jean-Yves Girard. Proof-nets: The parallel syntax for proof-theory. In Ursini and Agliano, editors, *Logic and Algebra*. Marcel Dekker, Inc., 1996.
- [HO00] Martin Hyland and Luke Ong. On full abstraction for PCF: I, II, and III. *Information and Computation*, 163(2):285–408, 2000.
- [HS04] Esfandiar Haghverdi and Philip J. Scott. A categorical model for the geometry of interaction. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Proceedings of the 31st International Colloquium on Automata, Languages and Programming (ICALP 2004)*, volume 3142 of *Lecture Notes in Computer Science*, pages 708–720. Springer, 2004.
- [Hyl76] Martin Hyland. A syntactic characterization of the equality in some models of the lambda calculus. *J. London Math. Society*, 2(12):361–370, 1976.
- [KKSdV97] Richard Kennaway, Jan Willem Klop, Ronan Sleep, and Fer-Jan de Vries. Infinitary lambda calculus. *Theoretical Computer Science*, 137(1):69–101, 1997.
- [Laf90] Yves Lafont. Interaction nets. In *Conference Record of POPL'90*, pages 95–108. ACM Press, 1990.
- [Laf95] Yves Lafont. From proof nets to interaction nets. In J.-Y. Girard, Y. Lafont, and L. Regnier, editors, *Advances in Linear Logic*, pages 225–247. Cambridge University Press, 1995.
- [Laf97] Yves Lafont. Interaction combinators. *Information and Computation*, 137(1):69–101, 1997.
- [Lon05] Giuseppe Longo. Savoir critique et savoir positif: l'importance des résultats négatifs. *Intellectica*, 40(1), 2005.

- [LS04] Stefania Lusin and Antonino Salibra. The lattice of lambda theories. *Journal of Logic and Computation*, 14(3):373–394, 2004.
- [Mac04] Ian Mackie. Efficient lambda evaluation with interaction nets. In *Proceedings of RTA '04*, LNCS, pages 155–169. Springer, 2004.
- [Mac05] Ian Mackie. An interaction net implementation of additive and multiplicative structures. *Journal of Logic and Computation*, 15(2):219–237, 2005.
- [Maz06] Damiano Mazza. *Interaction Nets: Semantics and Concurrent Extensions*. Ph.D. Thesis, Université de la Méditerranée/Università degli Studi Roma Tre, 2006.
- [Maz07a] Damiano Mazza. A denotational semantics for the symmetric interaction combinators. *Mathematical Structures in Computer Science*, 17(3):527–562, 2007.
- [Maz07b] Damiano Mazza. Observational equivalence for the interaction combinators and internal separation. *Electronic Notes in Theoretical Computer Science*, 176(1):113–137, 2007.
- [Mor68] J. H. Morris. *Lambda calculus models of programming languages*. Ph.D. Thesis, Massachusetts Institute of Technology, 1968.
- [MP02] Ian Mackie and Jorge Sousa Pinto. Encoding linear logic with interaction combinators. *Information and Computation*, 176(2):153–186, 2002.
- [Nak75] Reiji Nakajima. Infinite normal forms for the λ -calculus. In Corrado Böhm, editor, *Lambda-Calculus and Computer Science Theory*, LNCS, pages 62–82. Springer, 1975.
- [Pag07] Michele Pagani. Proofs, denotational semantics and observational equivalences in multiplicative linear logic. *Mathematical Structures in Computer Science*, 17(2):341–361, 2007.
- [Plo72] Gordon Plotkin. A set-theoretical definition of application. Technical Report MIP-R-95, School of A.I., Univ. of Edinburgh, 1972.
- [Sco76] Dana Scott. Data types as lattices. *SIAM Journal on Computing*, 5(3):522–587, 1976.
- [SS71] Dana Scott and Christopher Strachey. Toward a mathematical semantics for computer languages. In J. Fox, editor, *Proc. Symp. Computers and Automata*. Polytechnic Inst. of Brooklyn Press, 1971. Also Technical Monograph PRG-6, Programming Research Group, Oxford University.
- [Vis80] A. Visser. Numerations, lambda calculus, and arithmetic. In Hindley and Seldin, editors, *To H. B. Curry: Essays on Combinatory Logic, Lambda-Calculus, and Formalism*, pages 259–284. Academic Press, 1980.
- [Wad76] Christopher Wadsworth. The relation between computational and denotational properties for Scott’s D_∞ models. *Siam J. Comput.*, 5(3):488–521, 1976.
- [Win82] Glynn Winskel. Event structure semantics for CCS and related languages. In Mogens Nielsen and Erik Meineche Schmidt, editors, *Proceedings of ICALP '82*, volume 140 of *Lecture Notes in Computer Science*, pages 561–576. Springer-Verlag, 1982.


APPENDIX A. INVARIANCE OF THE TRACE

This appendix is devoted to the proof of Proposition 3.13, which we recall below:

Proposition 3.13. *Let $\mu \rightarrow_\beta \mu'$, and let $\mu = \sigma[\nu]$ and $\mu' = \sigma'[\nu']$ according to the Decomposition Lemma 1.6. Then, $\text{Tr}_\sigma(\mathfrak{E}(\nu)) = \text{Tr}_{\sigma'}(\mathfrak{E}(\nu'))$.*


The result is basically a corollary of the following:

Lemma A.1. *Let ν and σ be resp. the cut-free net and feedback drawn below:*




Then:

(1) if $\alpha = \beta$, and if ν' and σ' are the following cut-free net and feedback



then $\text{Tr}_{\sigma'}(\nu') = \text{Tr}_{\sigma}(\nu)$;

(2) if $\alpha \neq \beta$, and if ν' and σ' are the following cut-free net and feedback



then $\text{Tr}_{\sigma'}(\nu') = \text{Tr}_{\sigma}(\nu)$;

Proof. The proof is a bit technical, and not very interesting. We shall only address point (2), point (1) being analogous and technically simpler. We assume $\alpha = \delta$ and $\beta = \zeta$, the other configuration being perfectly symmetric, and start by showing the inclusion $\text{Tr}_{\sigma}(\nu) \subseteq \text{Tr}_{\sigma'}(\nu')$. Let $\mathfrak{s} = (u_k @ i_k \frown v_k @ j_k)_{1 \leq k \leq n}$ be a trace sequence of $\mathfrak{C}(\nu)$ along σ , such that $i_1 \notin \text{dom } \sigma$. We shall build by induction on n a trace sequence $\mathfrak{s}' = (u'_k @ i'_k \frown v'_k @ j'_k)_{1 \leq k \leq n'}$ of $\mathfrak{C}(\nu')$ along σ' such that:

- $u'_1 @ i'_1 = u_1 @ i_1$;
- $v'_{n'} @ j'_{n'} = v_n @ j_n$ in case $j_n \notin \text{dom } \sigma$, *i.e.*, in case \mathfrak{s} is visible;
- otherwise, j_n is one of the free ports denoted by p_1, p_2 in the above picture of $\sigma[\nu]$, and v_n is of the form $t_h w$, where $1 \leq h \leq 4$ and $t_1 = \mathbf{p} \otimes \mathbf{1}$, $t_2 = \mathbf{q} \otimes \mathbf{1}$, $t_3 = \mathbf{1} \otimes \mathbf{p}$, $t_4 = \mathbf{1} \otimes \mathbf{q}$, the value of h depending on the cell and auxiliary port “used” by the observable axiom/path of ν inducing \mathfrak{s}_n . In that case, \mathfrak{s}' will be such that $v'_{n'} @ j'_{n'} = w @ q_h$, where q_h is one of the free ports of ν' as shown in the picture of $\sigma'[\nu']$, point (2).

It is obvious that the above is sufficient to prove the inclusion, because a visible trace sequences \mathfrak{s} yield a visible trace sequence \mathfrak{s}' such that $\mathbf{a}(\mathfrak{s}') = \mathbf{a}(\mathfrak{s})$.

The base case is $n = 1$, in which \mathfrak{s} consists of a single arch $u @ i \frown v @ j$. If $j \notin \text{dom } \sigma$, then j is a free port of $\sigma[\nu]$, and \mathfrak{s} is also a visible trace sequence of $\mathfrak{C}(\nu')$ along σ' , so we take $\mathfrak{s}' = \mathfrak{s}$. Otherwise, $j \in \{p_1, p_2\}$, and $v = t_h v'$, with $1 \leq h \leq 4$ and t_h as described above. In this case, the sequence \mathfrak{s}' is defined to consist of the sole arch $u @ i \frown x \otimes y @ q_h$; this is clearly in $\mathfrak{C}(\nu')$, and \mathfrak{s}' satisfies the desired requirements.

Let now $n > 1$. We write $\mathfrak{s}_{n-1} = \xi \frown v_{n-1} @ j_{n-1}$, and observe that $j_{n-1} \in \text{dom } \sigma$, because of the chain condition. Then, we have $j_{n-1} \in \{p_1, p_2\}$ and $v_n = t_h w$ for some $w \in \mathcal{C} \times \mathcal{C}$ and $1 \leq h \leq 4$, with t_h is as described above. We shall assume $h = 1$; the other three cases are perfectly similar. So we have $j_{n-1} = p_1$, and the chain condition forces $i_n = \sigma(j_{n-1}) = p_2$, so \mathfrak{s}_n is of the form $su @ p_2 \frown v_n @ j_n$ for some $u \in \mathcal{C} \times \mathcal{C}$ and $s \in \{\mathbf{1} \otimes \mathbf{p}, \mathbf{1} \otimes \mathbf{q}\}$. We make the choice $s = \mathbf{1} \otimes \mathbf{p}$, the other cases being again analogous.

So, to resume, we know that the last two arches of \mathfrak{s} are of the form

$$\begin{aligned}\mathfrak{s}_{n-1} &= \xi \frown \mathbf{p}x \otimes y @ p_1, \\ \mathfrak{s}_n &= x' \otimes \mathbf{p}y' @ p_2 \frown v_n @ j_n,\end{aligned}$$

for some $x, y, x', y' \in \mathcal{C}$ and $\xi \in \mathcal{P}$; observe that the match condition implies $x' = \mathbf{p}x$ and $y = \mathbf{p}y'$.

Now, by the induction hypothesis applied to $\mathfrak{s}_1, \dots, \mathfrak{s}_{n-1}$, we know how to build a sequence $\mathfrak{s}' = (u'_k @ i'_k \frown v'_k @ j'_k)_{1 \leq k \leq n'}$ such that $u'_1 @ i'_1 = u_1 @ i_1$ and such that the last arch is of the form

$$\mathfrak{s}'_{n'} = \xi' \frown x \otimes y @ q_1$$

Remark that $\sigma'(q_1) = q_8$, $\sigma'(q_6) = q_3$, and that

$$x'' \otimes \mathbf{p}y'' @ q_8 \frown \mathbf{p}x'' \otimes y'' @ q_6 \in \mathfrak{E}(\nu')$$

for all $x'', y'' \in \mathcal{C}$. We now have two possibilities:

- $j_n \notin \text{dom } \sigma$, *i.e.*, j_n is a free port of $\sigma[\nu]$ and $\sigma'[\nu']$. In this case, $\mathfrak{s}_n \in \mathfrak{E}(\nu)$ implies $x' \otimes y' @ q_3 \frown v_n @ j_n \in \mathfrak{E}(\nu')$;
- $j_n \in \text{dom } \sigma$, which implies $j_n \in \{p_1, p_2\}$. We assume $j_n = p_1$, again the case $j_n = p_2$ being perfectly similar. Then, we have $v_n = \mathbf{p}z \otimes z'$ for some $z, z' \in \mathcal{C}$, which implies $x' \otimes y' @ q_3 \frown z \otimes z' @ p_1 \in \mathfrak{E}(\nu')$.

Then, define

$$\mathfrak{s}'_{n'+1} = x \otimes \mathbf{p}y' @ q_8 \frown \mathbf{p}x \otimes y' @ q_6,$$

and

$$\mathfrak{s}'_{n'+2} = x' \otimes y' @ q_3 \frown v_n @ j_n$$

in case $j_n \notin \text{dom } \sigma$, and

$$\mathfrak{s}'_{n'+2} = x' \otimes y' @ q_3 \frown z \otimes z' @ p_1$$

in case $j_n \in \text{dom } \sigma$. In both cases, by the arguments given above we have $\mathfrak{s}'_{n'+1}, \mathfrak{s}'_{n'+2} \in \mathfrak{E}(\nu')$, and $(\mathfrak{s}'_k)_{1 \leq k \leq n'+2}$ is a trace sequence of $\mathfrak{E}(\nu')$ along σ' satisfying the desired requirements.

We are left with proving that $\text{Tr}_{\sigma'}(\mathfrak{E}(\nu')) \subseteq \text{Tr}_{\sigma}(\mathfrak{E}(\nu))$. We use a similar argument, but this time we build a trace sequence \mathfrak{s} of $\mathfrak{E}(\nu)$ along σ only starting from a trace sequence $\mathfrak{s}' = (u'_k @ i'_k \frown v'_k @ j'_k)_{1 \leq k \leq n'}$ of $\mathfrak{E}(\nu')$ along σ' such that $j'_{n'} \notin \{q_5, q_6, q_7, q_8\}$ (the induction is on n'). This will be enough for the inclusion to be proved, because visible trace sequences of $\mathfrak{E}(\nu')$ along σ' do not end with any of those free ports of ν' , as they are not free in $\sigma'[\nu']$. The sequence $\mathfrak{s} = (u_k @ i_k \frown v_k @ j_k)_{1 \leq k \leq n}$ will have to satisfy the following:

- $u_1 @ i_1 = u'_1 @ i'_1$;
- $v_n @ j_n = v'_{n'} @ j'_{n'}$ if $j'_{n'} \notin \text{dom } \sigma'$, *i.e.*, in case \mathfrak{s}' is visible;
- otherwise, we must have $j'_{n'} \in \{q_1, q_2, q_3, q_4\}$; then, \mathfrak{s} will satisfy $v_n @ j_n = sv'_{n'} @ p_1$ with $s = \mathbf{p} \otimes \mathbf{1}$ (resp. $s = \mathbf{q} \otimes \mathbf{1}$) if $j'_{n'} = q_1$ (resp. $j'_{n'} = q_2$), or $v_n @ j_n = sv'_{n'} @ p_2$ with $s = \mathbf{1} \otimes \mathbf{p}$ (resp. $s = \mathbf{1} \otimes \mathbf{q}$) if $j'_{n'} = q_3$ (resp. $j'_{n'} = q_4$).

The base case is $n' = 1$, in which \mathfrak{s}' consists of exactly one arch $u' @ i' \frown v' @ j'$, and obviously $j' \notin \{q_5, q_6, q_7, q_8\}$, because $i' \notin \text{dom } \sigma'$. If $j' \notin \text{dom } \sigma'$, we take $\mathfrak{s} = \mathfrak{s}'$. Otherwise, supposing $j' = q_1$, we take \mathfrak{s} to be made of the sole arch $u' @ i' \frown (\mathbf{p} \otimes \mathbf{1})v' @ p_1$; the other three possible values of j' are handled similarly, prefixing v' with $\mathbf{q} \otimes \mathbf{1}, \mathbf{1} \otimes \mathbf{p}, \mathbf{1} \otimes \mathbf{q}$ as appropriate.

Let now $n' > 1$. We put $j = j'_{n'}$, and observe that, by $j \notin \{q_5, q_6, q_7, q_8\}$ and by the chain condition, we have $i'_{n'} \in \{q_1, q_2, q_3, q_4\}$. We have again four cases; as above, each

time we shall have a choice in the sequel, we shall analyze only one arbitrary case, all cases being easily recoverable every time from each other. So we assume, for instance, $i'_{n'} = q_1$. We have $j'_{n'-1} = \sigma'(q_1) = q_8$, so $\mathfrak{s}'_{n'-1}$ is based at q, q_8 , where $q \in \{q_5, q_6\}$; in both cases, q is not a free port of $\sigma'[\nu']$, so the sequence must contain a previous arch of the form $\mathfrak{s}'_{n'-2} = \xi' \frown v' @ q'$, with $q' \in \{q_1, q_2, q_3, q_4\}$. We choose $q = q_6$ and $q' = q_3$, so we can write, by using the match condition, that the last three arches of \mathfrak{s}' are of the form

$$\begin{aligned}\mathfrak{s}'_{n'-2} &= \xi \frown x \otimes y @ q_3, \\ \mathfrak{s}'_{n'-1} &= \mathbf{p}x' \otimes y' @ q_6 \frown x' \otimes \mathbf{p}y' @ q_8, \\ \mathfrak{s}'_{n'} &= x'' \otimes y'' @ q_1 \frown v @ j,\end{aligned}$$

for some $x, y, x', y', x'', y'' \in \mathcal{C}$, $v \in \mathcal{C} \times \mathcal{C}$, and $\xi \in \mathcal{P}$. Furthermore, by the match condition, we know that $x = \mathbf{p}x'$, $y = y'$, $x'' = x'$, and $y'' = \mathbf{p}y'$. If we apply the induction hypothesis to the sequence $\mathfrak{s}'_1, \dots, \mathfrak{s}'_{n'-2}$, we obtain a sequence $\mathfrak{s} = (u_1 @ i_1 \frown v_n @ j_n)_{1 \leq k \leq n}$ satisfying all the requirements mentioned above; in particular, we have

$$\mathfrak{s}_n = \xi' \frown x \otimes \mathbf{p}y @ p_2.$$

Now, suppose $j \in \text{dom } \sigma'$; we choose for example $j = q_1$. In that case, define

$$\mathfrak{s}_{n+1} = \mathbf{p}x'' \otimes y'' @ p_1 \frown (\mathbf{q} \otimes \mathbf{1})v @ p_1.$$

Otherwise, j is a free port of $\sigma'[\nu']$; then, we set

$$\mathfrak{s}_{n+1} = \mathbf{p}x'' \otimes y'' @ p_1 \frown v @ j.$$

In both cases, it is easy to see that $\mathfrak{s}_{n+1} \in \mathfrak{E}(\nu)$, and that $\mathfrak{s}_1, \dots, \mathfrak{s}_n, \mathfrak{s}_{n+1}$ is a trace sequence of $\mathfrak{E}(\nu)$ along σ , which is visible iff \mathfrak{s}' is. \square

To prove Proposition 3.13, observe first of all that trace sequences never use ε cells, so the only interesting interaction rules are those addressed by Lemma A.1. Then, we can always write $\sigma[\nu] = \sigma_1[\sigma_0[\nu]]$, where $\sigma_0[\nu]$ is of the form given in Lemma A.1; similarly, we can write $\sigma'[\nu'] = \sigma_1[\sigma'_0[\nu']]$, where $\sigma'_0[\nu']$ is of one of the forms given in point (1) or (2) of Lemma A.1, depending on whether the interaction rule is an annihilation or commutation. Note that σ_1, σ_0 and σ_1, σ'_0 are disjoint feedbacks. Then, by Lemma A.1 and the associativity of the trace (Lemma 3.8), we have

$$\text{Tr}_\sigma(\mathfrak{E}(\nu)) = \text{Tr}_{\sigma_1}(\text{Tr}_{\sigma_0}(\mathfrak{E}(\nu))) = \text{Tr}_{\sigma_1}(\text{Tr}_{\sigma'_0}(\mathfrak{E}(\nu'))) = \text{Tr}_{\sigma'}(\mathfrak{E}(\nu')).$$