EDIFICES AND FULL ABSTRACTION FOR THE SYMMETRIC INTERACTION COMBINATORS

DAMIANO MAZZA

Preuves, Programmes et Systèmes, Université Paris 7
e-mail address: Damiano.Mazza@pps.jussieu.fr

Abstract. The symmetric interaction combinators are a variant of Lafont's interaction combinators. They are a graph-rewriting model of parallel deterministic computation. We define 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 metric spaces, called (pre-)edifices, which play roughly the same role as Böhm trees in the theory of the lambda-calculus. However, edifices have a richer structure: they are able to give a nice topological account of phenomena like infinite eta-expansion, and there is a notion of trace on them, which is reminiscent of the notion of "play" in games semantics, or of the execution formula of the geometry of interaction.

Introduction

Lafont's interaction nets [Laf90] are a powerful and versatile model of parallel deterministic computation, derived from the proof-nets of Girard's linear logic [Gir96, Laf95]. Interaction nets are characterized by the atomicity and locality of their rewriting rules. They can be seen as "parallel Turing machines": computational steps are elementary enough to be considered as constant-time operations, but several steps can be executed at the same time.

Several interesting applications of interaction nets exist. The most notable ones are implementations of optimal evaluators for the λ-calculus [GAL92, 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.

2000 ACM Subject Classification: Don't know.
Key words and phrases: Interaction nets, denotational semantics, full abstraction, Böhm trees.
Supported by a post-doctoral fellowship of the Fondation Sciences Mathématiques de Paris.

LOGICAL METHODS
IN COMPUTER SCIENCE
DOI:10.2168/LMCS-???

This work aims precisely at studying and expanding the theory of interaction nets, in particular of the symmetric interaction combinators. These latter are particularly interesting because of their *universality*: any interaction net system can be translated in the symmetric interaction combinators [Laf97]. Therefore, a semantic study of the symmetric combinators applies, modulo a translation, to all interaction net systems.

The main contribution of this paper is the introduction of *observable paths*, a notion inspired by the Geometry of Interaction [Gir89, DR95]. The central idea is to see a net as a collection of observable paths, each being like one thread of a parallel computation. Using observable paths, we develop the notion of *observable net* and *totally observable net*. Intuitively, the first is a net in which at least one parallel thread of the computation terminates; the second is a net in which *all* parallel threads terminate.

It is useful to keep in mind an analogy with the λ-calculus: observable nets are similar to solvable λ-terms, *i.e.*, terms having a head normal form, and totally observable nets are akin to normalizable λ-terms. The correspondence is somewhat looser in the second case, because the symmetric combinators already have a notion of "being normalizable", and it does not coincide with being "totally observable". However, there are several phenomena supporting this analogy.

In fact, the notions of observable and totally observable net can be used to define two context-based equivalences on nets: *observational equivalence*, denoted by $\simeq$, and *total equivalence*, denoted by $\cong$. The first one, based on observable nets, 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, based on totally observable nets, 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 total equivalence is strictly included in observational 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, as is the case for the λ-calculus. It should be noted in passing that these two equivalences are both different from the one introduced by Fernández and Mackie [FM03]: the latter is in fact based on *interface normal forms*, which seem to be related to λ-calculus *weak* head normal forms.

In the λ-calculus, hnf- and nf-equivalence were semantically characterized in the early '70s by the independent results of Wadsworth and Hyland [Wad76, Hyl76]: two λ-terms are hnf-equivalent iff their Böhm tree has the same *infinite* η-normal form, and two λ-terms are nf-equivalent iff their Böhm tree has the same η-normal form. Shortly after, Nakajima introduced a similar characterization of hnf-equivalence in terms of what are now known as Nakajima trees [Nak75].

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 with respect to observational equivalence and total equivalence. Edifices are compact (hence complete) metric spaces, related to Cantor spaces. When nets are interpreted as edifices, phenomena similar to infinite η-expansion, which, as mentioned above, are also present in the symmetric combinators, receive a precise topological explanation.

Apart from characterizing the interactive behavior of nets, edifices show other interesting aspects, not developed in this paper. They have many common features with the *strategies* of game semantics, and are related to the Geometry of Interaction interpretation of nets [DR95, Laf97]. In particular, one can define a notion of *trace* for edifices, which may be used to compose them as strategies are composed in games semantics, or as operators are evaluated according to the execution formula in the Geometry of Interaction. This feature is a novelties with respect to Böhm trees, and it may be of help in improving the theory of interaction nets. For example, one may think of defining a category out of edifices, which may serving as the base for a typed semantics of the symmetric combinators. Also, edifices may suggest additive (or non-deterministic) extensions of interaction nets. They may also turn out to be useful in defining alternative models of other systems, like proof-nets, or the λ-calculus itself, as these can all be encoded in the symmetric combinators.

1. The Symmetric Interaction Combinators

1.1. **Nets and reduction.** The symmetric interaction combinators, or, more simply, the symmetric combinators, are an interaction net system [Laf90, Laf97]. An interaction net is the union of two structures: a labelled, directed hypergraph, and an undirected multigraph.

We start with a formal definition of interaction net, adapted to the special case of the symmetric combinators. After that, we give a more informal, graphical definition, which will be the one actually used in the paper (as in most of the existing literature).

In what follows, we fix a denumerably infinite set of *ports*, which we assume contains the positive integers. We denote by $[a_1, \ldots, a_n]$ the multiset containing the elements $a_1, \ldots, a_n$, which may not be pairwise distinct.

Definition 1.1 (Wire, cell, net)**.** A *wire* is a set of exactly two ports. We fix three *symbols* $\delta, \varepsilon, \zeta$; we say that δ and ζ are binary, while ε is nullary. A *cell* is a tuple $(\alpha, p_0, p_1, \ldots, p_n)$ where α is a symbol, $p_0, p_1, \ldots, p_n$ 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 $p_1, \ldots, p_n$ are the *auxiliary ports.*

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

- if a port appears in $\mathsf{Cells}(\mu)$ or $\mathsf{Wires}(\mu)$, then it must appear in at least one wire;
- each port appears at most twice in $\mathsf{Cells}(\mu)+\mathsf{Wires}(\mu)$, counting multiplicities (in this union, $\mathsf{Cells}(\mu)$ is considered as a multiset in which all elements have multiplicity 1).

The set of ports appearing in μ is denoted by $\mathsf{Ports}(\mu)$. A port appearing only once in $\mathsf{Cells}(\mu) + \mathsf{Wires}(\mu)$ 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 $\{1, \ldots, n\}$.

Definition 1.2 (Structural equivalence)**.** A *renaming* for a net with n free ports is an injective function from ports to ports which is the identity on $\{1, \ldots, n\}$. Two nets are *α-equivalent* iff they are equal modulo a renaming. Two nets μ, μ' are *ω_0-equivalent* iff $\mathsf{Wires}(\mu) = W + [\{p, q\}, \{q, r\}]$ and $\mathsf{Wires}(\mu') = W + [\{p, r\}]$, with $p \neq r$. *ω-equivalence* is the reflexive-transitive closure of ω_0-equivalence. *Structural equivalence* is the transitive closure of the union of α- and ω-equivalence.

It is natural to always consider nets up to α-equivalence. Sometimes it may be of interest to consider ω-equivalence as part of reduction;[1] however, we do not do so in this paper, and we shall always work modulo structural equivalence.

[1]Basically, interaction nets are generalized multiplicative proof nets with implicit axiom and cut links [Laf95]. The acquainted reader will then see that, once properly oriented into a rewriting step, is easy to see that ω_0-equivalence can be oriented into a rewriting step, which corresponds to the axiom step of cut-elimination in proof nets.

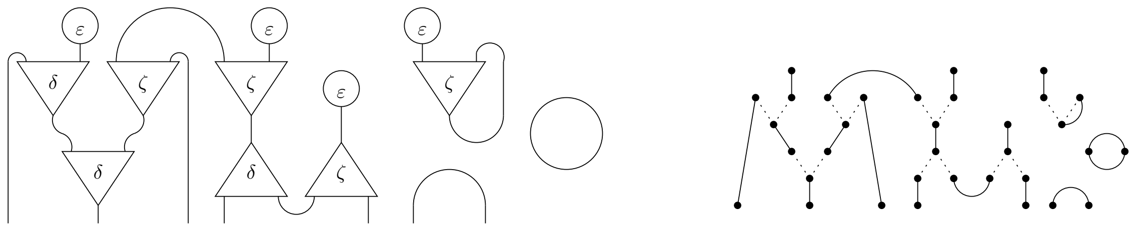

Figure 1: A net (left) and its port graph (right, internal edges dotted).

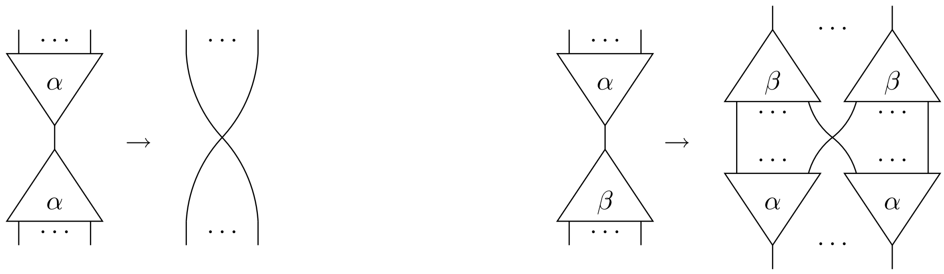

Figure 2: The interaction rules: annihilation (left) and commutation (right). In the annihilation, the right member is empty in case $\alpha = \varepsilon$.

Most of the time it is convenient to present a net graphically, as in Fig. 1. In these representations, only cells and wires are drawn, and ports are left implicit. For a binary cell, the principal port is represented by one of the "tips" of the triangle representing it. A wire is represented as… a wire, and the free ports appear as extremities of "pending" wires. For example, the net in Fig. 1 has 7 free ports. Note that graphical representations equate exactly structurally equivalent nets.

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

Reducing an active pair inside a net means removing it and replacing it with the net given by the corresponding rule. If a net μ is transformed into μ' after such an operation, we write $\mu \to \mu'$. We define $\mu \simeq_\beta \nu$ iff there exists o such that $\mu \to^* o$ and $\nu \to^* o$.

Proposition 1.3 (Strong confluence)**.** *If $\mu \to \mu'$ and $\mu \to \mu''$ with $\mu' \neq \mu''$, there exists ν such that $\mu' \to \nu$ and $\mu'' \to \nu$. Hence, the relation $\to^*$ is confluent, and $\simeq_\beta$ is an equivalence relation.* □

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

A *tree* is a net defined by induction as follows. A single ε cell is a tree with no leaf, denoted by ε; a single wire is a tree with one leaf (it is arbitrary which of the two extremities is the root and which is the leaf), denoted by $\bullet$; if τ_1, τ_2 are two trees with resp. n_1, n_2 leaves, and if α is a binary symbol, the net

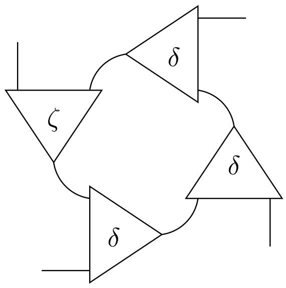

Figure 3: A vicious circle.

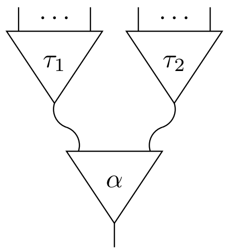

is a tree with $n_1 + n_2$ leaves, denoted by $\alpha(\tau_1, \tau_2)$.

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

Lemma 1.4. *Let τ be a tree. Then, we have*

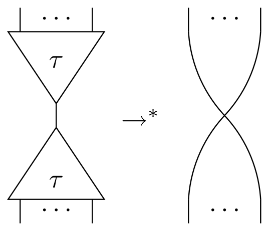

Proof. By induction on τ. ☐

There is also a generalization of the commutation rule:

Lemma 1.5. *Let Δ be a tree not containing ζ cells, and let Z be a tree not containing δ cells. Then, we have*

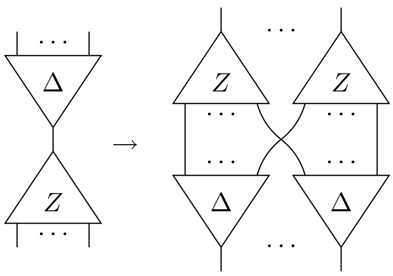

Proof. By induction. ☐

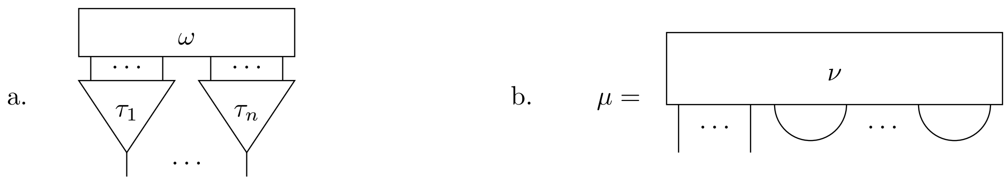

Figure 4: a. General form of a cut-free net; b. Decomposition of a generic net μ.

A *vicious circle* is either a cyclic wire, or a configuration consisting of n binary cells $c_1, \ldots, c_n$ such that, for all $i \in \{1, \ldots, n-1\}$, the principal port of c_i is connected to an auxiliary port of c_{i+1}, and the principal port of c_n is connected to an auxiliary port of c_1. Such configurations are stable under reduction, because cells can interact only through their principal port: they are sort of deadlocks. The loop at the far right of the net in Fig. 1 is the simplest example of vicious circle; another example is given in Fig. 3.

A net is *cut-free* if it contains no active pair and no vicious circle. A *total* net is a net admitting a cut-free form. In contrast to the λ-calculus, 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. Total nets are important because the Separation Theorem 1.11 applies to them, and because cut-free nets are the "true" normal forms: a non-total net represents either a diverging or an error-bound computation, *i.e.*, one that generates deadlocks. Other useful notions of convergence will be introduced in Sect. 3, but none of them will coincide with simple normalization.

One can prove by a simple induction that every cut-free net uniquely decomposes in terms of trees and wirings as in Fig. 4a. Hence, any net μ with n free ports and k active pairs and/or vicious circles can be written as in Fig. 4b, where ν is a cut-free net with $n+2k$ free ports. The net ν is unique as soon as μ does not contain vicious circles.

The following is an easy corollary of Lemmas 1.4 and 1.5:

Lemma 1.6 (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*

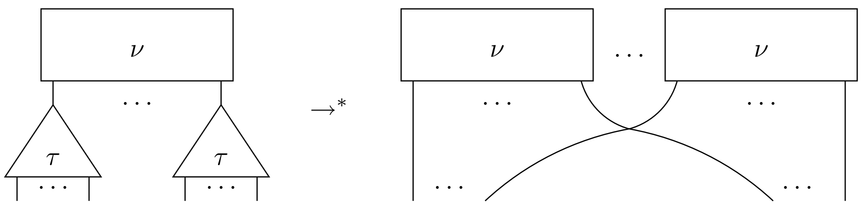

Any cut-free net can be freely erased, as shown again by an easy induction:

Lemma 1.7 (Erasing)**.** *For any cut-free net ν, we have*

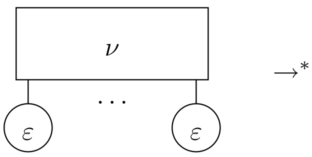

1.3. **Expressiveness.** The interest of the symmetric combinators is given by the following result:

Theorem 1.8 (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 combinators[2] is given by Mackie and Pinto [MP02]. We refer the reader to Lafont's paper [Laf97] for a proper formulation and proof of Theorem 1.8.

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 equations of the form

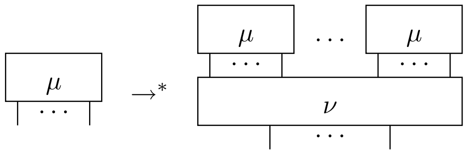

may be solved. In the λ-calculus, the above equation would correspond to

$$M \to^* N[M/x]$$

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, $\Theta T \to^* T(\Theta T)$. Then, a solution to the above equation 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 1.6, so we do not have a fixpoint combinator at our disposal. To compensate for this, we use a fundamental construction due to Lafont [Laf97].

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. The net $\S\mu$ is called the *cut-free code* of μ. Recovering μ from its cut-free code can be done as in Fig. 8. Rigorously speaking, $\S\mu$ is not well defined because, given a net μ, the cut-free net ν is not unique in general. However, the reader can check that 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 $!\S\mu$ as the *code* of μ. Decoding is done by composing the nets **D** and **R** of Fig. 6 and 8, respectively; we denote by **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 1.6

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

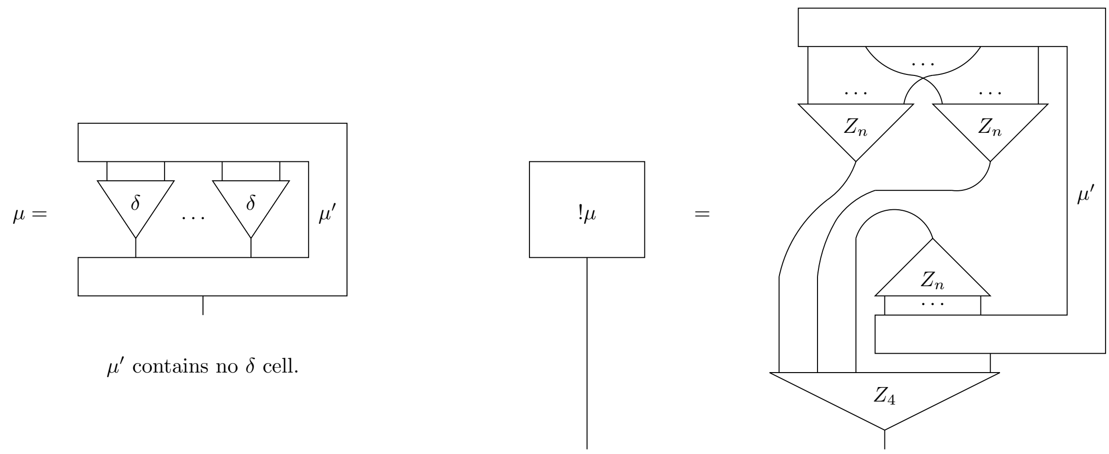

Figure 5: The Lafont code of a net.

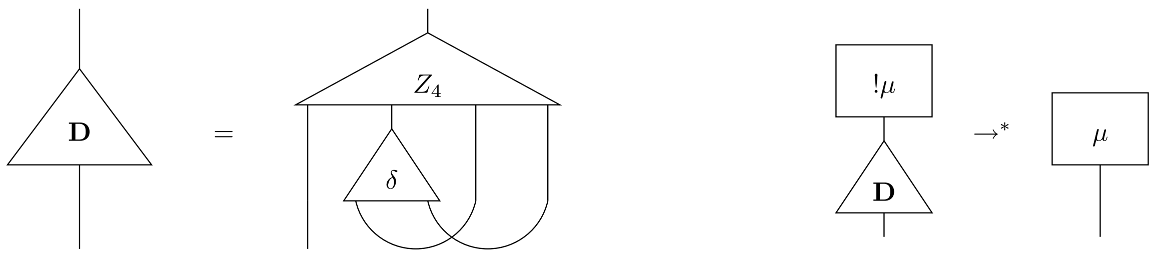

Figure 6: The universal decoder for the Lafont code.

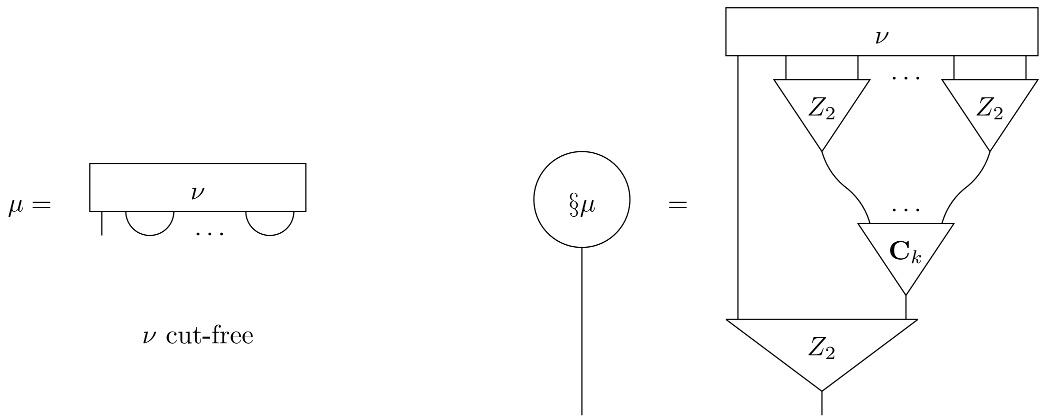

Figure 7: The cut-free code of a net.

applies to it, and it can be freely duplicated by means of trees of δ cells, *i.e.*, the nets we called $\mathbf{C}_n$ in Fig. 7.

We leave it as an instructive exercise for the reader to check that, using Lemmas 1.4, 1.5, and 1.6, the net μ of Fig. 9 is a solution to the recursive equation introduced above (we

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

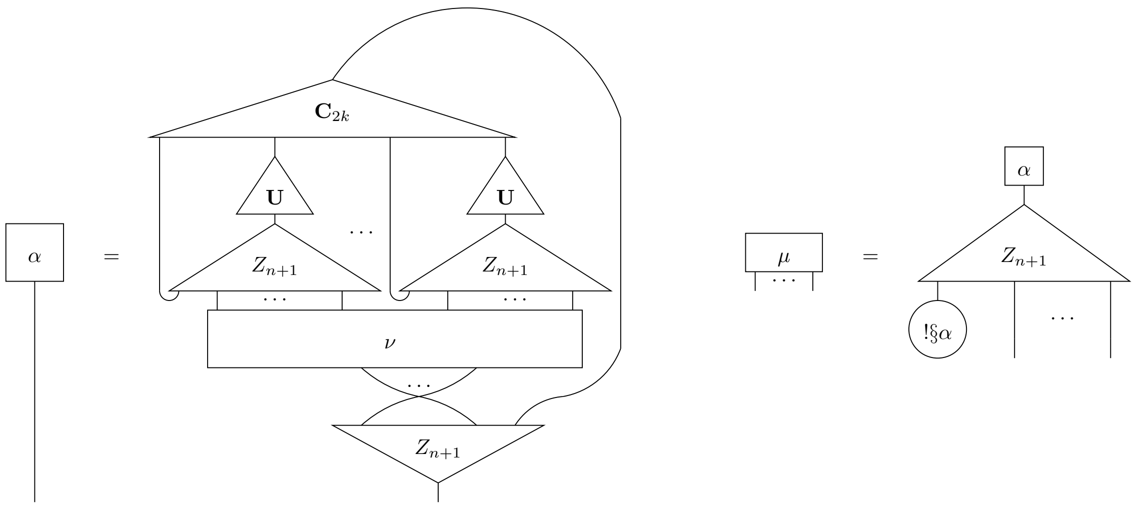

Figure 9: Solving recursive equations.

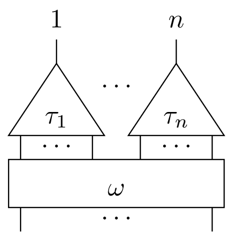

Figure 10: A principal context. The wiring ω does not connect lower ports.

have supposed that μ has n free ports and that there are k copies of μ on the right hand side of the equation).

1.4. **η-equivalence and internal separation.** In this section we recall an internal separation result similar to Böhm's theorem for the λ-calculus [Maz07b]. It will be fundamental in guiding us towards a definition of observational equivalence (Sect. 3).

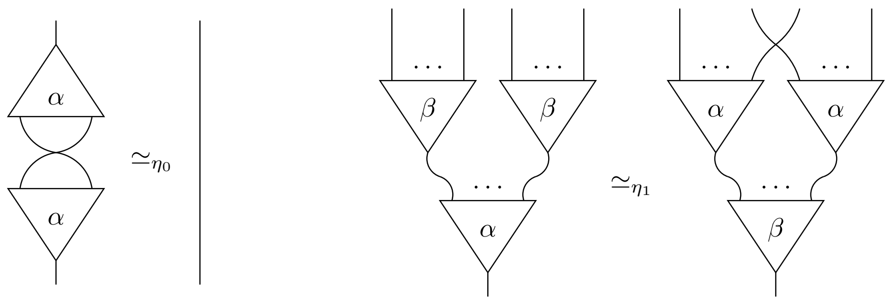

Figure 11: The equations defining η-equivalence. In the left equation, α ranges over binary symbols; in the right equation, $\alpha \neq \beta$.

Definition 1.9 (Context, principal context, test)**.** Let μ be a net with n free ports. A *context* for μ is a net C with at least n free ports. We denote by $C[\mu]$ the *application* of C to μ, which is the net obtained by plugging the free port i of μ to the free port i of C, with $i \in \{1, \dots, n\}$. A *principal context* is a context of the form given in Fig. 10. A *test* is a principal context in which ω is the identity, *i.e.*, it is a forest of n trees $\tau_1, \dots, \tau_n$ such that the root of each τ_i is the free port i.

Definition 1.10 (η- and $\beta\eta$-equivalence)**.** We define the relations $\simeq_{\eta_0}$ and $\simeq_{\eta_1}$ as the reflexive, transitive, and contextual closure of respectively the left and right equation of Fig. 11, 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 $\simeq_{\beta\eta} = (\simeq_\beta \cup \simeq_\eta)^+$.

Theorem 1.11 (Separation [Maz07b])**.** *Let μ, ν be two total nets with the same interface, such that $\mu \not\simeq_{\beta\eta} \nu$. Then, there exists a test θ such that*

$\theta[\mu] \quad \to^*$ 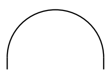 $\qquad \theta[\nu] \quad \to^*$ 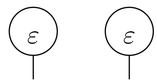

or vice versa. □

The two discriminating nets used in the theorem are not arbitrary: it is possible to prove that if a congruence containing β-equivalence equates them, then it equates all total nets. Hence, the Separation Theorem implies that any congruence on total nets containing β-equivalence is either contained in $\simeq_{\beta\eta}$, or is trivial.

The following result is the analogue of Lemma 1.4 for η-equivalence, and will be used in Sect. 4.2:

Lemma 1.12. *Let τ be a tree without ε cells. Then, we have*

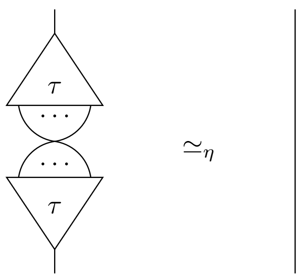

Proof. By induction on τ. □

Corollary 1.13. *For any net ν and for any trees without ε cells $\tau_1, \dots, \tau_n$, there exists a net ν' such that*

$\nu \quad \simeq_\eta$

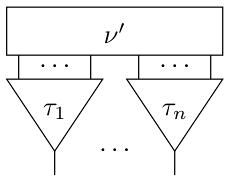

Proof. Simply "η-expand" the wires connected to the free ports of ν as in Lemma 1.12. □

2. Paths

2.1. **Straight paths.** The Separation Theorem distinguishes two nets by sending one to a net presenting a direct connection between its free ports, and the other to a net in which no such direct connection will ever form. This points out that the an appropriate notion of "connection" may be the key to a nice definition of observational equivalence.

Definition 2.1 (Port graph)**.** The *port graph* of a net μ, denoted $\mathsf{PG}(\mu)$, is the undirected multigraph whose vertices are the elements of $\mathsf{Ports}(\mu)$ and such that there is an edge between two ports p, q iff one of the following (non mutually exclusive) conditions hold:

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

Internal edges: p and q are principal and auxiliary ports of a cell of μ.

Definition 2.2 (Straight path, Danos and Regnier [DR95])**.** A *straight path* of a net μ is a path (not necessarily simple) of $\mathsf{PG}(\mu)$ which does not contain two consecutive external edges. We say that a straight path *crosses an active pair* iff it contains a wire connecting two principal ports. A *maximal path* is a non-empty straight path connecting two free ports of μ.

Definition 2.3 (Residue and lift of a maximal path)**.** Let $\mu \to \mu'$, and let ϕ be a maximal path of μ. The *residue* ϕ' of ϕ in μ' is, if it exists, the maximal path 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.

Conversely, if ϕ' is a maximal path of μ', then it is the residue of exactly one maximal path of μ, which is called the *lift* of ϕ' in μ.

Proposition 2.4. *Let $\mu \to \mu_1 \to \cdots \to \mu_{n-1} \to \mu_n$ and $\mu \to \mu'_1 \to \cdots \to \mu'_{n-1} \to \mu_n$, let ϕ_n be a maximal path of μ_n, and let $\phi_{n-1}, \dots, \phi_1, \phi$ (resp. $\phi'_{n-1}, \dots, \phi'_1, \phi'$) be the successive lifts of ϕ_n in $\mu_{n-1}, \dots, \mu_1, \mu$ (resp. $\mu'_{n-1}, \dots, \mu'_1, \mu$). Then, $\phi = \phi'$, so it makes sense to speak of "the lift" of ϕ_n in μ, independently of the reduction used to go from μ to μ_n. This will be denoted by ϕ_n^μ.*

Proof. This is a consequence of strong confluence: all reductions from μ to μ_n are just permutations of one another. □

The presence of maximal paths is preserved under anti-reduction: if $\mu \to \mu'$ and μ' has a straight path ϕ between its free ports i an j, then μ has at least one straight path between the same free ports: the lift of ϕ.

The following material will be needed in Sect. 3.4 and Sect. 4.3; the reader may safely safely skip to the next section and come back to this point when these notions are used.

Definition 2.5. Let ϕ and c be resp. a straight path and a cell of a net μ.

(1) We say that ϕ *crosses* c iff ϕ contains an internal edge of c.
(2) We say that ϕ *starts from* c iff one of the edges at the extremities of ϕ connects the principal port of c to a port other than an auxiliary port of c itself.

Definition 2.6 (Weight of a finite straight path)**.** Let ϕ be a finite straight path crossing a cell c, and let ϕ' be a subpath of ϕ starting from c and ending where ϕ ends. We define the quantity $w_\phi^{\phi'}(c)$ to be equal to zero if ϕ' crosses no active pair; otherwise, it is equal to the number of cells crossed by ϕ', *counting multiplicities*, *i.e.*, if the same cell is crossed twice by ϕ', then it is counted twice in $w_\phi^{\phi'}(c)$. Moreover, c itself is counted as any other cell in computing $w_\phi^{\phi'}(c)$. The *weight of c in ϕ* is $w_\phi(c) = \sum w_\phi^{\phi'}(c)$, where ϕ' ranges over all subpaths of ϕ starting from c. The *weight* of ϕ is $w(\phi) = \sum w_\phi(c)$, where c ranges over all cells crossed by ϕ. Note that this time *we do not count multiplicities*, *i.e.*, even if c is crossed multiple times by ϕ, its weight is still taken once in the sum.

Lemma 2.7. *Let $\mu \to \mu'$, and let ϕ' be a maximal path of μ' such that its lift ϕ in μ crosses the active pair reduced. Then, $w(\phi') < w(\phi)$.*

Proof. We only check this in the case of a commutation rule, the annihilation being simpler. Observe that ε cells need not be taken into account since ϕ and ϕ' are maximal, so we are dealing with the commutation of two binary cells. Suppose we have the following situation:

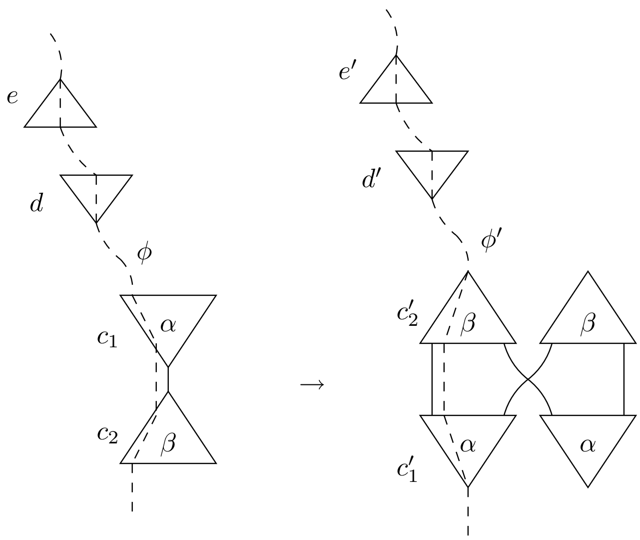

We have assumed that ϕ uses the "left" port of the α cell and the "right" port of the β cell of the active pair; this of course is just a graphical convenience, and will not affect the generality of our argument. Also, we have assumed that ϕ is not cyclic, *i.e.*, it does not come back several times to the active pair. This need not be the case; however, what is important to compute $w(\phi)$ is not ϕ itself, but its subpaths starting from the cells it crosses. Any such subpath can always be "linearized" by duplicating the cells it crosses more than once, so the above drawing is not unsound. The same holds for ϕ' of course.

Let us start by computing the weight of ϕ. We clearly have $w_\phi(c_1) = C_1 + 1$ and $w_\phi(c_2) = C_2 + 1$, where C_1, C_2 are suitable non-negative integers. For what concerns the cells crossed by ϕ other than c_1, c_2, we distinguish three kinds: those "above" c_1, those "above" c_2, and the "neutral" cells. A cell is "neutral" if the subpath of ϕ starting from it never crosses the active pair shown (like e in the picture), and is "above" c_i if it crosses the active pair by crossing c_i first (for example, d in the picture is "above" c_1). We then have $w_\phi(e) = E_e$ for any "neutral" cell e and $w_\phi(d) = D_d + 2$ for any cell "above" c_1 or c_2, where D_d, E_e are suitable non-negative integers depending on d and e. Therefore, we have $w(\phi) = C_1 + C_2 + 2 + 2n + \sum D_d + \sum E_e$, where d ranges over the cells "above" c_1 or c_2, e ranges over "neutral" cells, and n is the number of cells "above" c_1 or c_2.

Let us now turn to the weight of ϕ'. Clearly, if e' is a cell of μ' corresponding to a "neutral" cell e of μ, then nothing changed for it, so $w_{\phi'}(e') = E_e$. The situation of a cell d' corresponding to a cell d "above" c_1 is also unchanged: the presence of such cells in μ implies that the commutation has created new active pairs in μ', and these active pairs are crossed by ϕ', so $w_{\phi'}(d') = D_d + 2$, the minimal two cells crossed being c_1' and c_2'. On the other hand, for these latter two cells the situation has changed with respect to c_1 and c_2. For instance, it may happen that the subpath of ϕ' starting from c_1' does not cross any active pair, as suggested in the picture, so $w_{\phi'}(c_1') = 0$. Or, if there are cells "above" c_1 in μ, as in the picture, the subpath of ϕ' starting from c_2' certainly crosses active pairs, but the weight of the cell now is $w_{\phi'}(c_2') = C_2$. In both cases, we have $w(\phi') \leq C_1 + C_2 + 2n + \sum D_d + \sum E_e < w(\phi)$, as desired. $\square$

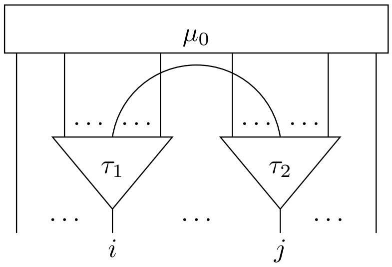

Figure 12: An observable path.

2.2. **Observable paths.** A maximal path is preserved under reduction precisely when it does not cross any active pair. In that case, we say that it is *observable*:

Definition 2.8 (Observable path)**.** Let μ be a net. An observable path of μ is a maximal path crossing no active pair.

It is perhaps useful to visualize observable paths. A net μ contains an observable path between its free ports i and j iff it is of the shape given in Fig. 12. If $i = j$, then $\tau_1 = \tau_2$, and the wire shown connects two leaves of the same tree. The actual observable path, if seen from i to j, takes the branch of τ_1 leading to the leaf connected by the wire shown, follows this wire, and descends to the root of τ_2 through the only possible branch.

Observable paths may be succinctly described by assigning *addresses* to the leaves of trees. In the following, we let a, b range over the set $\{\mathbf{p}, \mathbf{q}\}^*$ of finite binary words, and we denote by $\mathbf{1}$ the empty word. Pairs of finite words are denoted by $a \otimes b$, and ranged over by s, t. The concatenation of two finite words a, b 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 2.9 (Address of a leaf)**.** Let τ be a tree, and l a leaf of τ. The *address* of l in τ, denoted by $\mathsf{addr}_\tau(l)$, is a pair of finite binary words defined by induction on τ:[3]

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

An observable path ϕ may now be denoted by the unordered pair $\{(s, i), (t, j)\}$, where i, j are the free ports connected by the path, and s, t are the addresses of the leaves of the trees rooted at i, j connected by ϕ. In fact, from here on we shall take this as the definition of observable path, *i.e.*, any object of the form $\{(s, i), (t, j)\}$ where s, t are pairs of words and i, j ports will be referred to as an observable path, regardless of any particular net to which it may belong to.

In the following, we denote by $\mathsf{op}(\mu)$ the set of observable paths of a net μ, and we define

$$\mathsf{op}^*(\mu) = \bigcup_{\mu \to^* \mu'} \mathsf{op}(\mu').$$

[3]For the acquainted reader, $\mathsf{addr}_\tau(l)$ is nothing but the GoI weight of the path going down from l to the root of τ [Laf97]. This justifies our notations for binary words.

Proposition 2.10. *Let $\mu \to^* \mu'$. Then, $\mathsf{op}(\mu) \subseteq \mathsf{op}(\mu')$, and $\mathsf{op}^*(\mu) = \mathsf{op}^*(\mu')$.*

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

Note that a net can only have a finite number of observable paths; then, by Proposition 2.10, $\mathsf{op}^*(\mu)$ is finite whenever μ has a normal form. However, normalizability is not necessary: non-normalizable nets producing a finite number of observable paths exist, as shown by the following example (the only observable path ever to be found is the one already present in the net starting the reduction sequence, *i.e.* $\{(\mathbf{1} \otimes \mathbf{pp}, 1), (\mathbf{1} \otimes \mathbf{pq}, 1)\}$):

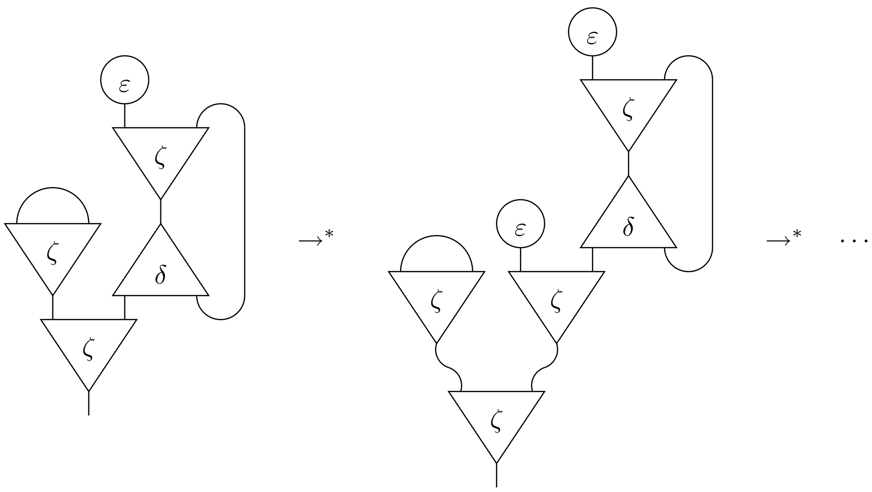

3. Observational Equivalences

3.1. **Path-based observational equivalences.** The stability of observable paths under reduction is the main reason for considering them as the base of observational equivalence.

Definition 3.1 (Observability predicates)**.** We say that μ is *immediately observable*, and we write $\mu\downarrow$, iff $\mathsf{op}(\mu) \neq \emptyset$. We say that μ is *observable*, and we write $\mu\Downarrow$, iff $\mathsf{op}^*(\mu) \neq \emptyset$, or, equivalently, $\mu \to^* \mu'\downarrow$. We say that μ is *totally observable*, and we write $\mu\Downarrow\!\!\Downarrow$, iff $\mathsf{op}^*(\mu)$ is non-empty and finite. We write $\mu\Uparrow$ and $\mu\Uparrow\!\!\Uparrow$ for the negations of $\mu\Downarrow$ and $\mu\Downarrow\!\!\Downarrow$, respectively. In particular, if $\mu\Uparrow$ we say that μ is *blind*.

Definition 3.2 (Observational equivalences)**.** Two nets μ, ν with the same interface are *observationally equivalent* (resp. *totally equivalent*), and we write $\mu \simeq \nu$ (resp. $\mu \cong \nu$), iff for all contexts C, $C[\mu]\Downarrow$ iff $C[\nu]\Downarrow$ (resp. $C[\mu]\Downarrow\!\!\Downarrow$ iff $C[\nu]\Downarrow\!\!\Downarrow$).

It is immediate to see that any two totally equivalent nets are also observationally equivalent. The converse is false, although we shall be able to prove this only at the end of the paper. It is also easy to verify the following, which is a consequence of 2.10:

Proposition 3.3. *$\mu \simeq_\beta \nu$ implies $\mu \cong \nu$ (and hence $\mu \simeq \nu$).* $\square$

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 path can in fact be extended to any interaction net system, in particular to sharing graphs [GAL92]; then, one can adapt the definition of observable net so as to obtain that a λ-term is in head normal form iff its corresponding net is immediately observable. This adaptation, which we do not detail here, takes into account only the observable paths starting from the free port representing the "root" of the term, and iteratively using the "root" of each subterm.

The existence of a "root" (*i.e.*, a distinguished free port in sharing graphs) is what allows one to define the notion of *principal* head normal form, of which no meaningful equivalent exists for nets. This is because nets, like proof nets, are "classical", as opposed to λ-terms, which are "intuitionistic". This is also the reason why the symmetric combinators equivalent of Böhm trees will not be trees (*cf.* Sect. 4.1).

Lemma 3.4. *$\mu \simeq \nu$ iff, for every* principal *context C, $C[\mu]\Downarrow$ iff $C[\nu]\Downarrow$.*

Proof. The important part of the proof is to show that one can restrict to contexts in normal form and without vicious circles. Suppose that $C[\mu]\Downarrow$ but $C[\nu]\Uparrow$. We have that $C[\mu] \to^* \mu'$ such that μ' contains an observable path. Let k be the length of the reduction leading from $C[\mu]$ to μ', and let $C_1, \ldots, C_m$ be all the k-step reducts of C, or the normal form of C if C does not admit a reduction of length k. Then, by strong confluence, there exits a C_i such that $C_i[\mu] \to^* \mu''\downarrow$ without reducing any active pair that may appear in C_i before its application to μ. Now, it can be shown (by induction on the number of cells) that every net can be decomposed in terms of trees and wirings as follows:

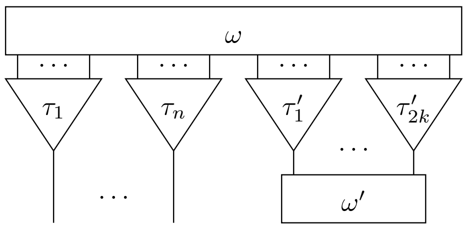

The wiring ω' accounts for the active pairs and the vicious circles; if we remove it, we obtain a normal net with no vicious circle. If we replace ω' with $2k$ ε cells in the decomposition of C_i, we obtain a net which can be shown to also have a normal form without vicious circles. Let C'_i be such a normal form. Since no active pair of C_i is reduced to obtain an observable path from $C_i[\mu]$, we also have $C'_i[\mu]\Downarrow$. On the other hand, since $C \to^* C_i$, we still have $C_i[\nu]\Uparrow$, and also $C'_i[\nu]\Uparrow$, because observable paths cannot use ε cells. Therefore, we have found a normal and vicious-circle-free context discriminating between μ and ν:

$C'_i =$ 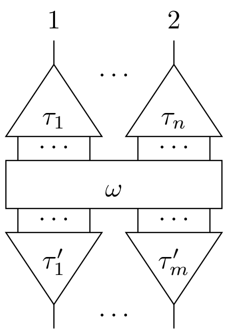

Now, any wire in ω connecting a τ'_i with a τ'_j yields an observable path which is present regardless of the application of C'_i to any net, and therefore adds no discriminative power to the context. Similarly, all the τ'_i can be completely removed, because they do not alter the observability of the net. $\square$

3.2. **Bisimilarity.**

It is possible to find an equivalent reformulation of Definition 3.2 which is useful for proving that two nets are observationally equivalent.

Definition 3.5 (Bisimulation and bisimilarity)**.** Let $\mathcal{B}$ be a binary relation relating nets with the same interface. We say that $\mathcal{B}$ is a *bisimulation* iff, whenever $(\mu, \nu) \in \mathcal{B}$, we have:

(1) $\mu\downarrow$ implies $\nu\Downarrow$;
(2) $\mu \to \mu'$ implies $\nu \to^* \nu'$ with $(\mu', \nu') \in \mathcal{B}$;
(3) $\nu\downarrow$ implies $\mu\Downarrow$;
(4) $\nu \to \nu'$ implies $\mu \to^* \mu'$ with $(\mu', \nu') \in \mathcal{B}$.

Bisimilarity, denoted $\approx$, is the union of all bisimulations.

Lemma 3.6. *$\mu \approx \nu$ iff μ and ν are either both observable, or both blind.*

Proof. The forward implication is a straightforward consequence of the definition; for the converse, we invite the reader to check that $\{(\mu, \nu) \ ; \ \mu\Downarrow \text{ iff } \nu\Downarrow\}$ is a bisimulation. $\square$

Corollary 3.7. *$\mu \simeq \nu$ iff for every context C, $C[\mu] \approx C[\nu]$.*

This makes possible the use of coinductive techniques for proving that two nets are observationally equivalent. In fact, thanks to Corollary 3.7, to prove $\mu \simeq \nu$ it is enough to show that there exists a bisimulation containing $(C[\mu], C[\nu])$ for all C. This may be of help because, instead of having to prove $C[\mu]\Downarrow \Leftrightarrow C[\nu]\Downarrow$, one can first prove the weaker implications $C[\mu]\downarrow \Rightarrow C[\nu]\Downarrow$ and $C[\nu]\downarrow \Rightarrow C[\mu]\Downarrow$, and then show that $C[\mu]$ and $C[\nu]$ are able to simulate the one-step reductions of each other. Of course there is still the quantification over all contexts which may render the task difficult, but in many interesting cases this is just what will be needed.

An example is given by the following. For all $n \in \mathbb{N}$, we define $\mathbf{E}_n$ to be the net with n free ports consisting of n ε cells, ε_0 being the empty net. These nets are the prototypical blind nets: they contain no observable path, and they do not develop any, since they are normal. In fact, we shall see that any blind net with n free ports is observationally equivalent to $\mathbf{E}_n$.

Let C be a net, and let I be a subset of its interface. We say that C is *relatively blind* on I no reduct of C has an observable path connecting two ports of I. A context C for nets with n free ports will be said to be relatively blind if its interface is $\{1, \ldots, n\} \uplus I$ and C is relatively blind on I. Principal contexts are the typical examples of relatively blind contexts. Consider now the relation

$$\mathcal{B} = \{(C[\mu], \mathbf{E}_n) \mid \forall \mu\Uparrow \text{ with } n \text{ free ports}, \ \forall \text{ relatively blind context } C\}.$$

Lemma 3.8. *$\mathcal{B}$ is a bisimulation.*

Proof. Parts 1, 3, and 4 of Definition 3.5 vacuously hold, so let us concentrate on part 2. We have three cases, depending on the nature of the active pair reduced in $C[\mu]$. If $C[\mu] \to C'[\mu]$, *i.e.*, the active pair is within C, then $(C'[\mu], \mathbf{E}_n) \in \mathcal{B}$, because by definition C' is still relatively blind. If $C[\mu] \to C[\mu']$, *i.e.*, the active pair is in μ, then we still have

$\mu'\Uparrow$, and so $(C[\mu'], \mathbf{E}_n) \in \mathcal{B}$. We are left with the case in which $C[\mu]$ reduces because of an active pair created by plugging μ into C. There are several different subcases; we only check the most difficult one, leaving the others to the reader. Suppose we have

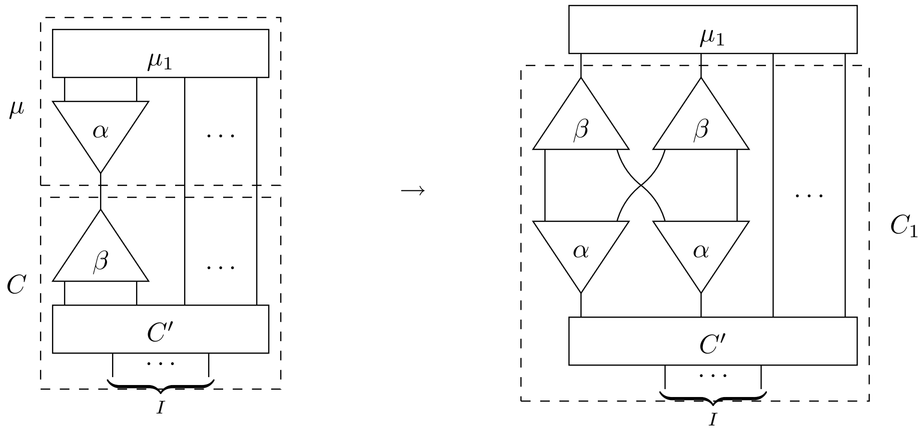

We start by observing that $\mu\Uparrow$ implies $\mu_1\Uparrow$, and that C relatively blind on I implies C' relatively blind on I too. We claim that C_1 is also relatively blind on I. In fact, if any of its reducts contained an observable path between two of the free ports of I, then such a path would have a lift ϕ in C_1. Now, if ϕ were entirely contained in C', then C' would not be relatively blind; hence ϕ must go through the subnet of C_1 which does not belong to C', i.e., the reduct of the active pair. But such subnet contains no straight path allowing to "go up" from I and then "go down" again towards I. Therefore, ϕ cannot exist, and $(C_1[\mu_1], \mathbf{E}_n) \in \mathcal{B}$. □

The above result shows that blind nets are "resistant" to contexts:

Corollary 3.9. *Let μ be a net. Then, $\mu\Uparrow$ implies $C[\mu]\Uparrow$ for any principal context C.* □

Since, for C principal, $C[\mathbf{E}_n] \to^* \mathbf{E}_m$ (this is a variant of Lemma 1.7), by Lemma 3.4 we obtain $\mu \simeq \mathbf{E}_n$ for all $\mu\Uparrow$ with n free ports, as anticipated. In the next section we shall see that we actually also have $\mu \cong \mathbf{E}_n$.

A useful variant of the notion of bisimulation is *bisimulation up to reduction*:

Definition 3.10 (Bisimulation up to reduction)**.** Let $\mathcal{B}$ be a binary relation relating nets with the same interface. We say that $\mathcal{B}$ is a *bisimulation up to reduction* iff, whenever $(\mu, \nu) \in \mathcal{B}$, we have:

(1) $\mu\downarrow$ implies $\nu\Downarrow$;
(2) $\mu \to \mu'$ implies $\nu \to^* \nu'$ and $\mu' \to^* \mu''$ such that $(\mu'', \nu') \in \mathcal{B}$;
(3) $\nu\downarrow$ implies $\mu\Downarrow$;
(4) $\nu \to \nu'$ implies $\mu \to^* \mu'$ and $\nu' \to^* \nu''$ such that $(\mu', \nu'') \in \mathcal{B}$.

Lemma 3.11. *If $\mathcal{B}$ is a bisimulation up to reduction, then $\to^* \mathcal{B}\ {}^*\!\!\leftarrow$ is a bisimulation.*

Proof. Let $\mathcal{B}$ be a bisimulation up to reduction. We shall only check that properties 1 and 2 of Definition 3.5 hold, the other two being perfectly symmetrical. So set $\mathcal{B}' = \to^* \mathcal{B}\ {}^*\!\!\leftarrow$, and let $(\mu, \nu) \in \mathcal{B}'$, which means that $\mu \to^* \mu_1$ and $\nu \to^* \nu_1$ such that $(\mu_1, \nu_1) \in \mathcal{B}$. For what concerns property 1, by Proposition 2.10, we have that $\mu\downarrow$ implies $\mu_1\downarrow$, so $\nu_1\Downarrow$ because $\mathcal{B}$ is a bisimulation up to reduction, and therefore $\nu\Downarrow$ again by Proposition 2.10. Now let $\mu \to \mu'$. If μ_1 is a reduct of μ', then we immediately obtain $(\mu', \nu) \in \mathcal{B}'$. Otherwise, by strong confluence (Proposition 1.3), we have that there exists μ_2 such that $\mu' \to^* \mu_2$, and $\mu_1 \to \mu_2$ (all reduction sequences from μ to μ_2 have the same length). In this case, $(\mu_1, \nu_1) \in \mathcal{B}$ implies that $\nu_1 \to^* \nu'$ such that $(\mu_2, \nu') \in \mathcal{B}'$; but μ_2 is a reduct of μ', and ν_1 is a reduct of ν, hence $\nu \to^* \nu'$ with $(\mu', \nu') \in \mathcal{B}'$. □

Corollary 3.12. *To prove $\mu \simeq \nu$, it is enough to find a bisimulation up to reduction $\mathcal{B}$ such that $(C[\mu], C[\nu]) \in \mathcal{B}$ for all C.*

We shall use bisimulations up to reduction to prove that $\beta\eta$-equivalent nets are observationally equivalent, which we do next. In the following, we write $C[\mu_1, \ldots, \mu_n]$ for a net containing $\mu_1, \ldots, \mu_n$ as subnets, i.e., C is a "multi-hole context".

Lemma 3.13. *We define the relation $\rightleftharpoons$ as follows:*

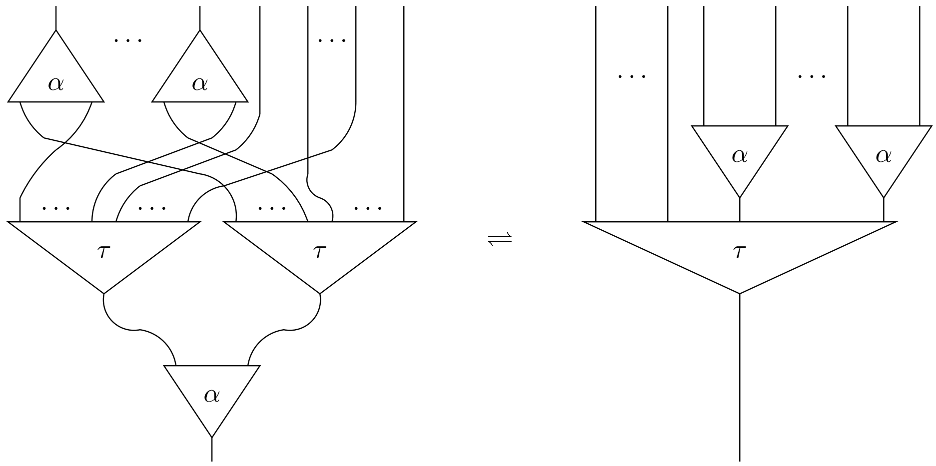

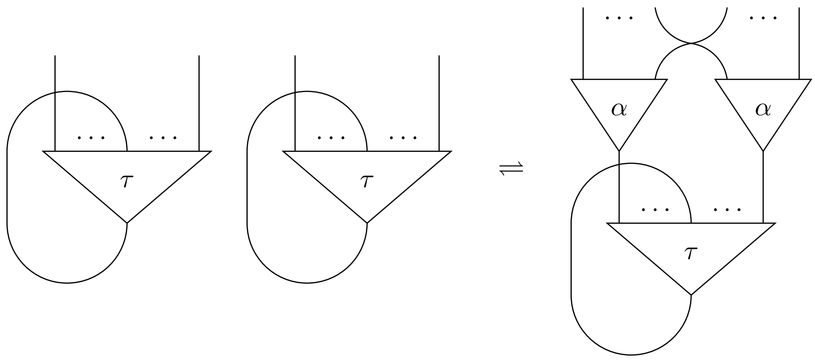

In both equations, α ranges over binary symbols and τ ranges over trees not containing α cells. Then, the relation

$$\mathcal{B} = \{(C[\mu_1, \ldots, \mu_k], C[\nu_1, \ldots, \nu_k]) \ ; \ \mu_i \rightleftharpoons \nu_i, \text{ for all multi-hole contexts } C\}$$

is a bisimulation up to reduction.

Proof. The proof is extremely long, and not particularly interesting. We only check a few cases, leaving the others to the reader. Let $(\mu, \nu) \in \mathcal{B}$, so that $\mu = C[\mu_1, \ldots, \mu_k]$ and $\nu = C[\nu_1, \ldots, \nu_k]$ for some k-hole context C, with $\mu_i \rightleftharpoons \nu_i$ for all $1 \le i \le k$. Observe that, if $k = 0$, then $\mu = \nu$, and there is nothing to check. So we can write for short $\mu = C[\mu_0, \overrightarrow{\mu}]$ and $\nu = C[\nu_0, \overrightarrow{\nu}]$. Suppose that $\mu\downarrow$. First of all, note that the second equation defining $\rightleftharpoons$ concerns blind nets, so it does not affect observable paths. Additionally, the first equation obviously preserves observable paths when applied from left to right, so the only case we need to check is that in which ν is obtained from μ by using the first equation from right to left. We can then suppose that the observable path in μ passes through μ_0, and that μ_0 and ν_0 are resp. the right and left member of the first equation. The interesting situation is

$\mu =$ 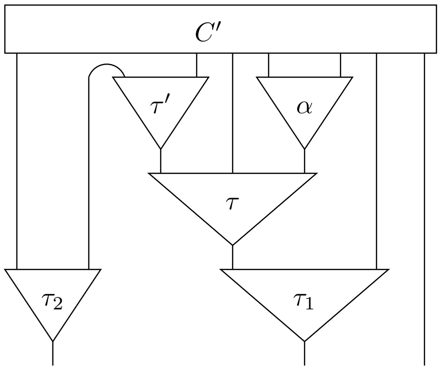

and

$\nu =$ 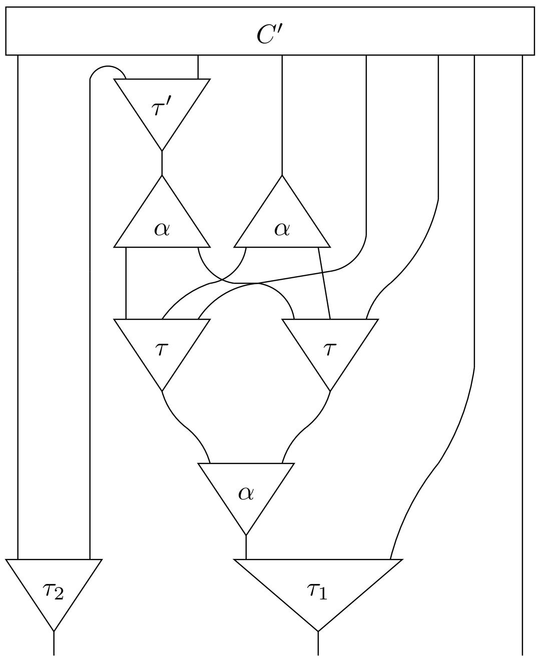

The observable path in μ is the one shown going from the root of τ_1 to the root of τ_2. To economize on notations, we have limited as much as possible the arities of all trees in the picture; this will not affect the generality of our argument. This case is the interesting one because ν contains an active pair which "breaks" the observable path that is present in μ, so that ν may no longer be immediately observable. However, we still have $\nu\Downarrow$. This is proved by induction on τ'. Observe that we must have $\tau' \neq \varepsilon$, so the base case is $\tau' = \bullet$, in which obviously $\nu\downarrow$. In case $\tau' = \alpha(\tau_1', \tau_2')$, it is easy to see that $\nu \to \nu'\downarrow$; in case $\tau' = \beta(\tau_1', \tau_2')$ with $\beta \neq \alpha$, we have

$\nu \to^*$

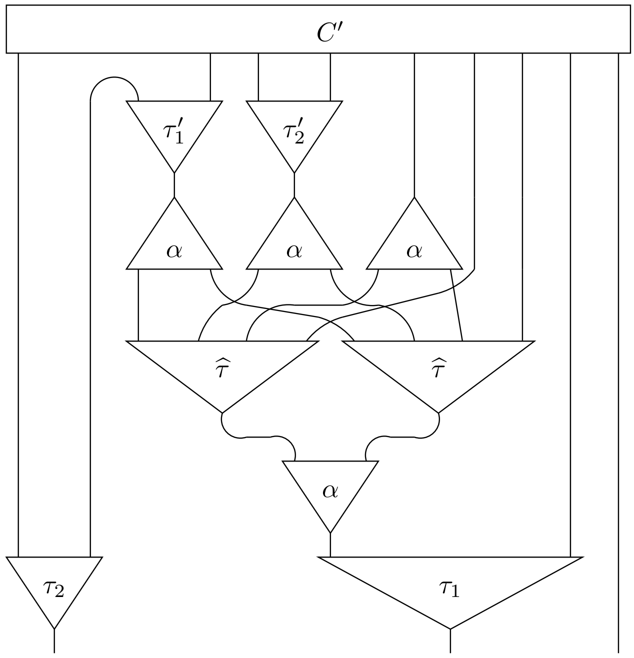

where $\widehat{\tau}$ is a tree formed by τ plus a β cell, and therefore does not contain α cells. We have assumed that the observable path uses a leaf of τ_1', but nothing changes if a leaf of τ_2' is used instead: in any case, we conclude $\nu\Downarrow$ by applying the induction hypothesis. This is all as far as property 1 of Definition 3.10 is concerned; property 3 is completely symmetric, and is established in the same way.

We thus turn to verify property 2. Suppose that $\mu \to \mu'$. If the reduction takes place inside C, ν has no problem in simulating it, so we may assume that one of the μ_i has interacted. Again, we put $\mu = C[\mu_0, \overrightarrow{\mu}]$ and $\nu = C[\nu_0, \overrightarrow{\nu}]$, with μ_0 and ν_0 resp. left and right side of the first equation above; the second equation does not need to be checked, because the configurations related by it cannot be involved in any active pair.

We start by supposing that a cell of C has interacted with the "root" of μ_0, i.e., the principal port of the α cell drawn at the bottom of the left hand side of the first equation. We have three cases, one for each possible cell. We briefly go over them, leaving the details to the reader:

- ε cell: we have $\mu' \to^* C'[\mu_0', \overrightarrow{\mu}]$ and $\nu \to^* C'[\nu_0', \overrightarrow{\nu}]$, with $\mu_0' \rightleftharpoons \nu_0'$.
- α cell: if we put $\mu' = C'[\overrightarrow{\mu}]$, using Lemma 1.5 we get that $\nu \to^* C'[\overrightarrow{\nu}]$.
- β cell, with $\beta \neq \alpha$: we reason by analyzing the structure of τ.
  - $\tau = \varepsilon$: this is one case in which multi-hole contexts are necessary. We have $\mu' \to^* C'[\mu_0, \mu_0, \overrightarrow{\mu}]$ and $\nu \to C'[\nu_0, \nu_0, \overrightarrow{\nu}]$.
  - $\tau = \bullet$: here we have $\mu' = C'[\mu_0', \overrightarrow{\mu}]$, and $\nu = C'[\nu_0', \overrightarrow{\nu}]$, with $\mu_0' \rightleftharpoons \nu_0'$.
  - $\tau = \beta(\tau_1, \tau_2)$: here multi-hole contexts are needed again. We have $\mu' \to^* C'[\mu_0', \mu_0'', \overrightarrow{\mu}]$ and $\nu \to C'[\nu_0', \nu_0'', \overrightarrow{\nu}]$, with $\mu_0' \rightleftharpoons \nu_0'$ and $\mu_0'' \rightleftharpoons \nu_0''$.

  The case $\tau = \alpha(\tau_1, \tau_2)$ is not possible, since in the first equation τ is a tree not containing α cells.

In all cases, one sees that $\mu' \to^* \mu''$ and $\nu \to^* \nu''$ such that $(\mu'', \nu'') \in \mathcal{B}$, as required by point 2 of Definition 3.10.

The next cases to be checked are those in which $\mu \to \mu'$ because of an interaction with one of the principal ports of the α cells at the "top" of μ_0. In all cases, one can see that $\mu' = C'[\mu'_0, \overrightarrow{\mu}]$, and that ν is actually equal to $C'[\nu'_0, \overrightarrow{\nu}]$, with $\mu'_0 \rightleftharpoons \nu'_0$.

It is also possible that $\mu \to \mu'$ because μ_0 has interacted with itself; in that case, we have $\mu' = C'[\mu'_0, \overrightarrow{\mu}]$ and $\nu = C'[\nu'_0, \overrightarrow{\nu}]$, with $\mu'_0 \rightleftharpoons \nu'_0$ because of the second equation (this is in fact the only time this equation is ever used).

The last cases to be checked are those in which $\mu \to \mu'$ because a μ_i interacts with a μ_j, with $i \neq j$. There is quite a big number of subcases to be considered; none of them poses any particular problem, so we leave them to the reader.

It remains to prove that $\mathcal{B}$ satisfies also property 4 of Definition 3.10; this is done by following the same arguments used above, and is in fact easier, because many of the above cases do not exist when we consider the reduction $\nu \to \nu'$. $\square$

Proposition 3.14. *Let μ, ν be two nets. Then, $\mu \simeq_{\beta\eta} \nu$ implies $\mu \simeq \nu$.*

Proof. Let $\leftrightarrows$ be the contextual closure of the relation $\rightleftharpoons$ defined in Lemma 3.13. Observe that the equations of Fig. 11 are special cases of the first equation defining $\rightleftharpoons$: the η_0 equation is obtained by taking $\tau = \bullet$; the η_1 equation of Fig. 11 is obtained by taking τ to be equal to either ε or $\beta(\bullet, \bullet)$ with $\beta \neq \alpha$, and taking no α cells at the "top" of the left hand side. Therefore, for all μ, ν, if μ and ν are obtained one from the other by applying one equation of Fig. 11, we have $\mu \leftrightarrows \nu$, which by Lemma 3.13 implies $\mu \approx \nu$ (in fact, $\leftrightarrows \subseteq \mathcal{B}$, where $\mathcal{B}$ is the bisimulation up to reduction of Lemma 3.13). But $\approx$ is transitive, so actually for all μ, ν, $\mu \simeq_\eta \nu$ implies $\mu \approx \nu$, which in turn implies $\mu \simeq \nu$, because $\simeq_\eta$ is a congruence. Now, we have already seen that $\simeq_\beta \subseteq \simeq$ (Proposition 3.3), so the proposition is a consequence of the definition of $\beta\eta$-equivalence and the transitivity of $\simeq$. $\square$

In the next section we shall see that $\beta\eta$-equivalent nets are actually totally equivalent. Hence, thanks to the Separation Theorem 1.11, $\cong$, $\simeq$ and $\simeq_{\beta\eta}$ all coincide on total nets.

3.3. **ε-reduction.** The fact that all blind nets with n free ports are observationally equivalent to $\mathbf{E}_n$ suggests the following rewriting rule:

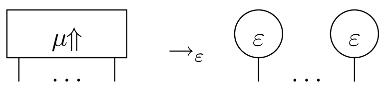

We write $\to_\varepsilon$ for the relation obtained from $\to$ plus the above rule applied under any context, *i.e.*, whenever μ is a blind subnet of a larger net. As usual, we write $\to_\varepsilon^*$ for the reflexive transitive closure of $\to_\varepsilon$. Of course the ε-step is not recursive, because it is undecidable whether a net is blind. In this sense, ε-reduction is very similar to $\beta\Omega$-reduction in the λ-calculus [Bar84]. The interest of ε-reduction is its relationship with total equivalence.

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 3.15. *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. We actually prove something stronger: if $\mu \rightsquigarrow^* \mu_1$ in n_1 steps, and $\mu \rightsquigarrow^* \mu_2$ in n_2 steps, then there exists ν such that $\mu_1 \rightsquigarrow^* \nu$ in at most n_2 steps, and $\mu_2 \rightsquigarrow^* \nu$ in at most n_1 steps. We do this by induction on $n_1 + n_2$, which we refer to as the *degree* of the triple (μ, μ_1, μ_2). If the degree is zero, then everything is trivial. If the degree is non-zero but one of n_1, n_2 is zero, then again the proof is trivial. So let us suppose $n_1, n_2 > 0$. We then have $\mu \rightsquigarrow \mu_1' \rightsquigarrow^* \mu_1$ and $\mu \rightsquigarrow \mu_2' \rightsquigarrow^* \mu_2$. By the quasi-diamond property, there exists ν' such that $\mu_1' \rightsquigarrow^* \nu'$ and $\mu_2' \rightsquigarrow^* \nu'$ in at most one step, *i.e.*, ν' may be equal to μ_1' or μ_2'. Consider now the triple (μ_1', ν', μ_1). Its degree is $n_1 - 1 + c$, where $c \leq 1$, so the degree is at most $n_1 < n_1 + n_2$, because we have supposed $n_2 \neq 0$. So we can apply the induction hypothesis, and obtain a net ν_1 such that $\nu' \rightsquigarrow^* \nu_1$ in at most $n_1 - 1$ steps and $\mu_1 \rightsquigarrow^* \nu_1$ in at most one step. The same holds for the triple (μ_2', ν', μ_2), for which we obtain a net ν_2 such that $\nu' \rightsquigarrow^* \nu_2$ in at most $n_2 - 1$ steps and $\mu_2 \rightsquigarrow^* \nu_2$ in at most one step. Then, if we consider the triple (ν', ν_1, ν_2), its degree is at most $n_1 + n_2 - 2 < n_1 + n_2$, so the induction hypothesis applies again, giving us a net ν such that $\nu_1 \rightsquigarrow^* \nu$ and $\nu_2 \rightsquigarrow^* \nu$. Composing the reductions yields the diamond for the original triple (μ, μ_1, μ_2). □

Lemma 3.16. *The relation $\to_\varepsilon$ satisfies the quasi-diamond property.*

Proof. Let $\mu \to_\varepsilon \mu_1$ and $\mu \to_\varepsilon \mu_2$. 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 1.3. Otherwise, suppose without loss of generality that $\mu \to \mu_1$ and $\mu \to_\varepsilon \mu_2$ by means of an ε-step, *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:

- $\nu \to \nu'$ 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 $\mu_1 \to_\varepsilon \mu_2$.
- 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 principal context C_0, and we can write $\mu = C'[C_0[\nu]]$ for a suitable context C'. Now $C_0[\nu] \to \nu'$ and $\mu_1 = C'[\nu']$, while $\mu_2 = C'[C_0[\mathbf{E}_n]]$. By Corollary 3.9, we have both $C_0[\nu]\Uparrow$ and $C_0[\mathbf{E}_n]\Uparrow$, 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

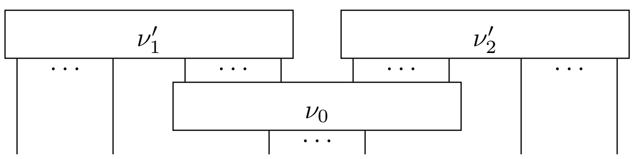

and ν_1 is equal to ν_1' plus ν_0, while ν_2 is equal to ν_2' plus ν_0. Now, if we put

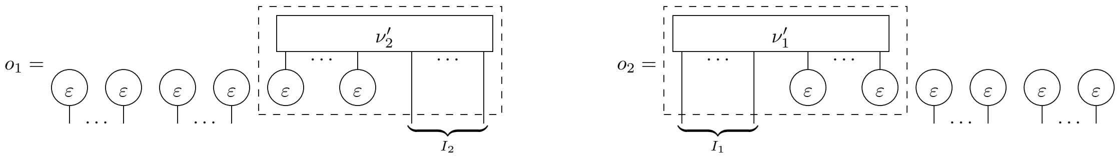

we have $\mu_1 = C[o_1]$ and $\mu_2 = C[o_2]$. But ν_1' and ν_2' must be relatively blind on I_1 and I_2, respectively, because ν_1 and ν_2 are blind. Hence, by Lemma 3.8, the subnets marked by the dashed rectangles in the above picture are both blind, so μ_1 and μ_2 both reduce in at most one ε-step to $C[\mathbf{E}_n]$, where n is the number of free ports of ν. □

Proposition 3.17 (Confluence of ε-reduction). *The relation $\to_\varepsilon^*$ is confluent.*

Proof. Apply Lemmas 3.15 and 3.16. □

The confluence of ε-reduction allows us to define *$\beta\varepsilon$-equivalence* (which is actually a congruence) as $\mu \simeq_{\beta\varepsilon} \nu$ iff there exists o such that $\mu \to_\varepsilon^* o$ and $\nu \to_\varepsilon^* o$. Similarly, one defines *$\beta\varepsilon\eta$-equivalence* as $\simeq_{\beta\varepsilon\eta} = (\simeq_{\beta\varepsilon} \cup \simeq_\eta)^+$. Clearly ε-steps do not alter observability, so the following statement is an obvious consequence of Proposition 3.14:

Lemma 3.18. *Let $\mu \simeq_{\beta\varepsilon\eta} \nu$. Then, $\mu\Downarrow$ iff $\nu\Downarrow$.* □

$\beta\varepsilon$-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* as a tree with no leaves; the ε-tree ε is called *trivial*.

Proposition 3.19 ($\beta\varepsilon$-normal forms). *A net μ is $\beta\varepsilon$-normal iff it is cut-free, and each ε-tree contained in μ is trivial.* □

One the main structural difference between reduction and ε-reduction is that the second one, although confluent, is not *strongly* confluent. However, we can define a sub-reduction which has this useful property.

Definition 3.20 (Big step ε-reduction). We write $\mu \twoheadrightarrow_\varepsilon \nu$ iff $\mu \to_\varepsilon \nu$ and:

- if an active pair is reduced, then it is not contained in a blind subnet of μ;
- if $\mu = C[\mu_0]$ and $\nu = C[\mathbf{E}_n]$ with $\mu_0\Uparrow$ with n free ports, then μ_0 is not a proper subnet of a blind subnet of μ.

We write $\twoheadrightarrow_\varepsilon^*$ for the reflexive-transitive closure of $\twoheadrightarrow_\varepsilon$.

Proposition 3.21 (Strong confluence of big step ε-reduction). *The relation $\twoheadrightarrow_\varepsilon^*$ is strongly confluent.*

Proof. Simply observe, by looking at the proof of Lemma 3.16, that the only cases in which $\to_\varepsilon$ does not satisfy the full diamond property are precisely those explicitly forbidden by the definition of big step ε-reduction. □

Clearly, a net is $\beta\varepsilon$-normalizable iff it is normalizable under the big step reduction. Thanks to strong confluence, weak normalization and strong normalization coincide in the big step reduction, which is not the case for usual ε-reduction. Hence, given a $\beta\varepsilon$-normalizable net μ, we can define the quantity $\|\mu\|$ as the number of steps needed to get from μ to its $\beta\varepsilon$-normal form under the big step reduction.

We shall now prove a result similar to Lemma 3.18, but with $\beta\varepsilon$-normalizability instead of observability.

Lemma 3.22. *Let $\mu \simeq_{\eta_1} \nu$. Then, μ is $\beta\varepsilon$-normalizable iff ν is.*

Proof. It is enough to verify that, given a $\beta\varepsilon$-normalizable net μ, if ν is obtained from μ by applying exactly one η_1 equation, then ν is also $\beta\varepsilon$-normalizable. First of all, observe that taking α or β to be equal to ε in the η_1 equation yields $\mu \simeq_{\beta\varepsilon} \nu$, so the result is trivial. Then, we have $\mu = C[\tau]$ and $\nu = C[\tau']$, where τ, τ' are two trees of binary cells related by the η_1 equation. We reason by induction on $\|\mu\|$. If μ is already $\beta\varepsilon$-normal, then obviously also ν is, so the base case holds. Suppose now $\|\mu\| > 0$. If C is not $\beta\varepsilon$-normal, then we have $\mu \to_\varepsilon C'[\tau]$ and $\nu \to_\varepsilon C'[\tau']$, so the lemma holds by induction hypothesis. Otherwise, the only steps reducing μ concern τ. If τ is involved in an active pair, then we invite the reader to check that $\mu \simeq_\beta \nu$; if, on the other hand, τ is contained in a blind subnet of μ which is the subject of an ε-step, then by Proposition 3.14 this step can also be applied to ν, and we obtain $\mu \simeq_{\beta\varepsilon} \nu$. In both cases, ν is clearly $\beta\varepsilon$-normalizable. □

Let ν_α be the left hand side of the η_0 equation, with α a fixed binary symbol. In what follows, we write $\nu \rightsquigarrow^{\alpha}_{\eta_0} \nu'$ iff ν' is obtained from ν by replacing some wires with ν_α.

Lemma 3.23. *Let ν be a $\beta\varepsilon$-normal net, and let $\nu \rightsquigarrow^{\alpha}_{\eta_0} \nu'$. Then, ν' is $\beta\varepsilon$-normalizable.*

Proof. We prove the statement by induction on the number of active pairs of ν'. If ν' is normal, the lemma holds trivially. Otherwise, ν' contains a subnet μ_0 such that

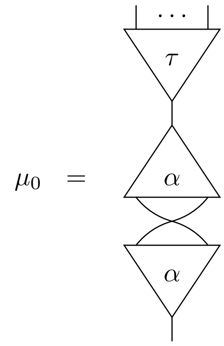

where τ is "maximal", *i.e.*, its leaves of τ are either free or connected to an auxiliary port. This is not possible only if one of the leaves of τ is connected to the principal port of the α cell shown at the bottom of the picture; but in this case ν would contain a vicious circle, which is forbidden by Proposition 3.19. The tree τ can always be decomposed as follows:

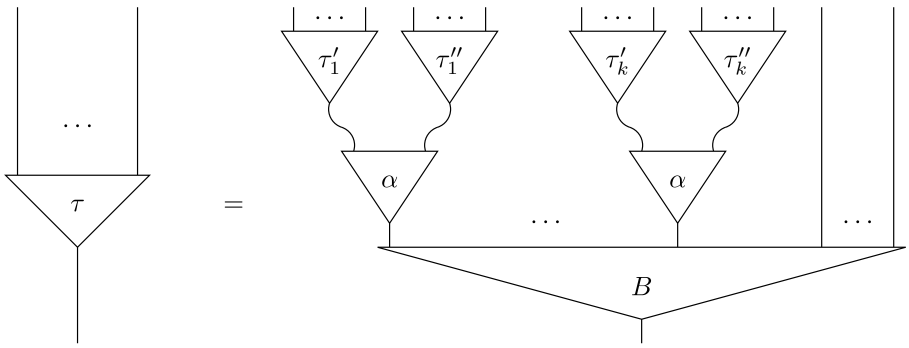

where B is a tree not containing α cells, and $\tau'_1, \ldots, \tau'_k$, $\tau''_1, \ldots, \tau''_k$ are trees (k may be equal to zero). Now, by Lemma 1.5, we have

$\mu_0 \quad \to^*$

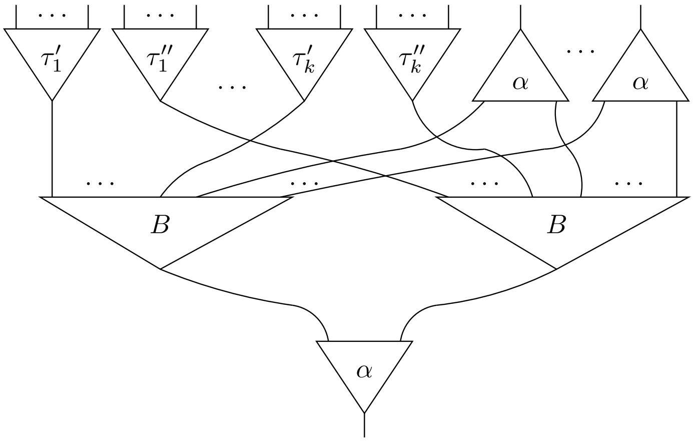

$= \quad \mu_0'$

By induction on B we can now prove that

$\mu_0' \quad \simeq_{\eta_1}$

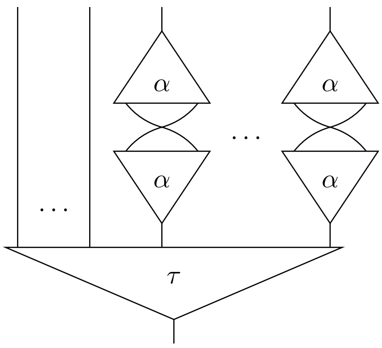

$= \quad \mu_0''$

To sum up, we have $\nu' = C[\mu_0] \to^* C[\mu_0'] \simeq_{\eta_1} C[\mu_0''] = \nu''$. Observe now that $\nu \rightsquigarrow^{\alpha}_{\eta_0} \nu''$, and ν'' has one active pair less than ν' (this is because of the "maximality" of τ). Hence, we obtain the result by applying the induction hypothesis and Lemma 3.22. $\square$

Lemma 3.24. *Let $\mu \rightsquigarrow^{\alpha}_{\eta_0} \nu$. Then, μ is $\beta\varepsilon$-normalizable iff ν is.*

Proof. We start by supposing μ $\beta\varepsilon$-normalizable, and proceed by induction on $\|\mu\|$. If $\|\mu\| = 0$, we apply Lemma 3.23. So let us suppose $\|\mu\| > 0$. Observe that, if we write ω for a single wire, we have $\mu = C[\omega, \dots, \omega]$ and $\nu = C[\nu_\alpha, \dots, \nu_\alpha]$ for some multi-hole context C. If C is not $\beta\varepsilon$-normal, then we have $\mu \to_\varepsilon C'[\omega, \dots, \omega] = \mu'$, $\nu \to_\varepsilon C'[\nu_\alpha, \dots, \nu_\alpha] = \nu'$, with $\mu' \rightsquigarrow^{\alpha}_{\eta_0} \nu'$ and $\|\mu'\| < \|\mu\|$, so we conclude by applying the induction hypothesis. We must then suppose that all reduction steps that can be applied to μ concern at least one of the wires replaced by ν_α in ν. Let us start with the case in which one of such wires connects two principal ports of μ, and let β, γ be the symbols of the two cells involved in the active pair. Assume first that one of the two, say β, is equal to ε or to α:

- $\beta = \varepsilon$: we can easily check that $\mu \simeq_{\beta\varepsilon} \nu$, so ν is $\beta\varepsilon$-normalizable;
- $\beta = \alpha$: in this case we even have $\mu \simeq_\beta \nu$, so we can also conclude.

We may then assume that $\beta = \gamma$ is a binary symbol different from α. Let μ_0 be the active pair composed of two β cells, and let ν_0 be μ_0 in which the wire connecting the principal ports of the active pair is replaced by ν_α. Now, it is just a matter of going through a few reductions to check that $\mu = C'[\omega, \dots, \omega, \mu_0] \to C'[\omega, \dots, \omega, \omega, \omega] = \mu'$, while $\nu = C'[\nu_\alpha, \dots, \nu_\alpha, \nu_0] \to^* C'[\nu_\alpha, \dots, \nu_\alpha, \nu_\alpha, \nu_\alpha] = \nu'$. Now $\|\mu'\| < \|\mu\|$, and $\mu' \rightsquigarrow^{\alpha}_{\eta_0} \nu'$, so we conclude again by applying the induction hypothesis. The other possibility is that some wires of μ which are replaced with ν_α in ν belong to a blind subnet, *i.e.*, we have $\mu = C'[\omega, \ldots, \omega, C_0[\omega, \ldots, \omega]]$ with $\mu_0 = C_0[\omega, \ldots, \omega]$ a blind net. Then, we have $\nu = C'[\nu_\alpha, \ldots, \nu_\alpha, C_0[\nu_\alpha, \ldots, \nu_\alpha]]$. If we put $\nu_0 = C_0[\nu_\alpha, \ldots, \nu_\alpha]$, we clearly have $\nu_0 \simeq_\eta \mu_0$, so $\nu_0\Uparrow$ by Proposition 3.14. Then, $\mu \to_\varepsilon \mu'$ and $\nu \to_\varepsilon \nu'$ with $\mu' \rightsquigarrow_{\eta_0}^{\alpha} \nu'$. But $\|\mu'\| < \|\mu\|$, so we can apply once more the induction hypothesis, and conclude. This is all for what concerns the "only if" part of the statement. We must now suppose ν $\beta\varepsilon$-normalizable, and prove that μ is also $\beta\varepsilon$-normalizable. This is done using similar arguments to the ones given above, and is actually easier (for example, the base of the induction is trivial). □

Proposition 3.25. *Let $\mu \simeq_{\beta\varepsilon\eta} \nu$. Then, μ is $\beta\varepsilon$-normalizable iff ν is.*

Proof. It is enough to check that the proposition holds when μ and ν are obtained from each other by applying a single equation of Fig. 11. If it is the η_1 equation, we conclude by Lemma 3.22. If it is the η_0 equation, then we have either $\mu \rightsquigarrow_{\eta_0}^{\alpha} \nu$ or $\nu \rightsquigarrow_{\eta_0}^{\alpha} \mu$, for some binary symbol α. In both cases, we conclude by Lemma 3.24. □

The following result shows that ε-reduction and total equivalence are related. We first need to extend Proposition 2.10 to ε-reduction, which is unproblematic:

Lemma 3.26. *$\mu \to_\varepsilon^* \mu'$ implies $\mathsf{op}(\mu) \subseteq \mathsf{op}(\mu')$.* □

Lemma 3.27. *μ is $\beta\varepsilon$-normalizable iff $\mathsf{op}^*(\mu)$ is finite.*

Proof. The forward implication is an immediate consequence of Lemma 3.26. For the converse, $\mathsf{op}^*(\mu)$ finite and Lemma 3.26 imply that any reduction starting from μ stumbles upon a net μ' such that $\mathsf{op}(\mu') = \mathsf{op}^*(\mu)$. This means that all subnets of μ' containing active pairs are blind, *i.e.*, they do not produce further observable paths. 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. □

We may now extend Proposition 3.14 to total equivalence, as anticipated at the end of the previous section.

Proposition 3.28. *$\mu \simeq_{\beta\varepsilon\eta} \nu$ implies $\mu \cong \nu$.*

Proof. Suppose $\mu \ncong \nu$. By definition, there exists a context C such that, for example, $C[\mu]\Downarrow$ and $C[\nu]\Uparrow$. This last statement means that $\mathsf{op}^*(C[\nu])$ is either empty, or infinite. Suppose that the latter is true. By Lemma 3.27, $C[\nu]$ is not $\beta\varepsilon$-normalizable; on the other hand, still by Lemma 3.27, $C[\mu]$ is $\beta\varepsilon$-normalizable, because $\mathsf{op}^*(C[\mu])$ is finite. Then, by Proposition 3.25, $C[\mu] \not\simeq_{\beta\varepsilon\eta} C[\nu]$, and so $\mu \not\simeq_{\beta\varepsilon\eta} \nu$, because $\simeq_{\beta\varepsilon\eta}$ is a congruence. Suppose now that $C[\nu]\Uparrow$. Observe that $C[\mu]\Downarrow$ implies $C[\mu]\Downarrow$, so we conclude by Lemma 3.18. □

3.4. **The Context Lemma.** This section is devoted to a proof of the following result, saying that, according to observational equivalence, tests suffice to discriminate nets:

Lemma 3.29 (Context)**.** *$\mu \simeq \nu$ iff, for every test θ, $\theta[\mu]\Downarrow$ iff $\theta[\nu]\Downarrow$.* □

This lemma is crucial in proving the adequacy of our denotational semantics with respect to observational equivalence (*cf.* Sect. 4.2). Its proof is a bit technical though; the reader may initially assume this result and skip to the next section, where the more interesting content of the paper starts.

It is evident from its definition that observational equivalence can be split into two halves:

Definition 3.30 (Observational preorder)**.** We define $\mu \preceq \nu$ iff, for every context C, $C[\mu]\Downarrow$ implies $C[\nu]\Downarrow$. Therefore, $\simeq\, =\, \preceq \cap \succeq$.

The same definition can be formulated restricting on tests:

Definition 3.31 (Test preorder)**.** We define $\mu \precsim \nu$ iff, for every test θ, $\theta[\mu]\Downarrow$ implies $\theta[\nu]\Downarrow$.

Lemma 3.32. *$\mu \precsim \nu$ implies that, for every test θ, $\theta[\mu] \precsim \theta[\nu]$.*

Proof. Simply observe that tests compose: if θ_1 is a test for μ and θ_2 is a test for $\theta_1[\mu]$ then $\theta_2[\theta_1[\mu]] = \theta[\mu]$ for a test θ. □

Proving the Context Lemma amounts to showing that $\precsim$ and $\preceq$ coincide. One inclusion is obvious; the non-trivial one is of course $\precsim\, \subseteq\, \preceq$. In what follows, we use the notions of lift and weight of a path, introduced in Sect. 2.1.

Definition 3.33 (Measure of a net)**.** We define the function $\sharp$ from nets to the ordinal $\omega+1$ as follows:

$$\sharp\mu = \inf\{w(\phi^\mu)\ ;\ \phi \in \mathsf{op}^*(\mu)\}.$$

$\sharp\mu$ is said to be the *measure* of μ.

Lemma 3.34. *$\sharp\mu < \omega$ iff $\mu\Downarrow$, and $\sharp\mu = 0$ iff $\mu\downarrow$.*

Proof. The first part is obvious. For the second part, simply observe that $w(\phi) = 0$ iff ϕ is observable. □

Lemma 3.35. *If $0 < \sharp\mu < \omega$, then there exists μ' such that $\mu \to \mu'$ and $\sharp\mu' < \sharp\mu$.*

Proof. From $\sharp\mu < \omega$, using Lemma 3.34 we deduce that $\mu\Downarrow$, which means that μ has a reduct containing an observable path, call it ϕ, and of course we can choose ϕ of minimum weight, *i.e.*, such that $w(\phi^\mu) = \sharp\mu$ (remember that ϕ^μ is the lift of ϕ in μ, *cf.* Definition 2.3). Now, if ϕ^μ were observable, again by Lemma 3.34 we would have $\sharp\mu = 0$, contrarily to our hypothesis. Hence, ϕ^μ is not observable, and it therefore crosses an active pair. If we reduce it, by Lemma 2.7 we obtain a net μ' containing a maximal path ϕ' of weight strictly smaller than ϕ^μ. But, by Proposition 2.4, $\phi' = \phi^{\mu'}$, and since, by Proposition 2.10, $\phi \in \mathsf{op}(\mu')$, we have $\sharp\mu' \leq w(\phi') < w(\phi^\mu) = \sharp\mu$. □

We arrive at last to the two crucial lemmas to prove that $\precsim\, \subseteq\, \preceq$:

Lemma 3.36. *Let $\mu \precsim \nu$. Then, for every wiring ω, $\omega[\mu]\Downarrow$ implies $\omega[\nu]\Downarrow$.*

Proof. (Sketch) Let n be the number of free ports of μ and ν. First of all, observe that if ω has only n free ports, then $\omega[\mu]$ and $\omega[\nu]$ are nets without interface, and the statement holds vacuously. Then, we shall always consider ω to have at least $n+1$ free ports.

By Lemma 3.34, it is enough to show that, for all $k \in \mathbb{N}$, for all nets μ, ν with n free ports, and for all wirings ω with $n+1$ free ports, $\mu \precsim \nu$ and $\sharp\omega[\mu] = k$ imply $\omega[\nu]\Downarrow$. We do this by induction on k. If $k = 0$, then by Lemma 3.34 we have $\omega[\mu]\downarrow$. We can assume without loss of generality that

$\omega[\mu] =$ 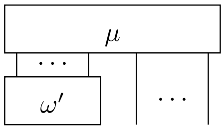

that is, ω connects some of the free ports of μ between them, and acts as the identity on all other ports. We shall actually make an even bigger simplification, which, as the reader may check, will not harm the generality of our arguments: we shall suppose that μ and ν have only 3 free ports, and that ω' is therefore just a wire. Now, if the observable path in $\omega[\mu]$ does not use ω', then $\mu\downarrow$, which implies $\nu\Downarrow$ by $\mu \precsim \nu$. Additionally, one can show the following:

Claim 1. In case $\mu \precsim \nu$, the presence of a path between the free ports i and j in $\mathsf{op}^*(\mu)$ implies the presence of a path between the same ports in $\mathsf{op}^*(\nu)$. □

So we also have $\omega[\nu]\Downarrow$, as desired. Now consider the case in which the observable path in $\omega[\mu]$ passes through ω'. Then, we must have the following situation:

$\omega[\mu] \quad =$ 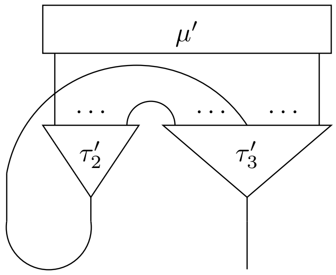

(We have assumed that the tree rooted at the free port 1 of μ is a wire. Of course this need not be the case; however, if it is not a wire, then τ_2' must be. Therefore, our assumption causes no loss of generality). By Claim 1, we have

$\omega[\nu] \;\to^*$ 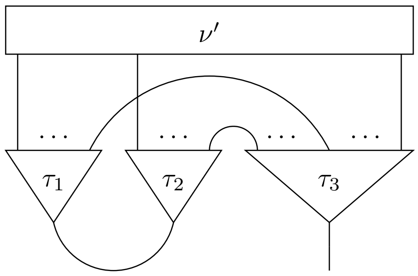

It can be shown that, whatever is the shape of τ_1 and τ_2, they always reduce to a normal net ν_0:

$\omega[\nu] \;\to^*$ 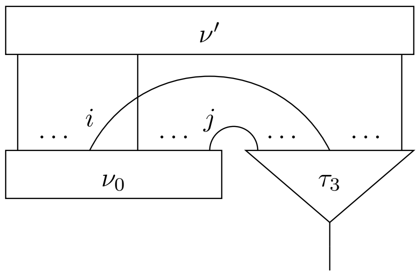

Now suppose $\omega[\nu]\Uparrow$. Then, there can be no observable path between the free ports of ν_0 marked by i and j in the above picture. Consider then the test

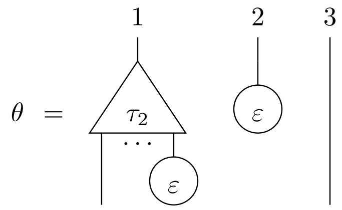

where the leaves of τ_2 are all connected to ε cells except the one which, when τ_2 interacts with τ_1, becomes the free port j in ν_0. It is easy to see that, by what observed above, we would obtain $\theta[\mu]\downarrow$ while $\theta[\nu]\Uparrow$, contradicting the hypothesis that $\mu \precsim \nu$. Therefore, $\omega[\nu]$ must be observable.

We now turn to the case $k > 0$. By Lemma 3.35, we can always find a net μ_0 such that $\omega[\mu] \to \mu_0$ and $\sharp\mu_0 < k$. If the reduction is performed inside μ, then we actually have $\mu_0 = \omega[\mu']$ with $\mu \to \mu'$. It is not hard to show the following:
Claim 2. $\mu \to^* o$ implies $o \precsim \mu$. $\square$

Therefore, we have $\mu' \precsim \mu$, which by the hypothesis that $\mu \precsim \nu$ yields $\mu' \precsim \nu$. But $\sharp\omega[\mu'] < k$, so the induction hypothesis applies, and we have $\omega[\nu]\Downarrow$, as desired. Otherwise, a wire of ω connects two principal ports, and the active pair induced by this connection is the one reduced to obtain μ_0 from $\omega[\mu]$. We have two cases, depending on the active pair: two binary cells with the same or with different symbols. We only check the first case, leaving the other to the reader. We have

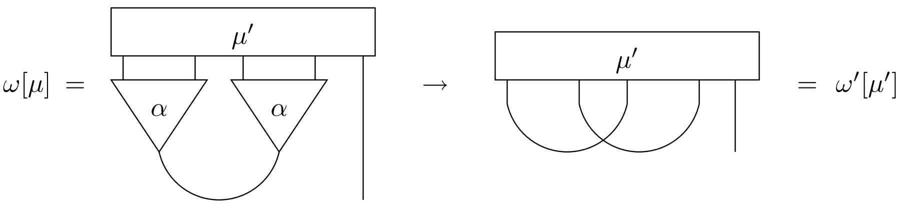

We know that $\sharp\omega'[\mu'] < k$, so if we ever find ν' such that $\mu' \precsim \nu'$, by induction hypothesis we can conclude $\omega'[\nu']\Downarrow$. Define the test

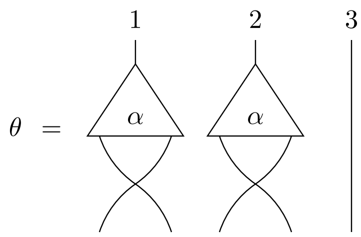

By Lemma 3.32, $\mu \precsim \nu$ implies $\theta[\mu] \precsim \theta[\nu]$. But $\theta[\mu] \to^* \mu'$, so by Claim 2, we have $\mu' \precsim \theta[\mu] \precsim \theta[\nu]$, and thus by induction hypothesis $\omega'[\theta[\nu]]\Downarrow$. But $\omega'[\theta[\nu]] \simeq_\eta \omega[\nu]$, and by Proposition 3.14 $\simeq_\eta \subseteq \simeq$, so $\omega'[\theta[\nu]]\Downarrow$ iff $\omega[\nu]\Downarrow$, and we are done. $\square$

Proposition 3.37. *$\mu \precsim \nu$ implies $\mu \preceq \nu$.*

Proof. By Lemma 3.4, it is enough to show that, for every principal context C, $C[\mu]\Downarrow$ implies $C[\nu]\Downarrow$. So let C be a principal context, and suppose that $C[\mu]\Downarrow$. By definition (*cf.* Fig. 10), we have $C[\mu] = \omega[\theta[\mu]]$ for a suitable wiring ω and test θ. Thanks to Lemma 3.32, we know that $\mu \precsim \nu$ implies $\theta[\mu] \precsim \theta[\nu]$, so by Lemma 3.36 $\omega[\theta[\mu]]\Downarrow$ implies $\omega[\theta[\nu]]\Downarrow$. But $\omega[\theta[\nu]] = C[\nu]$, so $\mu \preceq \nu$. □

4. Full Abstraction

4.1. **Edifices.** We shall now introduce the main mathematical objects of our paper, namely *edifices*. These will be used to develop a denotational semantics for nets, borrowing ideas from the path semantics of linear logic, *i.e.*, Girard's Geometry of Interaction as formulated by Danos and Regnier [DR95]. Although edifices and Böhm trees are technically quite different, there are strong analogies between the two. Also, the topology used to define edifices is the same used by Kennaway et al. to define the infinitary λ-calculus [KKSdV97].

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

Definition 4.1 (Pillar)**.** Given $I \subseteq \mathbb{N}$, set $\mathcal{P}_I = \mathcal{C} \times \mathcal{C} \times I$, equipped with the product topology. A *pillar* is an element of $\mathcal{P} = \mathcal{P}_{\mathbb{N}}$. Pillars are denoted by $u\,@\,i$, and are ranged over by ξ, υ. The pillar $u\,@\,i$ is said to be *based* at i.

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

Definition 4.2 (Arch)**.** Given $I \subseteq \mathbb{N}$, pose $\overrightarrow{\mathcal{A}}_I = \mathcal{P}_I \times \mathcal{P}_I$, equipped with the product topology, and set $(\xi, \upsilon) \sim (\xi', \upsilon')$ iff $\xi' = \upsilon$ and $\upsilon' = \xi$, or $\xi' = \xi$ and $\upsilon' = \upsilon$. We then define $\mathcal{A}_I = \overrightarrow{\mathcal{A}}_I / \sim$, equipped with the quotient topology. An *arch* is an element of $\mathcal{A} = \mathcal{A}_{\mathbb{N}}$. Arches are denoted by $\xi \frown \upsilon$ (which is the same as $\upsilon \frown \xi$), and ranged over by $\mathfrak{a}$; sets of arches are ranged over by $\mathfrak{E}$. An arch is said to be *based* at the unordered pair where its two pillars are based.

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

Proposition 4.3. *The space $\mathcal{A}$ is metrizable; if $\mathfrak{a} = \xi \frown \upsilon$ and $\mathfrak{a}' = \xi' \frown \upsilon'$, the function $D(\mathfrak{a}, \mathfrak{a}') = \min\{\max\{d(\xi, \xi'), d(\upsilon, \upsilon')\}, \max\{d(\xi, \upsilon'), d(\upsilon, \xi')\}\}$ is a distance inducing its topology.* □

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

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

Proposition 4.4. *$\mathfrak{E}$ is compact iff it is a closed subset of $\mathcal{A}_I$ for some finite I.*

Proof. If $\mathfrak{E}$ is compact, then it must be closed; suppose however that $\mathfrak{E} \not\subseteq \mathcal{A}_I$ for all finite I. Then, let $\mathfrak{a}_{i,j}$ be a sequence of arches spanning all of the i, j where the arches of $\mathfrak{E}$ are based, and set $U_{i,j} = \mathfrak{E} \cap \mathcal{B}^\circ(\mathfrak{a}_{i,j}, 2)$. These are all open sets in the relative topology, and since, for all i, j, $D(\mathfrak{a}_{i,j}, \mathfrak{a}) < 2$ iff $\mathfrak{a}$ 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 $\mathfrak{a}_{i,j}$ 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 4.3): when I is finite, there is identity between closed, compact, and complete subsets of $\mathcal{A}_I$.

Definition 4.5 (Edifice)**.** An *edifice* is a compact set of arches.

4.2. Nets as edifices.

The basic idea to assign an edifice to a net is that arches model observable paths.[4] In fact, we have already seen that an observable path is an unordered pair of couples (s, i), where s is an address and i a free port. A pillar contains roughly the same information; the need for infinite words arises from η-expansion (the η_0 equation of Fig. 11), which can be applied indefinitely, as in the pure λ-calculus.

Definition 4.6 (Edifice of an observable path)**.** Let $\phi = \{(s, i), (t, j)\}$ be an observable path. We define

$$\mathfrak{E}(\phi) = \{sw \,@\, i \frown tw \,@\, j \ ; \ \forall w \in \mathcal{C} \times \mathcal{C}\}.$$

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

Proposition 4.7. *Let ϕ be an observable path. Then, $\mathfrak{E}(\phi)$ is an edifice.*

Proof. If $\phi = \{(s, i), (t, j)\}$, clearly $\mathfrak{E}(\phi) \subseteq \mathcal{A}_{\{i,j\}}$. Now take an arch $\mathfrak{a} = u \,@\, i' \frown v \,@\, j'$ not belonging to $\mathfrak{E}(\phi)$. If $i' \neq i$ or $j' \neq j$, then obviously $\mathcal{B}^\circ(\mathfrak{a}, 1)$ is all outside of $\mathfrak{E}(\phi)$. 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 $\mathcal{B}^\circ(\mathfrak{a}, 2^{-(k+1)})$ is all outside of $\mathfrak{E}(\phi)$. So $\mathfrak{E}(\phi)$ is a closed subset of $\mathcal{A}_{\{i,j\}}$, and we conclude by Proposition 4.4. $\square$

[4] Graphically (Fig. 12), observable paths look like arches, hence the terminology.

Observe that $\mathfrak{E}(\phi)$ is not open: given an arch $\mathfrak{a} \in \mathfrak{E}(\phi)$, any open ball centered at $\mathfrak{a}$ contains arches which are not of the form $sw @ i \frown tw @ j$, *i.e.*, the word w completing s and t may not be the same in both pillars.

Definition 4.8 (Edifice of a net)**.** Let μ be a net. The *pre-edifice* of μ is the set

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

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

The soundness of the above definition can be checked as follows: by Proposition 4.7, all of the $\mathfrak{E}(\phi)$ are subsets of $\mathcal{A}_I$ for some finite I; arches based outside of I are "too far" to be adherent to $\mathfrak{E}_0(\mu)$, therefore its closure is still in $\mathcal{A}_I$. By Proposition 4.4, this is enough to ensure the compactness of $\mathfrak{E}(\mu)$.

In some cases, the pre-edifice of a net is already an edifice. This is exactly when the net is $\beta\varepsilon$-normalizable, as we shall now prove. If a is a finite binary word, we denote by $|a|$ its length. If $k \geq |a|$, we denote by $\mathsf{C}_k(a)$ 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 $\phi = \{(a_0 \otimes b_0, i), (c_0 \otimes d_0, j)\}$ be an observable path, and let $k_\phi = \max\{|a|, |b|, |c|, |d|\}$. We define the set of *centers* of ϕ as

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

where x_0 is some fixed infinite word. Then we set

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

The set $\mathcal{O}(\phi)$ is clearly open; additionally, we have the following easy result:

Lemma 4.9. *For every observable path ϕ, $\mathfrak{E}(\phi) \subseteq \mathcal{O}(\phi)$.* □

Lemma 4.10. *Let μ be a net, and let $\phi, \phi' \in \mathsf{op}(\mu)$, with $\phi \neq \phi'$. Then, $\mathcal{O}(\phi) \cap \mathcal{O}(\phi') = \emptyset$.*

Proof. Let $\phi = \{(s, i), (t, j)\}$, $\phi' = \{(s', i'), (t', j')\}$, and let $\phi, \phi' \in \mathsf{op}(\mu)$ for some μ. If $i \neq i'$ or $j \neq j'$, then the result is obvious. Otherwise, observe that we must have $s \neq s'$ and $t \neq t'$, because each leaf of a tree can be involved in at most one observable path. Moreover, s and t cannot be prefixes of s' or t', and vice versa. In fact, if, for example, s were a prefix of s', then s would not be the address of a leaf of μ. Now, the sets $\mathcal{O}(\phi), \mathcal{O}(\phi')$ are built precisely so that, whenever $u @ i \frown v @ j \in \mathfrak{E}(\phi)$ and $u' @ i \frown v' @ j \in \mathfrak{E}(\phi')$, 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. □

In particular, the above result shows that the edifices of two distinct observable paths belonging to the same net are always disjoint.

Proposition 4.11. *For all μ, $\mathfrak{E}_0(\mu)$ is an edifice iff μ is $\beta\varepsilon$-normalizable.*

Proof. The backward implication is a straightforward consequence of Lemma 3.27 (a finite union of compact sets is compact). Suppose now that μ is not $\beta\varepsilon$-normalizable. Again thanks to Lemma 3.27, we know that $\mathsf{op}^*(\mu)$ is infinite. Consider now the family of sets $\mathcal{O}(\phi) \cap \mathfrak{E}_0(\mu)$ as ϕ varies over $\mathsf{op}^*(\mu)$; by Lemma 4.9, this forms an infinite open cover of $\mathfrak{E}_0(\mu)$. By Lemma 4.10, removing any of these sets causes the family not to cover $\mathfrak{E}_0(\mu)$ anymore; hence, $\mathfrak{E}_0(\mu)$ is not compact. □

4.3. **The trace.** We shall now endow pre-edifices with a *trace* operation, which closely corresponds to the execution formula of the Geometry of Interaction [Gir89]. It is also reminiscent of the notion of *play* in games semantics.

In the following, we shall use directed arches, *i.e.*, the elements of $\overrightarrow{\mathcal{A}}$, denoted by $\xi \curvearrowright \upsilon$, where ξ, υ are pillars. Given a set of arches $\mathfrak{E}$, we associate with it a set of directed arches as follows:

$$\overrightarrow{\mathfrak{E}} = \{\xi \curvearrowright \upsilon, \upsilon \curvearrowright \xi \mid \xi \frown \upsilon \in \mathfrak{E}\}.$$

Definition 4.12 (Feedback, trace sequence)**.** A *feedback function* σ is a 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 $\sigma^2(i) = i$. Let σ be a feedback function and $\mathfrak{E} \subseteq \overrightarrow{\mathcal{A}}$. A *trace sequence* of $\mathfrak{E}$ along σ is a sequence $(\mathfrak{s}_n)_{n\in\mathbb{N}}$ of directed arches of $\mathfrak{E}$ such that $\mathfrak{s}_0 = u\,@\,i \curvearrowright \upsilon$ with $\sigma(i)$ undefined, and for all $n \in \mathbb{N}$, $\mathfrak{s}_n = \xi \curvearrowright \upsilon$ and $\mathfrak{s}_{n+1} = \upsilon' \curvearrowright \zeta$ implies $\upsilon = u\,@\,i$ and $\upsilon' = u\,@\,\sigma(i)$, or $\mathfrak{s}_{n+1} = \mathfrak{s}_n$ if $\sigma(i)$ is undefined. The *length* of a trace sequence $\mathfrak{s}$, denoted by $|\mathfrak{s}|$, is the smallest ordinal k such that $\mathfrak{s}_k = \xi \curvearrowright u\,@\,i$ with $\sigma(i)$ undefined. We denote by $\sigma[\mathfrak{E}]$ the set of trace sequences of $\mathfrak{E}$ along σ of finite length. If $\mathfrak{s} \in \sigma[\mathfrak{E}]$ such that $|\mathfrak{s}| = k$, we define the *arch generated by* $\mathfrak{s}$ as $\mathfrak{a}(\mathfrak{s}) = \xi \frown \zeta$, where $\mathfrak{s}_0 = \xi \curvearrowright \upsilon$ and $\mathfrak{s}_k = \upsilon' \curvearrowright \zeta$.

Observe that, if $\mathfrak{E}$ is a set of arches, σ a feedback function, and $\mathfrak{s} \in \sigma[\overrightarrow{\mathfrak{E}}]$ such that $|\mathfrak{s}| = k$, then we also have $\mathfrak{s}' \in \sigma[\overrightarrow{\mathfrak{E}}]$, where $\mathfrak{s}'$ is the "reverse" sequence of $\mathfrak{s}$, *i.e.*, such that, for $n \leq k$, $\mathfrak{s}'_n = \upsilon \curvearrowright \xi$ iff $\mathfrak{s}_{k-n} = \xi \curvearrowright \upsilon$, and, for $k > n$, $\mathfrak{s}'_n = \mathfrak{s}_0$. Note also that $\mathfrak{a}(\mathfrak{s}') = \mathfrak{a}(\mathfrak{s})$.

Definition 4.13 (Trace)**.** Let $\mathfrak{E}$ be a set of arches, and σ a feedback function. We define the *trace* of $\mathfrak{E}$ along σ as

$$\mathrm{Tr}_\sigma(\mathfrak{E}) = \{\mathfrak{a}(\mathfrak{s}) \mid \mathfrak{s} \in \sigma[\overrightarrow{\mathfrak{E}}]\}.$$

Let μ be a net. A *feedback context* for μ is a wiring connecting some of the free ports of μ between them. Note that if μ has n free ports and σ is a feedback function whose domain is included in $\{1, \ldots, n\}$, 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 $\{1, \ldots, n\}$. Hence, we shall use σ to range over both feedback functions and feedback contexts, and make confusion between the two, speaking more generally of a "feedback" σ for a net μ.

The aim of the rest of the section is to prove the following:

Proposition 4.14. *Let μ be a net, and σ a feedback for μ. Then,* $\mathfrak{E}_0(\sigma[\mu]) = \mathrm{Tr}_\sigma(\mathfrak{E}_0(\mu))$. □

We shall prove each inclusion separately, starting from the backward one.

Lemma 4.15. $\mathrm{Tr}_\sigma(\mathfrak{E}_0(\mu)) \subseteq \mathfrak{E}_0(\sigma[\mu])$.

Proof. Let $\mathfrak{s} \in \sigma[\overrightarrow{\mathfrak{E}_0(\mu)}]$. The directed arches appearing in $\mathfrak{s}$ come from arches of $\mathfrak{E}_0(\mu)$, which in turn come from observable paths of $\mathsf{op}^*(\mu)$. We say that an observable path $\phi \in \mathsf{op}^*(\mu)$ is *used* by $\mathfrak{s}$ iff, for some n, $\mathfrak{s}_n \in \overrightarrow{\mathfrak{E}(\phi)}$. Since $\mathfrak{s}$ is finite, only finitely many observable paths are used by it. Hence, by Proposition 2.10, we can take a net μ' such that $\mu \to^* \mu'$ and, whenever ϕ is used by $\mathfrak{s}$, $\phi \in \mathsf{op}(\mu')$. Now $\mathfrak{s}$ induces a sequence $\phi_0, \ldots, \phi_{k-1}$ of observable paths of μ', where $k = |\mathfrak{s}|$; note that the sequence may contain the same path multiple times. The concatenation of $\phi_0, \ldots, \phi_{k-1}$ is a maximal path of $\sigma[\mu']$: in fact, by definition of trace sequence, ϕ_0 and ϕ_{k-1} both have one extremity which is a free port of

$\sigma[\mu']$, while, for all $0 \leq n \leq k-2$, one extremity of ϕ_n is connected by the feedback σ to one extremity of ϕ_{n+1}. We call this maximal path ϕ, and we shall prove that $\mathfrak{a}(\mathfrak{s}) \in \mathfrak{E}_0(\sigma[\mu])$ by induction on $w(\phi)$ (*i.e.*, the weight of ϕ, *cf.* Definition 2.6, Sect. 2.1).

Suppose first that $w(\phi) = 0$. In that case, ϕ crosses no active pair, so it is an observable path of μ', and hence of $\mathsf{op}^*(\mu)$. The situation can be schematically depicted as follows:

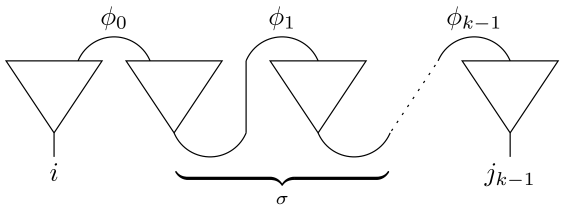

Let $\phi_0 = \{(s,i),(t_0,j_0)\}$ and, for all $n \geq 1$, $\phi_n = \{(\mathbf{1}, i_n),(t_n,j_n)\}$, with $\sigma(j_n) = i_{n+1}$ for all $n \geq 0$. Then, since $\mathfrak{s}$ is a trace sequence, $\mathfrak{s}_0$ must be of the form $sw @ i \curvearrowright t_0 w @ j_0$, $\mathfrak{s}_1$ must be of the form $t_0 w @ i_1 \curvearrowright t_1 t_0 w @ j_1$, $\mathfrak{s}_2$ must be of the form $t_1 t_0 w @ i_2 \curvearrowright t_2 t_1 t_0 w @ j_2$, and so on until $\mathfrak{s}_{k-1}$, which must be of the form

$$t_{k-2} \cdots t_0 w @ i_{k-1} \curvearrowright t_{k-1} \cdots t_0 w @ j_{k-1}.$$

Therefore, we have $\mathfrak{a}(\mathfrak{s}) = sw @ i \frown t_{k-1} \cdots t_0 w @ j_{k-1}$. Now ϕ is precisely the observable path $\{(s,i),(t_{k-1} \cdots t_0, j_{k-1})\}$, so $\mathfrak{a}(\mathfrak{s}) \in \mathfrak{E}_0(\mu)$.

We go on to the inductive case, and suppose $w(\phi) > 0$. By definition, we have the following subnet in μ':

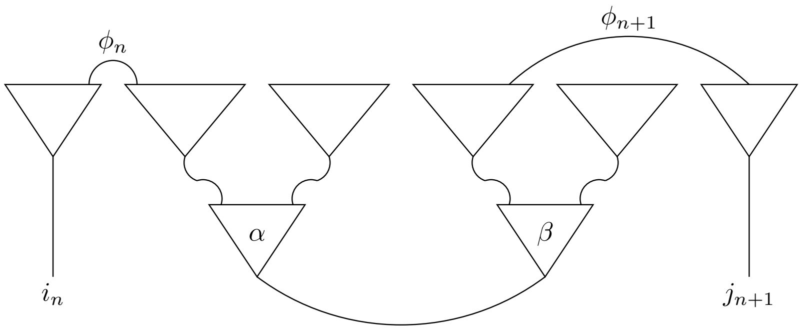

where ϕ_n, ϕ_{n+1} are two consecutive paths in the sequence induced by $\mathfrak{s}$. We made a convenient graphical assumption about the structure of ϕ_n and ϕ_{n+1}, *i.e.*, they both use the "left" auxiliary port of the cells involved in the active pair; we invite the reader to check that our arguments apply for any other choice of the structure of ϕ_n, ϕ_{n+1}. Also, although the picture may suggest so, i_n and j_{n+1}, which are free ports of μ', need not be free in $\sigma[\mu']$. We shall assume that $\alpha \neq \beta$, the case $\alpha = \beta$ being perfectly analogous and left to the reader. So, by taking the paths ϕ_n, ϕ_{n+1} as they are in the above picture, we must have

$$\begin{aligned} \mathfrak{s}_n &= u @ i_n \curvearrowright \mathbf{p}x \otimes \mathbf{p}y @ j_n \\ \mathfrak{s}_{n+1} &= \mathbf{p}x \otimes \mathbf{p}y @ i_{n+1} \curvearrowright v @ j_{n+1}, \end{aligned}$$

with $\sigma(j_n) = i_{n+1}$ and for some $u, v \in \mathcal{C} \times \mathcal{C}$, $x, y \in \mathcal{C}$. After reducing the active pair, we obtain a net μ_1 containing the subnet

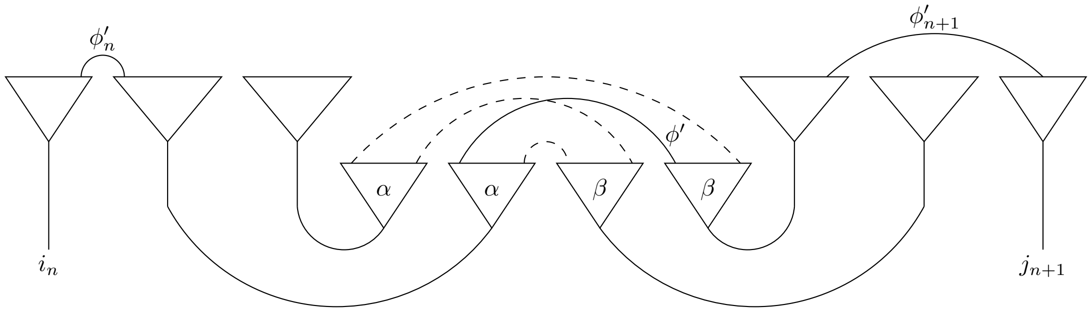

Observe that μ_1 is of the form $\sigma_1[\mu'']$ for a suitable net μ'' and feedback σ_1. The residue ϕ' of ϕ in μ_1 is formed by composing the same paths as ϕ, except for ϕ_n and ϕ_{n+1}, which are now replaced by $\phi'_n, \phi'', \phi'_{n+1}$, where ϕ'' is a new observable path of μ'' shown in the picture as the only solid line. Therefore, we have $\mathfrak{s}' \in \sigma_1[\overrightarrow{\mathfrak{E}_0(\mu'')}]$, where $\mathfrak{s}'_m = \mathfrak{s}_m$ for $m < n$ and $m > n+2$, and

$$\begin{aligned} \mathfrak{s}'_n &= u\,@\,i_n \curvearrowright x \otimes \mathbf{p}y\,@\,j_n \\ \mathfrak{s}'_{n+1} &= x \otimes \mathbf{p}y\,@\,i' \curvearrowright \mathbf{p}x \otimes y\,@\,j' \\ \mathfrak{s}'_{n+2} &= \mathbf{p}x \otimes y\,@\,i_{n+1} \curvearrowright v\,@\,j_{n+1} \end{aligned}$$

where i', j' are two new free ports of μ'' such that $\sigma_1(j_n) = i'$ and $\sigma_1(j') = i_{n+1}$, and we have supposed $\alpha = \zeta$ and $\beta = \delta$. Now, clearly $\mathfrak{a}(\mathfrak{s}') = \mathfrak{a}(\mathfrak{s})$, and since, by Lemma 2.7, $w(\phi') < w(\phi)$, by induction hypothesis we have $\mathfrak{a}(\mathfrak{s}) \in \mathfrak{E}_0(\sigma_1[\mu'']) = \mathfrak{E}_0(\mu_1) = \mathfrak{E}_0(\mu)$, and we are done. □

Lemma 4.16. $\mathfrak{E}_0(\sigma[\mu]) \subseteq \mathrm{Tr}_\sigma(\mathfrak{E}_0(\mu))$.

Proof. Let $\nu = \sigma[\mu]$, and let $\mathfrak{a} \in \mathfrak{E}_0(\nu)$. This means that $\nu \to^* \nu'$ such that $\phi \in \mathsf{op}(\nu')$ and $\mathfrak{a} \in \mathfrak{E}(\phi)$. We shall do the proof by induction on the length of the reduction sequence going from ν to ν'. If this is zero, then $\phi \in \mathsf{op}(\nu)$, which actually implies $\phi \in \mathsf{op}(\mu)$. Then we clearly have $\mathfrak{a} \in \mathrm{Tr}_\sigma(\mathfrak{E}_0(\mu))$, because ϕ connects two ports of μ which are outside of the domain of σ, and, if $\overrightarrow{\mathfrak{a}}$ is any directed version of $\mathfrak{a}$, the constant sequence $\mathfrak{s}_n = \overrightarrow{\mathfrak{a}}$ for all n is a trace sequence of $\sigma[\overrightarrow{\mathfrak{E}_0(\mu)}]$. Suppose now that $\nu \to \nu_1 \to^* \nu'$. If $\nu_1 = \sigma[\mu_1]$ with $\mu \to \mu_1$, then by induction hypothesis $\mathfrak{a} \in \mathrm{Tr}_\sigma(\mathfrak{E}_0(\mu_1))$; but $\mathfrak{E}_0(\mu_1) = \mathfrak{E}_0(\mu)$, so we are done. We may then assume than the active pair reduced in going from ν to ν_1 is not present in μ, and is created by σ. This time we analyze the case of an annihilation step, leaving to the reader the case of a commutation step, which is perfectly similar. We have

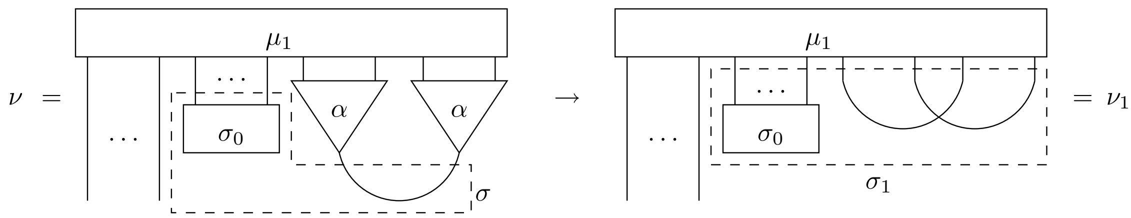

Of course we have $\mathfrak{E}_0(\nu_1) = \mathfrak{E}_0(\nu)$, so $\mathfrak{a} \in \mathfrak{E}_0(\nu_1)$. Then, by applying the induction hypothesis we get $\mathfrak{a} \in \mathrm{Tr}_{\sigma_1}(\mu_1)$, which means that there exists $\mathfrak{s} \in \sigma_1[\overrightarrow{\mathfrak{E}_0(\mu_1)}]$ such that $\mathfrak{a}(\mathfrak{s}) = \mathfrak{a}$. The reduction step leading from ν to ν_1 is necessary to the appearance of the observable path generating $\mathfrak{a}$, otherwise we would not have included it in the reduction sequence; hence,

$\mathfrak{s}$ uses at least one of the two connections created by the reduction step. In other words, we have, for some $n \in \mathbb{N}$,

$$\begin{aligned} \mathfrak{s}_n &= \xi \curvearrowright x \otimes y \,@\, j \\ \mathfrak{s}_{n+1} &= x \otimes y \,@\, j' \curvearrowright \zeta, \end{aligned}$$

with $\sigma_1(j) = j'$, and j, j' are two free ports of μ_1 which are both "left" or "right" auxiliary ports of the α cells forming the active pair in ν. Then, supposing we have $\alpha = \delta$, and supposing j, j' are "left" auxiliary ports, we can build a trace sequence $\mathfrak{s}'$ as follows: $\mathfrak{s}'_m = \mathfrak{s}_m$ for all m in which $\mathfrak{s}$ does not use the connection created by reducing the active pair; and

$$\begin{aligned} \mathfrak{s}'_n &= \xi \curvearrowright \mathbf{p}x \otimes y \,@\, j_0 \\ \mathfrak{s}'_{n+1} &= \mathbf{p}x \otimes y \,@\, j'_0 \curvearrowright \zeta \end{aligned}$$

where $\mathfrak{s}$ uses such connection; j_0, j'_0 are the free ports of μ which are the principal ports of the two α cells involved in the active pair. It is not hard to see that $\mathfrak{s}' \in \sigma[\overrightarrow{\mathfrak{E}_0(\mu)}]$, and that $\mathfrak{a}(\mathfrak{s}') = \mathfrak{a}(\mathfrak{s}) = \mathfrak{a}$. Hence, $\mathfrak{a} \in \mathrm{Tr}_\sigma(\mathfrak{E}_0(\mu))$, as desired. □

4.4. Characterizing total equivalence.

The following result, combined with Proposition 3.28, states that the pre-edifice (and, consequently, the edifice) of a net models $\simeq_{\beta\varepsilon\eta}$.

Proposition 4.17. *For all nets μ, ν, $\mu \cong \nu$ implies $\mathfrak{E}_0(\mu) = \mathfrak{E}_0(\nu)$.*

Proof. Suppose that $\mathfrak{E}_0(\mu) \neq \mathfrak{E}_0(\nu)$, and let $\mathfrak{a} \in \mathfrak{E}_0(\mu) \setminus \mathfrak{E}_0(\nu)$ (we are supposing w.l.o.g. that $\mathfrak{E}_0(\mu)$ is not contained in $\mathfrak{E}_0(\nu)$). We then have

$\mu \to^*$ 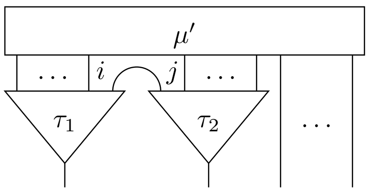

where the observable path shown generates $\mathfrak{a}$, whereas, by Corollary 1.13,

$\nu \simeq_\eta$ 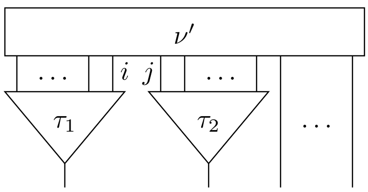

and no reduct of ν' develops a connection between the ports i, j generating $\mathfrak{a}$. Now consider the test

$\theta =$ 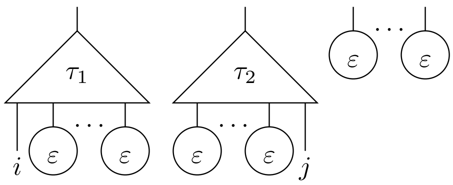

Now $\theta[\mu]$ ε-reduces to a wire, so $\theta[\mu]\Downarrow$. On the contrary, $\theta[\nu]$ reduces to a net with 2 free ports which cannot be $\beta\eta$-equivalent to a wire, otherwise we would have $\mathfrak{a} \in \mathfrak{E}_0(\nu)$. We have two possibilities: either $\theta[\nu]$ is $\beta\varepsilon$-normalizable, or it is not. In the latter case, by Lemma 3.27, 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 3.19, and conclude by applying the Separation Theorem 1.11. □

Corollary 4.18. *For all nets μ, ν, $\mu \simeq_{\beta\varepsilon\eta} \nu$ implies $\mu \cong \nu$, and hence $\mu \simeq \nu$.* □

The interpretation of nets as pre-edifices is a denotational semantics in the sense that it models $\beta\varepsilon\eta$-equivalence, and it induces a congruence, as shown by the following:

Lemma 4.19. *Let μ, ν be two nets such that $\mathfrak{E}_0(\mu) = \mathfrak{E}_0(\nu)$. Then, for every context C, $\mathfrak{E}_0(C[\mu]) = \mathfrak{E}_0(C[\nu])$.*

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 $1, \ldots, n$, 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 $i \in \{1, \ldots, n\}$, $\sigma(i) = i - n$ for $i \in \{n+1, \ldots, 2n\}$, and σ is undefined everywhere else. We clearly obtain $\sigma[C \bullet \mu] = C[\mu]$. Note furthermore that $\mathfrak{E}_0(C \bullet \mu) = \mathfrak{E}_0(C) \cup \mathfrak{E}_0(\mu)$, since the two nets are disjoint and do not share free ports by our assumption. The result is then an easy corollary of Proposition 4.14:

$$\begin{aligned}\mathfrak{E}_0(C[\mu]) &= \mathfrak{E}_0(\sigma[C \bullet \mu]) = \mathrm{Tr}_\sigma(\mathfrak{E}_0(C \bullet \mu)) = \\ &= \mathrm{Tr}_\sigma(\mathfrak{E}_0(C) \cup \mathfrak{E}_0(\mu)) = \mathrm{Tr}_\sigma(\mathfrak{E}_0(C) \cup \mathfrak{E}_0(\nu)) = \\ &= \mathrm{Tr}_\sigma(\mathfrak{E}_0(C \bullet \nu)) = \mathfrak{E}_0(\sigma[C \bullet \mu]) = \mathfrak{E}_0(C[\nu]).\end{aligned}$$

□

We now have our first full abstraction result:

Theorem 4.20 (Full abstraction for $\cong$)**.** *For all nets μ, ν, $\mu \cong \nu$ iff $\mathfrak{E}_0(\mu) = \mathfrak{E}_0(\nu)$.*

Proof. The forward implication is Proposition 4.17. For what concerns the converse, suppose $\mu \ncong \nu$. Then, there exists a context C such that, for example, $C[\mu]\Downarrow$ and $C[\nu]\Uparrow$. By Proposition 4.11, $\mathfrak{E}_0(C[\mu])$ is a non-empty edifice; on the contrary, $\mathfrak{E}_0(C[\nu])$ is either empty or, still by Proposition 4.11, it is not an edifice. In both cases, $\mathfrak{E}_0(C[\mu]) \neq \mathfrak{E}_0(C[\nu])$, which by Lemma 4.19 implies $\mathfrak{E}_0(\mu) \neq \mathfrak{E}_0(\nu)$. □

4.5. **Characterizing observational equivalence.** In the case of observational 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 4.21. *Let μ, ν be two nets with n free ports. Then, $\mathfrak{E}(\mu) \neq \mathfrak{E}(\nu)$ implies that there exist $i, j \in \{1, \ldots, n\}$, two pairs of finite words s, t, and two observable paths $\phi \in \mathsf{op}^*(\mu)$ and $\psi \in \mathsf{op}^*(\nu)$ such that, if we put $\mathfrak{a}_w = sw @ i \frown tw @ j$, either for all w, we have $\mathfrak{a}_w \in \mathfrak{E}(\phi) \setminus \mathfrak{E}(\nu)$, or for all w, we have $\mathfrak{a}_w \in \mathfrak{E}(\psi) \setminus \mathfrak{E}(\mu)$.*

Proof. Suppose, without loss of generality, that there exists $\mathfrak{a} \in \mathfrak{E}(\mu) \setminus \mathfrak{E}(\nu)$, based at $i, j \in \{1, \ldots, n\}$. Remember that $\mathfrak{E}(\mu)$ and $\mathfrak{E}(\nu)$ are defined as the closures of resp. $\mathfrak{E}_0(\mu)$ and $\mathfrak{E}_0(\nu)$, and that by Proposition 4.4 they are both compact, hence complete. Then, if

$\mathfrak{a} \in \mathfrak{E}(\mu) \setminus \mathfrak{E}_0(\mu)$, $\mathfrak{a}$ must be a "missing limit" of a Cauchy sequence $\mathfrak{a}_n \in \mathfrak{E}_0(\mu)$. Since a subsequence of a Cauchy sequence is still a Cauchy sequence, there must exists an integer m such that, for all $n \geq m$, $\mathfrak{a}_n \in \mathfrak{E}_0(\mu) \setminus \mathfrak{E}(\nu)$, otherwise $\mathfrak{a}$ would belong to $\mathfrak{E}(\nu)$ because of its completeness. Therefore, modulo replacing it by one of these $\mathfrak{a}_n$, we can always assume that $\mathfrak{a} \in \mathfrak{E}_0(\mu) \setminus \mathfrak{E}(\nu)$. If it is so, then by definition there exists an observable path $\phi \in \mathsf{op}^*(\mu)$ such that $\mathfrak{a} \in \mathfrak{E}(\phi)$, which means that $\mathfrak{a} = sw_0 \mathbin{@} i \frown tw_0 \mathbin{@} j$ and, for every $w \in \mathcal{C} \times \mathcal{C}$, $sw \mathbin{@} i \frown tw \mathbin{@} j \in \mathfrak{E}(\phi)$, where s and t are the addresses of two leaves in the reduct(s) of μ in which ϕ appears. Now let $s'_1, \ldots, s'_n, \ldots$ 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 $\mathfrak{a}_n = s_n u_n \mathbin{@} i \frown t_n v_n \mathbin{@} j \in \mathfrak{E}(\nu)$; it is not hard to verify that the arches $\mathfrak{a}_n$ would form a Cauchy sequence of limit $\mathfrak{a}$, and thus, by the completeness of $\mathfrak{E}(\nu)$, we would obtain $\mathfrak{a} \in \mathfrak{E}(\nu)$, a contradiction. Therefore, there must exist an integer n such that, for all w, $s_n w \mathbin{@} i \frown t_n w \mathbin{@} j \in \mathfrak{E}(\phi) \setminus \mathfrak{E}(\nu)$. □

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

Lemma 4.22.

(1) *Let τ be a tree, and let*

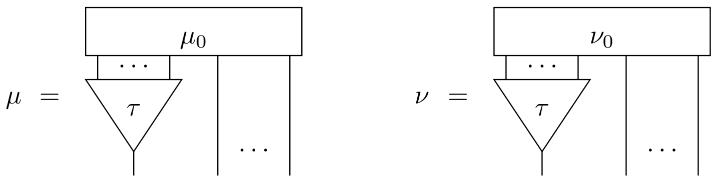

Then, $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$ iff $\mathfrak{E}(\mu_0) = \mathfrak{E}(\nu_0)$.

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

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

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

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

Proof.

(1) Easy.

(2) Simply consider the nets μ'', ν'' obtained from μ', ν' by adding a copy of τ to the one already existing in the two nets, so that each leaf l in one copy is connected to the same leaf l in the other copy. By Lemma 1.12, we have that $\mu'' \simeq_\eta \mu$ and $\nu'' \simeq_\eta \nu$; by point 1, we have $\mathfrak{E}(\mu') = \mathfrak{E}(\nu')$ iff $\mathfrak{E}(\mu'') = \mathfrak{E}(\nu'')$; but by Corollary 4.18, and by hypothesis, $\mathfrak{E}(\mu'') = \mathfrak{E}(\mu) = \mathfrak{E}(\nu) = \mathfrak{E}(\nu'')$.

(3) Call k the free port of μ to which the ε cell is connected in μ'. Observe that such ε cell can either disappear, or be duplicated, and that, in any case, ε cells cannot be used by observable paths. Hence, $\phi \in \mathsf{op}^*(\mu')$ iff $\phi \in \mathsf{op}^*(\mu)$ and ϕ connects two free ports of μ both different than k. Therefore, $\mathfrak{E}(\mu') = \{u \,@\, i \frown u \,@\, j \in \mathfrak{E}(\mu) \;;\; j, k \neq i\}$. The same holds for ν, so from $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$ it easily follows that $\mathfrak{E}(\mu') = \mathfrak{E}(\nu')$.

□

Corollary 4.23. *Let μ, ν be two nets with the same interface, and let θ be a test. Then, $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$ implies $\mathfrak{E}(\theta[\mu]) = \mathfrak{E}(\theta[\nu])$.* □

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

Lemma 4.24. *Let W be a net with two free ports connected by a wire, and let μ be a net with two free ports, such that $\phi \in \mathsf{op}^*(\mu)$ implies that ϕ does not connect the port 1 to the port 2. Then, there exists a test θ such that $\theta[W]\Downarrow$ and $\theta[\mu]\Uparrow$.*

Proof. If $\mu\Uparrow$, the identity test 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 $\mu \to^* \mu'$, and

$\mu' =$ 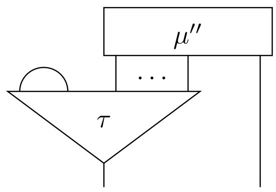

In the above picture, we have supposed that the observable path connects the free port 1 to itself, and that the leaves connected in the path are the two "leftmost" leaves of τ. These are just graphically convenient assumptions, causing no loss of generality: the observable path may as well connect port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

$\theta =$ 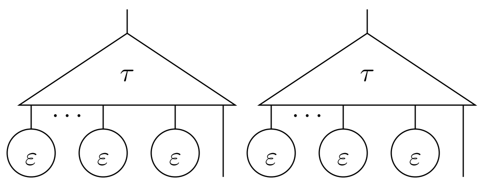

we have that, thanks to Lemma 1.4, $\theta[W] \to^* W$, while $\theta[\mu]$ reduces to a net whose free port 1 is connected to an ε cell. If this net is blind, we are done; otherwise, there is a reduct of $\theta[\mu]$ containing an observable path between the free port 2 and itself. This observable 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 $\theta'[W] \to^* W\downarrow$, while $\theta'[\mu]\Uparrow$, as desired. □

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

Theorem 4.25 (Full abstraction for $\simeq$)**.** *For all nets μ, ν, $\mu \simeq \nu$ iff $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$.*

Proof. Consider first the backward implication (also known as the adequacy property). We start by observing that, for any net o, $o\Downarrow$ iff $\mathsf{op}^*(o) \neq \emptyset$ iff $\mathfrak{E}(o) \neq \emptyset$. Now, suppose $\mathfrak{E}(\mu) = \mathfrak{E}(\nu)$, and let θ be a test. By Corollary 4.23, we have $\mathfrak{E}(\theta[\mu]) = \mathfrak{E}(\theta[\nu])$, so following the above remark $\theta[\mu]\Downarrow$ iff $\mathfrak{E}(\theta[\mu]) \neq \emptyset$ iff $\mathfrak{E}(\theta[\nu]) \neq \emptyset$ iff $\theta[\nu]\Downarrow$. Then $\mu \simeq \nu$ follows from the Context Lemma 3.29.

Now we turn to the actual full abstraction property. For this, we consider the contrapositive statement, and assume $\mathfrak{E}(\mu) \neq \mathfrak{E}(\nu)$. Let I be the interface of μ and ν. By Lemma 4.21, we know that there exist $i, j \in I$, $\phi \in \mathsf{op}^*(\mu)$, and two leaves in a reduct of μ of addresses s, t such that, for all w, $sw @ i \frown tw @ j \in \mathfrak{E}(\phi) \setminus \mathfrak{E}(\nu)$ (it could actually be that these arches belong to $\mathfrak{E}(\psi) \setminus \mathfrak{E}(\mu)$, where $\psi \in \mathsf{op}^*(\nu)$, but obviously our assumption causes no loss of generality). We shall suppose $i \neq j$; the reader is invited to check that the argument can be adapted to the case $i = j$. By Definition 4.6, and by the fact that $\phi \in \mathsf{op}^*(\mu)$, we have

$\mu \to^*$ 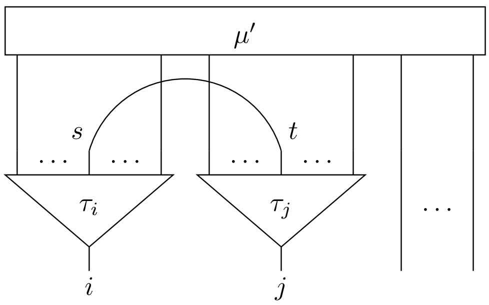

where we have explicitly drawn the connection between the two leaves of resp. addresses s and t. On the other hand, by Corollary 1.13, we have

$\nu \simeq_\eta$ 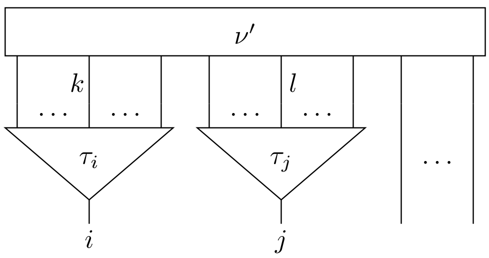

where we have called k and l the two free ports of ν' corresponding resp. to the addresses t and s in τ_i and τ_j. Observe that, by Corollary 4.18, the edifice of the net on the right is still $\mathfrak{E}(\nu)$. Now if, in any reduct of ν', there appeared an observable path between k and l, then we would contradict the fact that, for all w, $sw @ i \frown tw @ j \notin \mathfrak{E}(\nu)$. Therefore, no observable path ever appears between k and l in any reduct of ν'.

Consider then the test

$\theta =$ 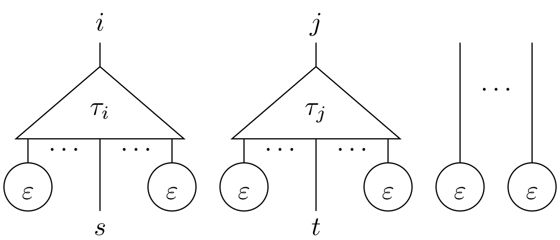

where we have left free only the leaves corresponding to the addresses s and t of τ_i and τ_j. Now, by Lemma 1.4, $\theta[\mu] \to^* W$, where W is a wire plus a net with no interface; on the other hand, we have

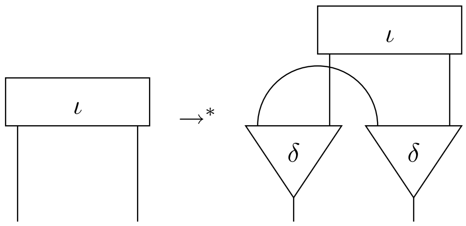

Figure 13: A non-$\beta\varepsilon$-normalizable net observationally equivalent to a wire.

$\theta[\nu] \simeq_{\beta\eta}$

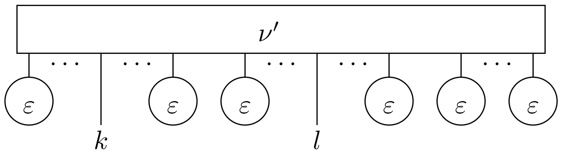

But ν' never develops observable paths between k and l, so Lemma 4.24 applies, and we obtain $\mu \not\simeq \nu$. □

4.6. **Examples.** As an immediate application of Theorem 4.25, we give an example of 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 $\lambda z.z$. It is an example proving that the inclusion $\cong \subsetneq \simeq$ is strict, as anticipated right after Definition 3.2.

Consider a net ι containing no observable paths, and reducing as in Fig. 13. Such a net exists by what we have shown in Sect. 1.3. We see that $\phi \in \mathsf{op}^*(\iota)$ iff

$$\mathfrak{E}(\phi) = \{\mathbf{q}^n\mathbf{p}x \otimes y \,@\, 1 \frown \mathbf{q}^n\mathbf{p}x \otimes y \,@\, 2 \ ; \ \forall x, y \in \mathcal{C}\}$$

for some non-negative integer n. On the other hand, if W denotes a wire,

$$\mathfrak{E}(W) = \mathfrak{E}_0(W) = \{u \,@\, 1 \frown u \,@\, 2 \ ; \ \forall u \in \mathcal{C} \times \mathcal{C}\}.$$

Now, if $\mathbf{q}^\infty$ denotes an infinite sequence of $\mathbf{q}$'s, all arches of the form

$$\mathfrak{a}_y = \mathbf{q}^\infty \otimes y \,@\, 1 \frown \mathbf{q}^\infty \otimes y \,@\, 2$$

are missing from $\mathfrak{E}_0(\iota)$, hence $\mathfrak{E}_0(\iota) \subsetneq \mathfrak{E}_0(W)$. But these arches are all adherent to $\mathfrak{E}_0(\iota)$: in fact, it is very easy to construct a Cauchy sequence in $\mathfrak{E}_0(\iota)$ of limit $\mathfrak{a}_y$, for any y. Therefore, $\mathfrak{E}(\iota) = \mathfrak{E}(W)$, and $\iota \simeq W$. On the other hand, $\iota \not\cong W$, and we do not need Theorem 4.20 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.

One my wonder whether the converse of Proposition 3.28 holds, *i.e.*, whether total equivalence reduces to $\beta\varepsilon\eta$-equivalence. This would be close in spirit to what happens in the λ-calculus: nf-equivalence, which we said is akin to total equivalence, exactly corresponds to $\beta\eta\Omega$-equivalence, which is morally analogous to our $\beta\varepsilon\eta$-equivalence. However, such analogy fails: the two nets given in Fig. 14 give an interesting example of this. They can be built by slightly twisting the constructions given in Sect. 1.3. It is not hard to show that $\mu_1 \not\simeq_{\beta\varepsilon\eta} \mu_2$;

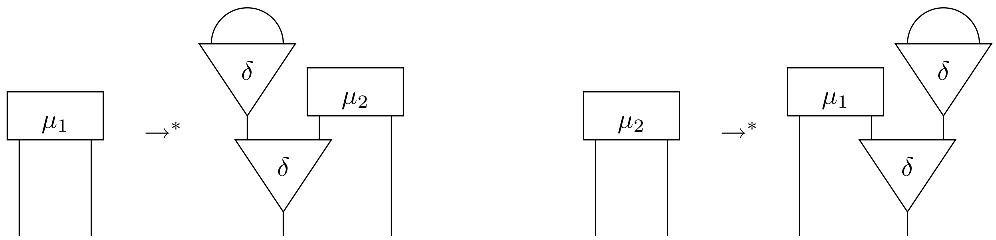

Figure 14: Nets showing that $\mu \cong \nu$ does not imply $\mu \simeq_{\beta\varepsilon\eta} \nu$.

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 paths, *i.e.*, $\mathsf{op}^*(\mu_1) = \mathsf{op}^*(\mu_2)$. Therefore, $\mathfrak{E}_0(\mu_1) = \mathfrak{E}_0(\mu_2)$, and $\mu_1 \cong \mu_2$ by Theorem 4.20.

References

[Bar84] Henk P. Barendregt. *The Lambda Calculus.* North Holland, revised edition, 1984.

[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.

[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.

[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.

[Gir96] Jean-Yves Girard. Proof-nets: The parallel syntax for proof-theory. In Ursini and Agliano, editors, *Logic and Algebra.* Marcel Dekker, Inc., 1996.

[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.

[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.

[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.

[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.

[Wad76] Christopher Wadsworth. The relation between computational and denotational properties for Scott's D_∞ models. *Siam J. Comput.*, 5(3):488–521, 1976.