

Fully Abstract Encodings of λ -Calculus in HOcore through Abstract Machines

Malgorzata Biernacka, Dariusz Biernacki, Sergueï Lenglet, Piotr Polesiuk,
Damien Pous, Alan Schmitt

► **To cite this version:**

Malgorzata Biernacka, Dariusz Biernacki, Sergueï Lenglet, Piotr Polesiuk, Damien Pous, et al.. Fully Abstract Encodings of λ -Calculus in HOcore through Abstract Machines. LICS 2017, Jun 2017, Reykjavik, Iceland. Proceedings of LICS 2017. <hal-01479035>

HAL Id: hal-01479035

<https://hal.inria.fr/hal-01479035>

Submitted on 28 Feb 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fully Abstract Encodings of λ -Calculus in HOcore through Abstract Machines

Małgorzata Biernacka^{*}, Dariusz Biernacki^{*}, Sergueï Lenglet[†], Piotr Polesiuk^{*}, Damien Pous[‡] and Alan Schmitt[◊]
^{*} University of Wrocław [†] Université de Lorraine [‡] Université de Lyon, CNRS, ENS Lyon, Inria [◊] Inria

Abstract—We present fully abstract encodings of the call-by-name λ -calculus into HOcore, a minimal higher-order process calculus with no name restriction. We consider several equivalences on the λ -calculus side—normal-form bisimilarity, applicative bisimilarity, and contextual equivalence—that we internalize into abstract machines in order to prove full abstraction.

I. INTRODUCTION

HOcore is a minimal process calculus with higher-order communication, meaning that messages are executable processes. It is a subcalculus of HO π [20] with no construct to generate names or to restrict the scope of communication channels. Even with such a limited syntax, HOcore is Turing complete [15]. However, as a higher-order calculus, it is less expressive than the name passing π -calculus: polyadic message sending cannot be compositionally encoded in monadic message sending in HO π [14], while it can be done in π [22].

Although HOcore is Turing complete, we initially thought a fully abstract encoding of the λ -calculus into HOcore was impossible. Indeed, a λ -term potentially has an unbounded number of redexes. A straightforward encoding would use communication to emulate β -reduction, but since HOcore does not provide means to restrict the scope of communication, one would need as many distinct names as there are redexes to avoid interference. Moreover, as new redexes may be created by β -reduction, we also need a way to generate new names on which to communicate. To circumvent these problems and illustrate the expressiveness of HOcore, we consider encodings where the *reduction strategy* is fixed, thus for which at most one redex is enabled at any time. In this setting, β -reduction can be emulated using communication on a single, shared, name. A first contribution of this paper is a novel encoding of the call-by-name λ -calculus, or more precisely the Krivine Abstract Machine (KAM) [13], into HOcore.

A faithful encoding not only reflects the operational semantics of a calculus, it should also reflect its equivalences. Ideally, an encoding is *fully abstract*: two source terms are behaviorally equivalent iff their translations are equivalent. On the HOcore side, we use *barbed equivalence with hidden names* [15], where a fixed number of names used for the translation cannot be observed. On the λ -calculus side, we consider three equivalences. First, we look at *normal-form bisimilarity* [16, 17], where normal forms are decomposed into subterms that must be bisimilar. Next, we turn to *applicative bisimilarity* [1], where normal forms must behave similarly when applied to identical arguments. And finally, we consider contextual equivalence, where terms must behave identically

when put into arbitrary contexts. Our second contribution is an *internalization* of these equivalences into extended abstract machines: these machines expand the KAM, which performs the evaluation of a term, with additional transitions with flags interpreting the equivalences. By doing so, we can express these different equivalences on terms by a simpler bisimilarity on these flags-generating machines, which can be seen as labeled transition systems. Finally, and as third contribution, we translate these extended machines into HOcore and prove full abstraction for all three equivalences. Altogether, this work shows that a minimal process calculus with no name restriction can faithfully encode the call-by-name λ -calculus.

The chosen equivalences: One may wonder why we study normal-form and applicative bisimilarities, as faithfully encoding contextual equivalence is enough to get full abstraction. Our motivation is twofold. First, we start with normal-form bisimilarity because it is the simplest to translate. Indeed, there is no need to inject terms from the environment to establish the equivalence. We next show how we can inject terms for applicative bisimilarity, and we then extend this approach to contexts for contextual equivalence. Second, the study of quite different equivalences illustrate the robustness of the internalization technique.

Related work: Since Milner’s seminal work [19], other translations of the λ -calculus or one of its variants into π -calculus have been proposed, e.g., to study connections with logic [2, 5, 23], termination [9, 4, 25], sequentiality [6], control [9, 12, 24], or Continuation-Passing Style (CPS) transforms [21, 22, 10]. These works use the more expressive *first-order* π -calculus, except for [21, 22], discussed below; full abstraction is proved w.r.t. contextual equivalence in [6, 25, 12], normal-form bisimilarity in [24], and normal-form and applicative bisimilarities in [22]. The definitions of the encodings and the equivalences of [6, 25, 12] are driven by types, and therefore cannot be compared to our untyped setting. In [24], van Bakel et al. establish a full abstraction result between the $\lambda\mu$ -calculus with normal-form bisimilarity and the π -calculus. Their encoding relies on an unbounded number of restricted names to evaluate several translations of λ -terms in parallel, while we rely on flags and on barbed equivalence to know which translated λ -term is being evaluated. We explain the differences between the two approaches in Section IV-B.

Sangiorgi translates the λ -calculus into a higher-order calculus as an intermediary step in [21, 22], but it is an abstraction-passing calculus, which is strictly more expressive than a process-passing calculus [14]. Like in our work, Sangiorgi

fixes the evaluation strategy in the λ -calculus, except that he uses CPS translations rather than abstract machines. In the light of Danvy et al.’s functional correspondence [3], the two approaches appear closely related, however it is difficult to compare our encoding with Sangiorgi’s, since we use different target calculi, and we internalize the normal-form and applicative bisimilarities in the abstract machines. Still, name restriction plays an important role in Sangiorgi’s encodings, since a local channel is used for each application in a λ -term. Full abstraction w.r.t. normal-form bisimilarity (called “applicative open bisimilarity”) is established in [22, Chapter 18]. In [22, Chapter 17], the encoding is proved to be sound but not complete w.r.t. applicative bisimilarity: there exist applicative bisimilar λ -terms whose translations are not bisimilar in π -calculus. It is then argued that completeness can be achieved for this encoding only by extending the λ -calculus with non-confluent constructs. We explain the discrepancy with our results in Remark 24.

Outline: Section II presents the syntax and semantics of HOcore. Section III shows how to encode the KAM; the machine and its encoding are then modified to get full abstraction with relation to normal-form bisimilarity (Section IV) and applicative bisimilarity (Section V). Section VI concludes this paper, and the appendix contains the main proofs.

II. THE CALCULUS HOCORE

Syntax and semantics: HOcore [15] is a simpler version of HO π [20] where name restriction is removed. We let $a, b, \text{etc.}$ range over channel names, and $x, y, \text{etc.}$ range over process variables. The syntax of HOcore processes is:

$$P, Q ::= a(x).P \mid \bar{a}\langle P \rangle \mid P \parallel Q \mid x \mid \mathbf{0}.$$

The process $a(x).P$ is waiting for a message on a which, when received, is substituted for the variable x in P . If x does not occur in P , then we write $a(_).P$. The process $\bar{a}\langle P \rangle$ is sending a message on a . Note that communication is higher order—processes are sent—and asynchronous—there is no continuation after a message output. The parallel composition of processes is written $P \parallel Q$, and the process $\mathbf{0}$ cannot perform any action. Input has a higher precedence than parallel composition, e.g., we write $a(x).P \parallel Q$ for $(a(x).P) \parallel Q$. We implicitly consider that parallel composition is associative, commutative, and has $\mathbf{0}$ as a neutral element. In an input $a(x).P$, the variable x is bound in P . We write $\text{fn}(P)$ for the channel names of P .

Informally, when an output $\bar{a}\langle P \rangle$ is in parallel with an input $a(x).Q$, a communication on a takes place, producing $[P/x]Q$, the capture avoiding substitution of x by P in Q . We define in Figure 1 the semantics of HOcore as a labeled transition system (LTS), omitting the rules symmetric to PAR and TAU. The labels (ranged over by l) are either τ for internal communication, $\bar{a}\langle P \rangle$ for message output, or $a(P)$ for process input. We use label a for an input where the received process does not matter (e.g., $a(_).P$).

Weak transitions allow for internal actions before and after a visible one. We write $\xRightarrow{\tau}$ for the reflexive and transitive closure $\xrightarrow{\tau}^*$, and \xRightarrow{l} for $\xrightarrow{l} \xrightarrow{\tau} \xrightarrow{l} \xrightarrow{\tau}$ when $l \neq \tau$.




Fig. 1: HOcore LTS

Barbed equivalence: We let γ range over names a and $\text{conames } \bar{a}$; we define observable actions as follows.

Definition 1. *The process P has a strong observable action on a (resp. \bar{a}), written $P \downarrow_a$ (resp. $P \downarrow_{\bar{a}}$), if $P \xrightarrow{a(Q)} R$ (resp. $P \xrightarrow{\bar{a}\langle Q \rangle} R$) for some Q, R . A process P has a weak observable action on γ , written $P \Downarrow_{\gamma}$, if $P \xRightarrow{\tau} P' \downarrow_{\gamma}$ for some P' . We write $\text{WkObs}(P)$ for the set of weak observable actions of P .*

Our definition of barbed equivalence depends on a set \mathbb{H} of hidden names, which allows some observable actions to be ignored. Instead of adding top-level name restrictions on these names, as in [15], we prefer to preserve the semantics of the calculus and simply hide some names in the equivalence. Hidden names are not a computational construct and are not required for the encoding of the KAM, but they are necessary to protect the encoding from an arbitrary environment when proving full abstraction. We emphasize that we do not need the full power of name restriction: the set of hidden names is finite and static—there is no way to create new hidden names.

Definition 2. *A symmetric relation \mathcal{R} is a barbed bisimulation w.r.t. \mathbb{H} if $P \mathcal{R} Q$ implies*

- $P \downarrow_{\gamma}$ and $\gamma \notin \mathbb{H}$ implies $Q \downarrow_{\gamma}$;
- for all R such that $\text{fn}(R) \cap \mathbb{H} = \emptyset$, we have $P \parallel R \mathcal{R} Q \parallel R$;
- if $P \xrightarrow{\tau} P'$, then there exists Q' such that $Q \xRightarrow{\tau} Q'$ and $P' \mathcal{R} Q'$.

Barbed equivalence w.r.t. \mathbb{H} , noted $\approx_{\text{HO}}^{\mathbb{H}}$, is the largest barbed bisimulation w.r.t. \mathbb{H} .

A strong barbed equivalence can be defined by replacing \Downarrow_{γ} with \downarrow_{γ} in the first item, and $\xRightarrow{\tau}$ with $\xrightarrow{\tau}$ in the third. From [15], we know that strong barbed equivalence is decidable when $\mathbb{H} = \emptyset$, but undecidable when \mathbb{H} is of cardinal at least 4. We lower this bound to 2 in Theorem 4.

III. ENCODING THE KRIVINE ABSTRACT MACHINE

We show in this section that HOcore may faithfully encode a call-by-name λ -calculus through an operationally equivalent encoding of the KAM.

A. Definition of the KAM

The KAM [13] is a machine for call-by-name evaluation of closed λ -calculus terms. We present a substitution-based variant of the KAM for simplicity, and to reuse the substitution

of HOCORE in the translation. A *configuration* C of the machine is composed of the term t being evaluated, and a stack π of λ -terms. Their syntax and the transitions are as follows.

$$\begin{aligned}
C &::= t * \pi && \text{(configurations)} \\
t, s &::= x \mid tt \mid \lambda x.t && \text{(terms)} \\
\pi &::= t :: \pi \mid [] && \text{(stacks)} \\
t * \pi &\mapsto t * s :: \pi && \text{(PUSH)} \\
\lambda x.t * s :: \pi &\mapsto [s/x]t * \pi && \text{(GRAB)}
\end{aligned}$$

A λ -abstraction $\lambda x.t$ binds x in t ; a term is closed if it does not contain any free variables. We use $[]$ to denote the empty stack. In PUSH, the argument s of an application is stored on the stack while the term t in function position is evaluated. If we get a λ -abstraction $\lambda x.t$, then an argument s is fetched from the stack (transition GRAB), and the evaluation continues with $[s/x]t$. If a configuration of the form $\lambda x.t * []$ is reached, then the evaluation is finished, and the result is $\lambda x.t$. Because we evaluate closed terms only, it is not possible to obtain a configuration of the form $x * \pi$.

B. Translation into HOCORE

The translation of the KAM depends essentially on how we push and grab terms on the stack. We represent the stack by two messages, one on name hd_c for its head, and one on name c (for *continuation*) for its tail (henceforth, a stack n is always encoded as a message on hd_n for its head and one on n for its tail). The empty stack can be represented by an arbitrary, non-diverging, deterministic process, e.g., $\mathbf{0}$; here we use a third name to signal that the computation is finished with $\bar{b}(\mathbf{0})$. As an example, the stack $1 :: 2 :: 3 :: 4 :: []$ is represented by $\overline{hd_c}(1) \parallel \bar{c}(\overline{hd_c}(2) \parallel \bar{c}(\overline{hd_c}(3) \parallel \bar{c}(\overline{hd_c}(4) \parallel \bar{c}(\bar{b}(\mathbf{0}))))))$.

With this representation, pushing an element e on a stack p is done by creating the process $\overline{hd_c}(e) \parallel \bar{c}(p)$, while grabbing the head of the stack corresponds to receiving on hd_c . With this idea in mind, we define the translations for the stacks, terms, and configurations as follows, where we assume the variable p does not occur in the translated entities.

$$\begin{aligned}
\llbracket t * \pi \rrbracket &\hat{=} \llbracket t \rrbracket \parallel \bar{c}(\llbracket \pi \rrbracket) \\
\llbracket [] \rrbracket &\hat{=} \bar{b}(\mathbf{0}) \\
\llbracket t :: \pi \rrbracket &\hat{=} \overline{hd_c}(\llbracket t \rrbracket) \parallel \bar{c}(\llbracket \pi \rrbracket) \\
\llbracket t s \rrbracket &\hat{=} c(p).(\llbracket t \rrbracket \parallel \bar{c}(\overline{hd_c}(\llbracket s \rrbracket) \parallel \bar{c}(p))) \\
\llbracket \lambda x.t \rrbracket &\hat{=} c(p).(hd_c(x). \llbracket t \rrbracket \parallel p) \\
\llbracket x \rrbracket &\hat{=} x
\end{aligned}$$

In the translation of a configuration $t * \pi$, the stack is in a message on c , meaning that before pushing on π or grabbing the head of π , we have to get $\llbracket \pi \rrbracket$ by receiving on c . For instance, in the application case $\llbracket t s \rrbracket$, we start by receiving the current stack p on c , and we then run $\llbracket t \rrbracket$ in parallel with the translation of the new stack $\overline{hd_c}(\llbracket s \rrbracket) \parallel \bar{c}(p)$. Similarly, in the λ -abstraction case $\llbracket \lambda x.t \rrbracket$, we get the current stack p on c , that we run in parallel with $hd_c(x). \llbracket t \rrbracket$. If p is not empty, then it is a process of the form $\overline{hd_c}(\llbracket s \rrbracket) \parallel \bar{c}(\llbracket \pi \rrbracket)$, and a communication

on hd_c is possible, realizing the substitution of x by s in t ; the execution then continues with $\llbracket [s/x]t \rrbracket \parallel \bar{c}(\llbracket \pi \rrbracket)$. Otherwise, p is $\bar{b}(\mathbf{0})$, and the computation terminates.

Formally, the operational correspondence between the KAM and its translation is as follows.

Theorem 3. *In the forward direction, if $C \mapsto^* C'$, then $\llbracket C \rrbracket \xrightarrow{\tau} \llbracket C' \rrbracket$. In the backward direction, if $\llbracket C \rrbracket \xrightarrow{\tau} P$, then there exists a C' such that $C \mapsto^* C'$ and either*

- $P = \llbracket C' \rrbracket$,
- or there exists P' such that $P \xrightarrow{\tau} P' = \llbracket C' \rrbracket$,
- or $C' = \lambda x.t * []$ and $P = hd_c(x). \llbracket t \rrbracket \parallel \bar{b}(\mathbf{0})$.

Sketch. The proof is straightforward in the forward direction. In the backward direction, we show that the translation is deterministic (if $\llbracket C \rrbracket \xrightarrow{\tau} P \xrightarrow{\tau} Q_1$ and $\llbracket C \rrbracket \xrightarrow{\tau} P \xrightarrow{\tau} Q_2$, then $Q_1 = Q_2$) and we rely on the fact that the translation of a PUSH step uses one communication, while we use two communications for a GRAB step. \square

We can then improve over the result of [15] about undecidability of strong barbed equivalence by hiding hd_c and c .

Theorem 4. *Strong barbed equivalence is undecidable in HOCORE with 2 hidden names.*

Proof. Assume we can decide strong barbed congruence with two hidden names, and let t be a closed λ -term. We can thus decide if $\llbracket t * [] \rrbracket$ is strong barbed-congruent to $c(x).(x \parallel \bar{c}(x)) \parallel \bar{c}(c(x).(x \parallel \bar{c}(x)))$ when hd_c and c are hidden. As the second term loops with no barbs, deciding congruence is equivalent to deciding whether the reduction of t converges, hence a contradiction. \square

IV. NORMAL-FORM BISIMILARITY

Our first full abstraction result is for normal-form bisimilarity [16]. We show how to internalize this equivalence in an extension of the KAM such that it may be captured by a simple barbed bisimilarity. We then translate this extended KAM into HOCORE, and we finally prove full abstraction.

A. Normal-Form Bisimilarity

Normal-form bisimilarity compares terms by reducing them to weak head normal forms, if they converge, and then decomposes these normal forms into subterms that must be bisimilar. Unlike the KAM, normal-form bisimilarity is defined on open terms, thus we distinguish free variables, ranged over by α , from bound variables, ranged over by x . The grammars of terms (t, s) and values (v) become as follows.

$$t ::= \alpha \mid x \mid \lambda x.t \mid tt \quad v ::= \alpha \mid \lambda x.t$$

Henceforth, we assume that λ -terms are well formed, i.e., all variables ranged over by x are bound: x is not a valid term but α is. We write $\text{fv}(t)$ for the set of free variables of t . A variable α is said *fresh* if it does not occur in any entities under consideration.

When evaluating an open term, we can obtain either a λ -abstraction, or a free variable in a stack. We inductively extend a relation \mathcal{R} on λ -terms to stacks by writing $\pi \mathcal{R} \pi'$ if $\pi = \pi' = []$, or if $\pi_1 = t :: \pi'_1$, $\pi_2 = s :: \pi'_2$, $t \mathcal{R} s$, and $\pi'_1 \mathcal{R} \pi'_2$.

Definition 5. A symmetric relation \mathcal{R} is a normal-form bisimulation if $t \mathcal{R} s$ implies:

- if $t \star [] \mapsto^* \lambda x.t' \star []$, then there exists s' such that $s \star [] \mapsto^* \lambda x.s' \star []$ and $[\alpha / x]t' \mathcal{R} [\alpha / x]s'$ for a fresh α ;
- if $t \star [] \mapsto^* \alpha \star \pi$, then there exists π' such that $s \star [] \mapsto^* \alpha \star \pi'$ and $\pi \mathcal{R} \pi'$.

Normal-form bisimilarity \approx_{nf} is the largest normal-form bisimulation.

Normal-form bisimilarity is not complete w.r.t. contextual equivalence in λ -calculus, but it characterizes Lévy-Longo tree equivalence [16].

B. Abstract Machine

We now describe how to extend the KAM so that it provides additional steps, identified by labeled transitions, that capture normal-form bisimilarity. These extended machines feature *flagged transitions*, *terminating transitions*, and a restricted form of *non-determinism*. Flagged transitions are usual transitions of the machine with some additional information to convey to the environment that a particular event is taking place. Machine bisimilarity, defined below, ensures that bisimilar machines have matching flags. Transitions without flags use a τ label. Terminating transitions are flagged transitions that indicate the computation has stopped. They are needed for bisimilarity: as machine bisimilarity ignores τ labels, we use terminating transitions to distinguish between terminated and divergent machine runs. Finally, we allow non-determinism in machines, i.e., a given configuration may take two different transitions, only if the transitions are flagged and have different flags. In other words, the non-deterministic choice is made explicit to the environment.

We define the NFB machine in Figure 2. When computation stops with a λ -abstraction and an empty stack, we have to restart the machine to evaluate the body of the abstraction with a freshly generated free variable (rule LAMBDA). To do so, we consider free variables as natural numbers, and we keep a counter n in the machine which is incremented each time a fresh variable is needed. For a configuration $\alpha \star \pi$, normal-form bisimilarity evaluates each of the t_i in the stack (rule VAR). To internalize this step, we could launch several machines in parallel, as in [24], where the translated t_i are run in parallel. This approach has two drawbacks: first, it is a further extension of abstract machines (a machine no longer steps to a single machine state but to a multiset of states). Second, when translating such extended machines into HOcore, we want to prevent them from interacting with each other, but we cannot rely on name restriction, as in [24], to encapsulate an unbounded number of translations. Alternatively, one could evaluate the elements of the stack sequentially, but this approach fails if one of the elements of the stack diverges, as the later elements will never be evaluated. We thus consider a third approach, built upon flagged non-determinism: the machine chooses arbitrarily an element of the stack to evaluate, and signals this choice using flags (rules ENTER, SKIP, and DONE). The burden of evaluating every element of the stack is thus relegated to the definition of machine bisimilarity: as




Fig. 2: NFB Machine

every flagged execution must be matched by an execution with the same flags, every possible choice is explored.

As before, we use C to range over configurations, which are now of two kinds. In *evaluation mode*, $\langle t, \pi, n \rangle_{\text{ev}}$ is reducing t within stack π and with counter n . The transitions PUSH and GRAB are as in the KAM, except for the extra parameter. If we reach a λ -abstraction in the empty context (transition LAMBDA), then the machine flags \circledast and then restarts to evaluate the body, replacing the bound variable by a fresh free variable, i.e., the current value n of the counter. If we reach a free variable α , i.e., a number, then we flag the value of α before entering the next mode (transition VAR).

In *continuation mode* $\langle \pi, n \rangle_{\text{ct}}$, the transition DONE simply finishes the execution if $\pi = []$, using the flag \star . Otherwise, $\pi = t :: \pi'$, and the machine either evaluates t with flag \blacktriangledown (and forgets about π'), or skips t with a flag \blacktriangleright to eventually evaluate a term in π' . The machine may skip the evaluation of all the terms in π , but it would still provide some information, as it would generate $m \blacktriangleright$ messages (followed by \star), telling us that π has m elements. Note that the counter n is stored in continuation mode just to be passed to the evaluation mode when one of the t_i is chosen with transition ENTER.

Example 6. To illustrate how the machine works, we show the transitions starting from the term $(\lambda x.x) (\lambda y.y \ 0 \ \Omega)$, where $\Omega \triangleq (\lambda x.x \ x) (\lambda x.x \ x)$. The term is executed in the empty context, and with a counter initialized to a value greater than its free variables.

$$\begin{aligned}
 & \langle (\lambda x.x) (\lambda y.y \ 0 \ \Omega), [], 1 \rangle_{\text{ev}} \\
 & \xrightarrow{\tau} \langle \lambda x.x, \lambda y.y \ 0 \ \Omega :: [], 1 \rangle_{\text{ev}} && \text{(PUSH)} \\
 & \xrightarrow{\tau} \langle \lambda y.y \ 0 \ \Omega, [], 1 \rangle_{\text{ev}} && \text{(GRAB)} \\
 & \xrightarrow{\circledast} \langle 1 \ 0 \ \Omega, [], 2 \rangle_{\text{ev}} && \text{(LAMBDA)} \\
 & \xrightarrow{\tau} \langle 1 \ 0, \Omega :: [], 2 \rangle_{\text{ev}} \xrightarrow{\tau} \langle 1, 0 :: \Omega :: [], 2 \rangle_{\text{ev}} && \text{(PUSH - PUSH)} \\
 & \xrightarrow{1} \langle 0 :: \Omega :: [], 2 \rangle_{\text{ct}} && \text{(VAR)}
 \end{aligned}$$

We then have three possibilities. First, we reduce the top of the stack, with the sequence $\langle 0 :: \Omega :: [], 2 \rangle_{\text{ct}} \xrightarrow{\blacktriangledown} \langle 0, [], 2 \rangle_{\text{ev}} \xrightarrow{0} \langle [], 2 \rangle_{\text{ct}} \xrightarrow{\star}$. Second, we evaluate Ω with the sequence $\langle 0 :: \Omega :: [], 2 \rangle_{\text{ct}} \xrightarrow{\blacktriangleright} \langle \Omega :: [], 2 \rangle_{\text{ct}} \xrightarrow{\blacktriangledown} \langle \Omega, [], 2 \rangle_{\text{ev}}$, and then the machine loops without generating any flag. Third, we skip both terms with $\langle 0 :: \Omega :: [], 2 \rangle_{\text{ct}} \xrightarrow{\blacktriangleright} \langle \Omega :: [], 2 \rangle_{\text{ct}} \xrightarrow{\blacktriangleright} \langle [], 2 \rangle_{\text{ct}} \xrightarrow{\star}$. Note that the three options generate different traces of flags.

Because the rules GRAB and PUSH are the same between the KAM and the NFB machine, there is a direct correspondence between the two.

Lemma 7. *For all $t, t', \pi, \pi', n, t \star \pi \mapsto t' \star \pi'$ iff $\langle t, \pi, n \rangle_{\text{ev}} \xrightarrow{\tau} \langle t', \pi', n \rangle_{\text{ev}}$.*

We finally show that a notion of bisimilarity between configurations of an NFB machine captures normal-form bisimilarity. To this end, we first define machine bisimilarity, where we denote the flags of the machine by f .

Definition 8. *A symmetric relation \mathcal{R} is a machine bisimulation if $C_1 \mathcal{R} C_2$ implies:*

- if $C_1 \xrightarrow{\tau} \xrightarrow{*} \xrightarrow{f} C'_1$, then there exists C'_2 such that $C_2 \xrightarrow{\tau} \xrightarrow{*} \xrightarrow{f} C'_2$ and $C'_1 \mathcal{R} C'_2$;
- if $C_1 \xrightarrow{\tau} \xrightarrow{*} \xrightarrow{f}$, then $C_2 \xrightarrow{\tau} \xrightarrow{*} \xrightarrow{f}$.

Machine bisimilarity \approx_m is the largest machine bisimulation.

Intuitively, machine bisimilarity ensures that every flag emitted by a machine is matched by an identical flag from the other machine, up to internal reductions. Note that a machine that diverges with τ labels can be related to any other diverging machine or any machine stuck without a flag. We make sure the latter case cannot occur in our machines by only having terminating transitions, which are flagged, as stuck transitions. We can now state that normal-form bisimilarity coincides with machine bisimilarity of NFB machines.

Theorem 9. *We have $t \approx_{\text{nf}} s$ iff there exists $n > \max(\text{fv}(t) \cup \text{fv}(s))$ such that $\langle t, [], n \rangle_{\text{ev}} \approx_m \langle s, [], n \rangle_{\text{ev}}$.*

Sketch. Machine bisimilarity implies \approx_{nf} because

$$\{(t, s) \mid \langle t, [], n \rangle_{\text{ev}} \approx_m \langle s, [], n \rangle_{\text{ev}}, n > \max(\text{fv}(t) \cup \text{fv}(s))\}$$

is a normal-form bisimulation, and the other direction is by showing that

$$\{(\langle t, [], n \rangle_{\text{ev}}, \langle s, [], n \rangle_{\text{ev}}) \mid t \approx_{\text{nf}} s, n > \max(\text{fv}(t), \text{fv}(s))\} \\ \cup \{(\langle \pi, n \rangle_{\text{ct}}, \langle \pi', n \rangle_{\text{ct}}) \mid \pi \approx_{\text{nf}} \pi', n > \max(\text{fv}(\pi), \text{fv}(\pi'))\}$$

is a machine bisimulation. \square

C. Translation into HOcore

In Figure 3, we present the translation of the NFB machine into HOcore, where we consider flags as channel names. Configurations now contain a counter n , which is represented by a message on k containing the value of n encoded as a process. We use $[\cdot]_{\text{Int}}$ to translate a natural number n into a process $\underbrace{\text{succ}(_)\dots\text{succ}(_)}_{n \text{ times}}.z(_).\text{init}(\mathbf{0})$; the role of

the final output on init is explained later. Free variables α are also numbers, since we cannot generate new names in HOcore, and we use the same translation for them. We also use non-deterministic internal choice, encoded as follows: $P + Q \triangleq \overline{ch}(P) \parallel \overline{ch}(Q) \parallel ch(x).ch(_).x$: both messages are consumed, and only one process is executed. This encoding supposes that at most one choice is active at a given time, as we use only one name ch to encode all the choices. We also use n -ary choices for $n > 2$ in Section V-C, which can be encoded in the same way.

$$\begin{aligned} \llbracket t s \rrbracket &\triangleq c(p).(\llbracket t \rrbracket \parallel \overline{c}\langle \overline{hd}_c(\llbracket s \rrbracket) \parallel \overline{c}(p) \rangle) \\ \llbracket \lambda x.t \rrbracket &\triangleq c(p).(p \parallel \overline{b}\langle \text{Restart} \rangle \parallel hd_c(x).b(_).\llbracket t \rrbracket) \\ \llbracket x \rrbracket &\triangleq x \\ \llbracket \alpha \rrbracket &\triangleq [\alpha]_{\text{Int}} \\ \text{Restart} &\triangleq \odot(_).k(x).\left(\overline{hd}_c\langle x \rangle \parallel \overline{k}\langle \text{succ}(_).x \rangle \parallel \overline{c}\langle [\cdot] \rangle \parallel \overline{b}\langle \mathbf{0} \rangle\right) \\ \text{Rec} &\triangleq \text{init}(_).\text{rec}(x).(x \parallel \overline{rec}\langle x \rangle \parallel \text{Cont}) \\ \text{Cont} &\triangleq c(p).(p \parallel \overline{b}\langle \star(_).\mathbf{0} \rangle \parallel hd_c(x).b(_).\text{Chce}(x)) \\ \text{Chce}(P_t) &\triangleq \blacktriangledown(_).c(_).(P_t \parallel \overline{c}\langle [\cdot] \rangle) + \blacktriangleright(_).\overline{\text{init}}(\mathbf{0}) \\ [\cdot] &\triangleq b(x).x \\ \llbracket t :: \pi \rrbracket &\triangleq \overline{hd}_c\langle \llbracket t \rrbracket \rangle \parallel \overline{c}\langle \llbracket \pi \rrbracket \rangle \\ [\mathbf{0}]_{\text{Int}} &\triangleq z(_).\overline{\text{init}}(\mathbf{0}) \\ \llbracket n + 1 \rrbracket_{\text{Int}} &\triangleq \text{succ}(_).\llbracket n \rrbracket_{\text{Int}} \\ \llbracket \langle t, \pi, n \rangle_{\text{ev}} \rrbracket &\triangleq \llbracket t \rrbracket \parallel \overline{c}\langle \llbracket \pi \rrbracket \rangle \parallel \overline{k}\langle \llbracket n \rrbracket_{\text{Int}} \rangle \parallel \text{Rec} \parallel \overline{rec}\langle \text{Rec} \rangle \\ \llbracket \langle \pi, n \rangle_{\text{ct}} \rrbracket &\triangleq \overline{c}\langle \llbracket \pi \rrbracket \rangle \parallel \overline{k}\langle \llbracket n \rrbracket_{\text{Int}} \rangle \parallel \text{Cont} \parallel \text{Rec} \parallel \overline{rec}\langle \text{Rec} \rangle \end{aligned}$$

Fig. 3: Translation of the NFB machine into HOcore

A stack is represented as in the KAM, by messages on hd_c and c , and the translation of an application $\llbracket t s \rrbracket$ is exactly the same as for the KAM. The encoding of the empty context $[\cdot]$ is different, however, because contexts are used to distinguish between execution paths at two points in the machine: when evaluating a function $\lambda x.t$ in evaluation mode, and when deciding whether the execution is finished in continuation mode. The empty context is thus encoded as $b(x).x$, waiting to receive the process to execute in the empty case. For the non-empty case, this input on b is absent and there are instead messages on hd_c and c . Thus the generic way to choose a branch is as follows:

$\overline{b}\langle \text{do this if empty} \rangle \parallel hd_c(x).c(y).b(_).\text{do this if non-empty}$.
In the non-empty case, the input on b discards the message for the empty behavior that was not used.

For λ -abstractions, the behavior for the empty case is described in the process Restart. More precisely, $\llbracket \lambda x.t \rrbracket$ receives the current stack $[\pi]$ on c to run it in parallel with $\overline{b}\langle \text{Restart} \rangle \parallel hd_c(x).b(_).\llbracket t \rrbracket$. If $[\pi]$ is of the form $\overline{hd}_c\langle \llbracket t' \rrbracket \rangle \parallel \overline{c}\langle \llbracket \pi' \rrbracket \rangle$, then we have the same behavior as with the KAM, with an extra communication on b to garbage collect the Restart process. Otherwise, $[\pi] = b(x).x$ and we obtain the following sequence of transitions.

$$\begin{aligned} &\xrightarrow{\tau} b(x).x \parallel \overline{b}\langle \text{Restart} \rangle \parallel hd_c(x).b(_).\llbracket t \rrbracket \parallel \overline{k}\langle \llbracket n \rrbracket_{\text{Int}} \rangle \\ &\xrightarrow{\odot} \odot(_).k(x).(\overline{hd}_c\langle x \rangle \parallel \overline{k}\langle \text{succ}(_).x \rangle \parallel \overline{c}\langle [\cdot] \rangle \parallel \overline{b}\langle \mathbf{0} \rangle) \\ &\quad \parallel hd_c(x).b(_).\llbracket t \rrbracket \parallel \overline{k}\langle \llbracket n \rrbracket_{\text{Int}} \rangle \\ &\xrightarrow{\text{Rec}} \overline{hd}_c\langle \llbracket n \rrbracket_{\text{Int}} \rangle \parallel \overline{k}\langle \text{succ}(_).\llbracket n \rrbracket_{\text{Int}} \rangle \parallel \overline{c}\langle [\cdot] \rangle \parallel \overline{b}\langle \mathbf{0} \rangle \\ &\quad \parallel hd_c(x).b(_).\llbracket t \rrbracket \\ &\xrightarrow{\tau} \overline{k}\langle \text{succ}(_).\llbracket n \rrbracket_{\text{Int}} \rangle \parallel \overline{c}\langle [\cdot] \rangle \parallel \overline{b}\langle \mathbf{0} \rangle \parallel b(_).\llbracket [n/x]t \rrbracket \\ &\xrightarrow{\tau} \overline{k}\langle \llbracket n + 1 \rrbracket_{\text{Int}} \rangle \parallel \overline{c}\langle [\cdot] \rangle \parallel \llbracket [n/x]t \rrbracket \end{aligned}$$

In the end, we have effectively restarted the machine to evaluate $[n/x]t$, as wished.

In continuation mode, the branching is done by the process Cont , which is executed after applying the transition VAR . More precisely, a free variable α is translated using $[\![\cdot]\!]_{\text{Int}}$, which signals first the value of α (with the names suc and z), and then sends a message on init to enter the continuation mode. The way the NFB machine chooses which t_i to evaluate in a stack $t_1 :: \dots :: t_m :: []$ is a recursive mechanism, and recursion can be encoded in a higher-order calculus: $\text{Rec} \parallel \overline{\text{rec}}(\text{Rec})$ reduces to $\text{Cont} \parallel \text{Rec} \parallel \overline{\text{rec}}(\text{Rec})$ when it receives a message on init . The process Cont is doing a case analysis on $[\![\pi]\!]$ when it is executed in parallel with $\overline{c}(\![\pi]\!)$: if $\pi = []$, then $[\![\pi]\!] = b(x).x$ receives the message on b which flags \star and the machine stops. Otherwise, $[\![\pi]\!] = \overline{hd}_c(\![t]\!) \parallel \overline{c}(\![\pi']\!)$, and we have the following reductions:

$$\begin{aligned} \text{Cont} \parallel \overline{c}(\![\pi]\!) &\xrightarrow{\tau} \overline{hd}_c(\![t]\!) \parallel \overline{c}(\![\pi']\!) \parallel \overline{b}(\star(_).0) \\ &\quad \parallel \overline{hd}_c(x).b(_).\text{Chce}(x) \\ &\xrightarrow{\tau}^2 \overline{c}(\![\pi']\!) \parallel \text{Chce}(\![t]\!) \end{aligned}$$

At this point, $\text{Chce}(\![t]\!)$ either evaluates t with flag \blacktriangledown , or flags \blacktriangleright and continues exploring π' . In the former case, the current stack π' is replaced by an empty stack, and in the latter, a message on init is issued to produce $\text{Cont} \parallel \overline{c}(\![\pi']\!)$ after some reduction steps.

D. Operational Correspondence and Full Abstraction

Establishing full abstraction requires first to state the correspondence between the NFB machine and its translation. In what follows, we let f range over the flags \odot , \blacktriangledown , \blacktriangleright , \star , and α , where α is used as a shorthand for the succession of α suc flags followed by a z flag. We let \hat{f} range over flags and τ .

Definition 10. A process P is a machine process if there exists a configuration C of the machine such that $[\![C]\!] \xrightarrow{\tau} P$.

To establish the correspondence between the machine and its translation, we first define a predicate $\text{next}(\hat{f}, P)$ such that, if $Q \in \text{next}(\hat{f}, P)$, then there is a weak reduction with a step \hat{f} from P to Q , Q is the translation of a machine, and there is no machine translation in between. Intuitively, it is the first translation reachable after a step \hat{f} .

Definition 11. We write $Q \in \text{next}(\hat{f}, P)$ if $P \xrightarrow{\tau} \xrightarrow{\hat{f}} \xrightarrow{\tau} Q$, $Q = [\![C']\!]$ for some C' , and for any P' such that $P' \neq P$, $P' \neq Q$, and $P \xrightarrow{\hat{f}} P' \xrightarrow{\tau} Q$ or $P \xrightarrow{\tau} P' \xrightarrow{\hat{f}} Q$, we have $P' \neq [\![C']\!]$ for any C' .

A machine process is either a translation of a configuration, or an intermediary state between two translations; $\text{next}(\hat{f}, P)$ gives the set of next processes which are translations. We now prove that, with a single exception, the translation is deterministic. The exception to determinism corresponds to the choice made in continuation mode, which is also not deterministic (but flagged) in the machine. We call *choice process* a process about to make that choice; such processes are of the form $\text{Chce}(P_t, P_\pi, P_n) \hat{=} \text{Chce}(P_t) \parallel \overline{c}(P_\pi) \parallel \overline{k}(P_n) \parallel \text{Rec} \parallel \overline{\text{rec}}(\text{Rec})$.

Lemma 12. Let P be a machine process which is not a choice process. If $P \xrightarrow{\hat{f}} P'$ and $P \xrightarrow{\hat{f}} P''$, then $P' = P''$. Let P be a choice process. If $P \xrightarrow{\tau} P'$, $P \xrightarrow{\tau} P''$, and $\text{WkObs}(P') = \text{WkObs}(P'')$, then $P' = P''$.

The choice process $\text{Chce}(P_t, P_\pi, P_n)$ may only reduce to the translation of a configuration after a flag \blacktriangleright or \blacktriangledown . Thus $\text{next}(\tau, \text{Chce}(P_t, P_\pi, P_n))$ is empty, and Lemma 12 implies that $\text{next}(\hat{f}, P)$ is a singleton if it is not empty. In the following, we write $\text{next}(\hat{f}, P)$ to assert that it is not empty and to directly denote the corresponding unique machine translation.

We can now state the correspondence between the NFB machine and its translation.

Lemma 13. The following assertions hold:

- $C \xrightarrow{\hat{f}} C'$ iff $[\![C]\!] \xrightarrow{\hat{f}} [\![C']\!]$ and $\text{next}(\hat{f}, [\![C]\!]) = [\![C']\!]$.
- $C \xrightarrow{\star} P$ iff $[\![C]\!] \xrightarrow{\tau} \xrightarrow{\star} P$ and $P \approx_{\text{HO}} \mathbf{0}$.

With this result, we can relate the bisimilarities of each calculus. Henceforth, we write \approx_{HO} for $\approx_{\text{HO}}^{\mathbb{H}}$ where \mathbb{H} contains the names of the translation that are not flags. Given a flag $f \neq \alpha$, we write \overline{f} for the process $\overline{f}(\mathbf{0})$. We also write $\overline{\alpha}$ for the process complementing α , defined as $\overline{\alpha} \hat{=} \overline{n+1} \hat{=} \overline{\text{suc}}(\mathbf{0}) \parallel \overline{n}$ and $\overline{0} \hat{=} \overline{z}(\mathbf{0})$.

Theorem 14. $C \approx_m C'$ iff $[\![C]\!] \approx_{\text{HO}} [\![C']\!]$.

Sketch. To prove that barbed equivalence implies machine equivalence, we show that $\mathcal{R} \hat{=} \{(C, C') \mid [\![C]\!] \approx_{\text{HO}} [\![C']\!]\}$ is a machine bisimulation. Let $C_1 \mathcal{R} C'_1$, so that $C_1 \xrightarrow{\tau}^* C_2 \xrightarrow{f} C_3$; then $[\![C_1]\!] \parallel \overline{f} \xrightarrow{\tau} [\![C_2]\!] \parallel \overline{f} \xrightarrow{\tau} [\![C_3]\!]$ by Lemma 13, which in turn implies that there exists P such that $[\![C'_1]\!] \parallel \overline{f} \xrightarrow{\tau} P$ and $[\![C_3]\!] \approx_{\text{HO}} P$. In particular, we have $\text{WkObs}([\![C_3]\!]) = \text{WkObs}(P)$, meaning that $\neg(P \downarrow_{\overline{f}})$. Consequently, there exists P' such that $[\![C'_1]\!] \xrightarrow{\tau} P'$ and $P' \downarrow_{\overline{f}}$. We can prove that $P \approx_{\text{HO}} \text{next}(f, P')$, but by definition of next , there exists C'_3 so that $\text{next}(f, P') = [\![C'_3]\!]$. As a result, we have $[\![C_3]\!] \approx_{\text{HO}} P \approx_{\text{HO}} [\![C'_3]\!]$, i.e., $C_3 \mathcal{R} C'_3$, and $[\![C'_1]\!] \xrightarrow{f} [\![C'_3]\!]$, which implies $C'_1 \xrightarrow{\tau}^* \xrightarrow{f} C'_3$, as wished. The case $C_1 \xrightarrow{\tau}^* C_2 \xrightarrow{\star} P$ is similar.

For the reverse implication, we prove that

$$\mathcal{R} \hat{=} \left\{ \begin{array}{l} (P_C \parallel R, P_{C'} \parallel R) \mid C \approx_m C', \\ \text{WkObs}(P_C) = \text{WkObs}(P_{C'}) \\ P_C \xrightarrow{\tau} [\![C]\!], P_{C'} \xrightarrow{\tau} [\![C']\!] \text{ or} \\ [\![C]\!] \xrightarrow{\tau} P_C, [\![C']\!] \xrightarrow{\tau} P_{C'}, \end{array} \right\} \cup \{(P \parallel R, P' \parallel R) \mid P \approx_{\text{HO}} \mathbf{0} \approx_{\text{HO}} P'\}$$

is a barbed bisimulation. The difficult part is to check that $P_C \parallel R \xrightarrow{\tau} P'$ with a communication on a flag f is matched by $P_{C'} \parallel R$. We have $P_C \downarrow_{\overline{f}}$, which implies that P_C cannot reduce: a machine process is either reducing or has an observable action. We are therefore in the case $[\![C]\!] \xrightarrow{\tau} P_C$, and $R \downarrow_{\overline{f}}$. Suppose $f \neq \star$; then $P_C \xrightarrow{f} P_{C_2} \xrightarrow{\tau} [\![C_2]\!]$ for some P_{C_2} and with $[\![C_2]\!] = \text{next}(f, P_C)$. Because a machine process is deterministic (Lemma 12), then in fact

$P' = P_{C_2} \parallel R'$ for some R' . We have $\llbracket C \rrbracket \xrightarrow{f} \llbracket C_2 \rrbracket$, so $C \xrightarrow{\tau^* f} C_2$ holds by Lemma 13. Because $C \approx_m C'$, there exists C'_2 such that $C' \xrightarrow{\tau^* f} C'_2$ and $C_2 \approx_m C'_2$. By Lemma 13, this implies $\llbracket C' \rrbracket \xrightarrow{f} \llbracket C'_2 \rrbracket$; in particular, there exists $P_{C'_2}$ such that $\llbracket C' \rrbracket \xrightarrow{\tau} \xrightarrow{f} P_{C'_2} \xrightarrow{\tau} \llbracket C'_2 \rrbracket$. By Lemma 12, we have $P_{C'} \xrightarrow{\tau} \xrightarrow{f} P_{C'_2}$, therefore we have $P_{C'} \parallel R \xrightarrow{\tau} P_{C'_2} \parallel R'$, with $P_{C_2} \parallel R' \mathcal{R} P_{C'_2} \parallel R'$, as wished. In the case $f = \star$, we show that we obtain processes in the second set of \mathcal{R} . \square

As a result, we can deduce full abstraction between HOcore and the λ -calculus with normal-form bisimilarity.

Corollary 15. *We have $t \approx_{nf} s$ iff there exists $n > \max(\text{fv}(t) \cup \text{fv}(s))$ such that $\llbracket \langle t, [], n \rangle_{ev} \rrbracket \approx_{HO} \llbracket \langle s, [], n \rangle_{ev} \rrbracket$.*

Proof. By Theorems 9 and 14. \square

V. APPLICATIVE BISIMILARITY

Proving full abstraction w.r.t. normal-form bisimilarity requires minimal interactions—synchronizations on flags—between a machine process and the outside. Achieving full abstraction w.r.t. applicative bisimilarity is intuitively more difficult, since this bisimilarity tests λ -abstractions by applying them to an arbitrary argument. Internalizing such bisimilarity is simple using higher-order flags: one may think of the following transition to test the result of a computation:

$$\lambda x.t \star [] \xrightarrow{s} [s/x]t \star []$$

Although HOcore has higher-order communications, we cannot use them to obtain a fully abstract encoding of such a machine for two reasons. First, allowing interactions where the environment provides a term may allow arbitrary processes to be received, including processes that are not in the image of the translation, thus potentially breaking invariants of the translation. Second, the translation of the KAM has to hide the names it uses for the translation to be fully abstract; it is thus impossible for the context to use such names and to provide translated λ -terms to be tested.

We thus propose in this section to internalize applicative bisimilarity using ordinary flags: when the abstract machine reaches a value, it switches to a different mode where it non-deterministically builds a test term step by step, using flags to indicate its choices so as to ensure that a bisimilar machine builds the same term. The translation of such a machine into HOcore is then similar to the translation of the NFB machine.

Using simple flags to generate terms step by step implies we need to deal with binders. In particular, and anticipating on the HOcore translation, we no longer can rely on the definition of binding and substitution from HOcore, as we cannot write a process that inputs a translation of t and outputs a translation of $\lambda x.t$ using an HOcore binding for x . We thus switch to a pure data description of bindings, using de Bruijn indices. As such terms still need to be executed, we first recall the definition of the KAM with de Bruijn indices and the definitions of contextual equivalence and applicative bisimilarity for λ -terms with de Bruijn indices. We then present the machine internalizing applicative bisimilarity, its

translation into HOcore, and show they are fully abstract. We finally conclude this section by showing how contextual equivalence is internalized in an abstract machine, generating contexts instead of terms.

A. The KAM and Behavioral Equivalences

In the λ -calculus with de Bruijn indices, a variable is a natural number, which indicates which encompassing λ is its binder. For example, $\lambda x.x$ is written $\lambda.0$ and $\lambda xy.x y$ is written $\lambda.\lambda.1 0$. The syntax of terms (t, s) , closures (η) , environments (e, d) , and values (v) is as follows.

$$t ::= n \mid t s \mid \lambda.t \quad \eta ::= (t, e) \quad e ::= \eta :: e \mid \epsilon \quad v ::= (\lambda.t, e)$$

A closure η is a pair (t, e) where e is an environment mapping the free variables of t to closures; environments are used in lieu of substitutions. A term t is closed if $\text{fv}(t) = \emptyset$, and a closure (t, e) is closed if the number of elements of e is bigger than the highest free variable of t .

The semantics is given by the original, environment-based KAM, where a configuration C is now composed of the closed closure (t, e) being evaluated, and a stack π of closures. The transitions rules are as follows.

$$\begin{aligned} C ::= \langle t, e, \pi \rangle_{ev} \text{ (configurations)} \quad \pi ::= \eta :: \pi \mid [] \text{ (stacks)} \\ \langle t s, e, \pi \rangle_{ev} &\xrightarrow{\tau} \langle t, e, (s, e) :: \pi \rangle_{ev} && \text{(PUSH)} \\ \langle 0, (t, e) :: d, \pi \rangle_{ev} &\xrightarrow{\tau} \langle t, e, \pi \rangle_{ev} && \text{(ZERO)} \\ \langle n + 1, (t, e) :: d, \pi \rangle_{ev} &\xrightarrow{\tau} \langle n, d, \pi \rangle_{ev} && \text{(ENV)} \\ \langle \lambda.t, e, \eta :: \pi \rangle_{ev} &\xrightarrow{\tau} \langle t, \eta :: e, \pi \rangle_{ev} && \text{(GRAB)} \end{aligned}$$

In PUSH, the argument s of an application is stored on the stack with its environment e while the term t in function position is evaluated. If we get a λ -abstraction $\lambda.t$ (transition GRAB), then an argument η is moved from the stack to the top of the environment to remember that η corresponds to the de Bruijn index 0, and the evaluation continues with t . Looking up the closure corresponding to a de Bruijn index in the environment is done with the rules ENV and ZERO. Because we evaluate closed closures only, it is not possible to obtain a configuration of the form $\langle n + 1, \epsilon, \pi \rangle_{ev}$. If a configuration of the form $\langle \lambda.t, e, [] \rangle_{ev}$ is reached, then the evaluation is finished, and the result is $(\lambda.t, e)$.

Behavioral equivalences: Contextual equivalence compares closed terms by testing them within all contexts. A context \mathcal{C} is a term with a hole \square at a variable position; plugging a term t in \mathcal{C} is written $\mathcal{C}[t]$. A context is closed if $\text{fv}(\mathcal{C}) = \emptyset$. Contextual equivalence is then defined as follows.

Definition 16. *Two closed terms t and s are contextually equivalent, written $t \approx_{ctx} s$, if for all closed contexts \mathcal{C} , $\langle \mathcal{C}[t], \epsilon, [] \rangle_{ev} \xrightarrow{\tau^*} \langle \lambda.t', e', [] \rangle_{ev}$ for some t' and e' iff $\langle \mathcal{C}[s], \epsilon, [] \rangle_{ev} \xrightarrow{\tau^*} \langle \lambda.s', d', [] \rangle_{ev}$ for some s' and d' .*

Contextual equivalence is characterized by applicative bisimilarity [1], which reduces closed terms to values that are then applied to an arbitrary argument.

Definition 17. *A symmetric relation \mathcal{R} on closed closures is an applicative bisimulation if $(t, e) \mathcal{R} (s, d)$ and $\langle t, e, [] \rangle_{ev} \xrightarrow{\tau^*} \langle \lambda.t', e', [] \rangle_{ev}$ implies that there exist s' and d' such that*

| | |
|--|-----------|
| $\langle \lambda.t, e, [] \rangle_{\text{ev}} \xrightarrow{\odot} \langle 0, 1, \odot, (t, e) \rangle_{\text{ind}}$ | (ARG) |
| $\langle n, \kappa, \rho, \eta \rangle_{\text{ind}} \xrightarrow{\boxplus} \langle n+1, \kappa+1, \rho, \eta \rangle_{\text{ind}}$ | (SUC) |
| $\langle n, \kappa, \rho, \eta \rangle_{\text{ind}} \xrightarrow{\boxminus} \langle n, \kappa, \rho, \eta \rangle_{\text{tm}}$ | (VAR) |
| $\langle t, \kappa+1, \rho, \eta \rangle_{\text{tm}} \xrightarrow{\lambda} \langle \lambda.t, \kappa, \rho, \eta \rangle_{\text{tm}}$ | (LAMBDA) |
| $\langle t, 0, \rho, \eta \rangle_{\text{tm}} \xrightarrow{\lambda} \langle \lambda.t, 0, \rho, \eta \rangle_{\text{tm}}$ | (LAMBDA0) |
| $\langle t, \kappa, \rho, \eta \rangle_{\text{tm}} \xrightarrow{\dashv} \langle 0, 1, (t, \kappa) :: \rho, \eta \rangle_{\text{ind}}$ | (APPFUN) |
| $\langle s, \kappa_1, (t, \kappa_2) :: \rho, \eta \rangle_{\text{tm}} \xrightarrow{\textcircled{a}} \langle t s, \max(\kappa_1, \kappa_2), \rho, \eta \rangle_{\text{tm}}$ | (APP) |
| $\langle t, 0, \odot, (s, e) \rangle_{\text{tm}} \xrightarrow{\star} \langle s, (t, \epsilon) :: e, [] \rangle_{\text{ev}}$ | (RESTART) |

Fig. 4: AB machine: argument generation

$\langle s, d, [] \rangle_{\text{ev}} \xrightarrow{\tau}^* \langle \lambda.s', d', [] \rangle_{\text{ev}}$, and for all closed t'' , we have $\langle t, (t'', \epsilon) :: e \rangle \mathcal{R} \langle s, (t'', \epsilon) :: d \rangle$.

Applicative bisimilarity \approx_{app} is the largest applicative bisimulation.

We can prove full abstraction between HOcore and the λ -calculus by either internalizing contextual equivalence or applicative bisimilarity. We choose the latter, as it is closer to normal-form bisimilarity, and we show in Section V-D the machine for contextual equivalence.

B. Argument Generation for the Applicative Bisimilarity

After evaluating a term thanks to the KAM, we want to produce a closed argument to pass it to the resulting value, and then restart the evaluation process. If we represent a λ -term as a syntax tree, then de Bruijn indices are at the leaves, and applications and λ -abstractions at the nodes. We start generating from the leftmost de Bruijn index, and we then go left to right, meaning that in an application, we create the term in function position before the argument. These choices are completely arbitrary, as doing the opposite—starting from the rightmost index and go right to left—is also possible. To be sure that we produce a valid, closed, λ -term, we have to check that each de Bruijn index n has at least $n+1$ λ -abstractions enclosing it, and that each application node has two children.

To do so, we consider machine states with four components: the term t being constructed, a counter κ giving the minimal number of λ -abstractions required to close the term, a stack ρ used to build applications, whose syntax is

$$\rho ::= (t, \kappa) :: \rho \mid \odot$$

and which is explained in more detail later, and finally the closure η for which the argument is being built. This last element is never modified by the building process, and is just used to restart the machine in evaluation mode when the argument is finished. We distinguish two kinds of states: the index state $\langle n, \kappa, \rho, \eta \rangle_{\text{ind}}$, where only de Bruijn indices can be built, and the term state $\langle t, \kappa, \rho, \eta \rangle_{\text{tm}}$, where any term can be produced. The transitions for these states are given in Figure 4.

The transition ARG starts the building process when we reach a λ -abstraction in evaluation mode with the empty continuation $[]$. We begin with the index 0, which requires

at least 1 λ -abstraction above it, and with the empty stack \odot . The value of the index can then be increased with the transition SUC, which accordingly also increases the value of κ . When we reach the needed value for the index, the transition VAR switches to the term mode; we use two modes to prevent a SUC transition on a term which is not an index.

In term mode, we can add λ -abstractions to the term, decreasing κ if $\kappa > 0$ with transition LAMBDA, or leaving κ at 0 with transition LAMBDA0; the abstractions we introduce when $\kappa = 0$ do not bind any variable. Once we are done building a term t in function position of an application, we use transition APPFUN to build the argument s . We start again in index mode, but we store on top of ρ the term t with its counter κ_2 . When we finish s with a counter κ_1 , we build the application with transition APP, which takes the maximum of κ_1 and κ_2 as the new minimal number of λ -abstractions needed above $t s$. Note that the APP transition is allowed only if ρ is not empty, meaning that at least one APPFUN has been done before. Finally, we can conclude the term building process with transition RESTART only if $\kappa = 0$, meaning that all the variables of the term are bound, and if ρ is empty, meaning that there is no application waiting to be finished.

Example 18. Figure 5 presents how we generate the term $\lambda.\lambda.(\lambda.0)$ ($1 \lambda.0$); we start with the underlined 0.

Any closed term t can be generated with the AB machine, and it is possible to define the sequence of flags $\text{Seq}(t)$ that will be raised in the process. We write $()$ for the empty sequence, and $(f_1, \dots, f_n, (f'_1, \dots, f'_m), f_{n+1}, \dots, f_l)$ for the sequence $(f_1, \dots, f_n, f'_1, \dots, f'_m, f_{n+1}, \dots, f_l)$.

Definition 19. Given a term t , we define $\text{Seq}(t)$ as

$$\begin{aligned} \text{Seq}(t) &\triangleq (\text{SeqTm}(t), \star) \\ \text{SeqTm}(t s) &\triangleq (\text{SeqTm}(t), \dashv, \text{SeqTm}(s), \textcircled{a}) \\ \text{SeqTm}(\lambda.t) &\triangleq (\text{Seq}(t), \lambda) \\ \text{SeqTm}(n) &\triangleq (\text{SeqInd}(n), \boxminus) \\ \text{SeqInd}(0) &\triangleq () \\ \text{SeqInd}(n+1) &\triangleq (\text{SeqInd}(n), \boxplus) \end{aligned}$$

We write $C \xrightarrow{\text{Seq}(t)} C'$ for $C \xrightarrow{f_1} \dots \xrightarrow{f_m} C'$ where $\text{Seq}(t) = (f_1, \dots, f_m)$.

Lemma 20. If t' is closed, then $\langle 0, 1, \odot, (t, e) \rangle_{\text{ind}} \xrightarrow{\text{Seq}(t')} \langle t, (t', \epsilon) :: e, [] \rangle_{\text{ev}}$.

This lemma allows us to prove the correspondence between the AB machine and applicative bisimilarity.

Theorem 21. $(t, e) \approx_{\text{app}} (s, d)$ iff $\langle t, e, [] \rangle_{\text{ev}} \approx_m \langle s, d, [] \rangle_{\text{ev}}$.

Sketch. To prove that machine bisimilarity implies applicative bisimilarity, we show that

$$\{((t, e), (s, d)) \mid \langle t, e, [] \rangle_{\text{ev}} \approx_m \langle s, d, [] \rangle_{\text{ev}}\}$$

is an applicative bisimulation. Roughly, (t, e) evaluating to $(\lambda.t', e')$ means that (s, d) evaluates to $(\lambda.s', d')$ thanks to \odot . Then for all closed t'' , the sequence $\langle 0, 1, \odot, (t', e') \rangle_{\text{ind}} \xrightarrow{\text{Seq}(t'')} \langle t', (t'', \epsilon) :: e', [] \rangle_{\text{ev}}$ can only be

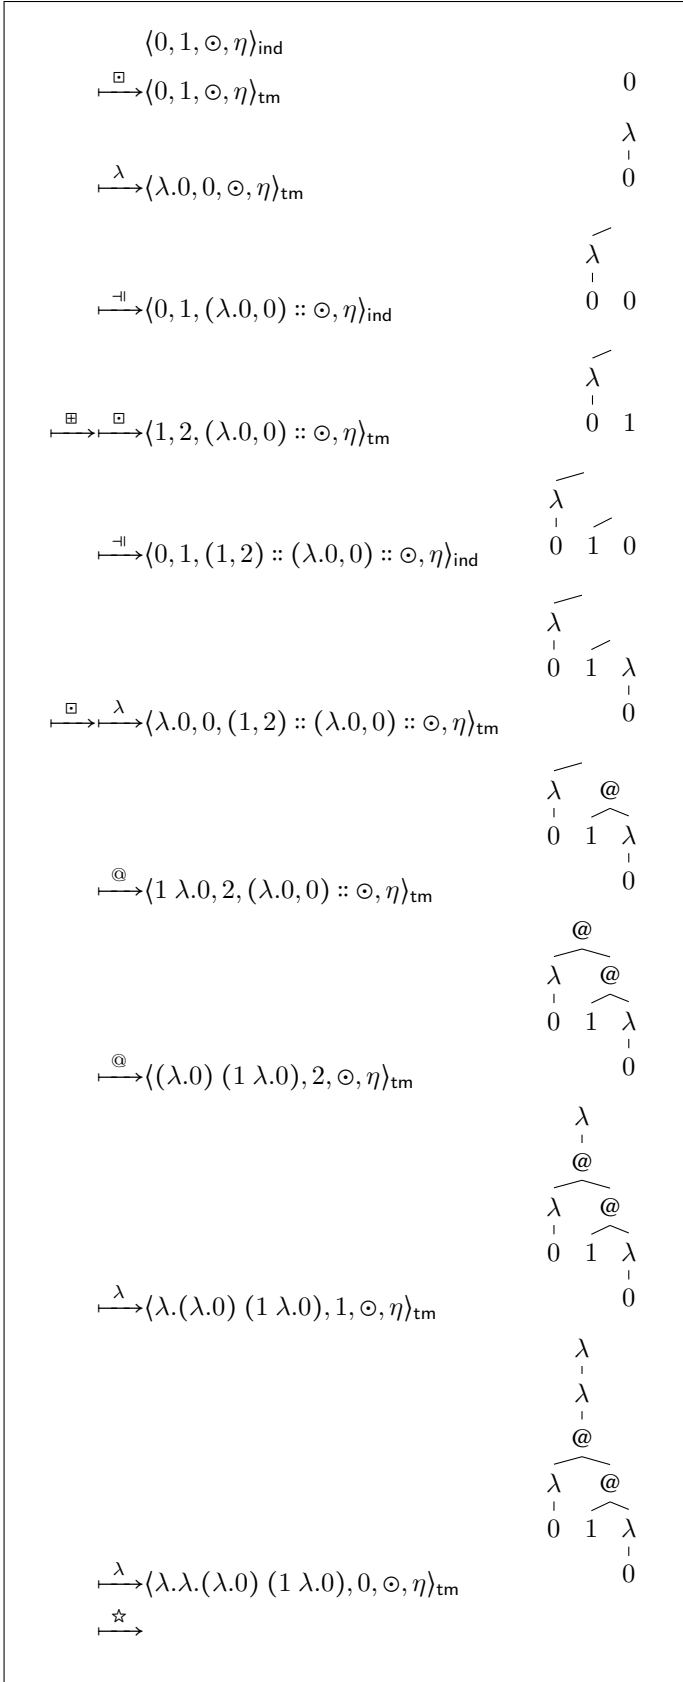


Fig. 5: Example of argument generation

matched by $\langle 0, 1, \odot, (s', d') \rangle_{\text{ind}} \xrightarrow{\text{Seq}(t'')} \langle s', (t'', \epsilon) :: d', [] \rangle_{\text{ev}}$, so we can conclude.

For the reverse implication, we show that

$$\begin{aligned} \mathcal{R} \triangleq & \{ \langle (t, e, [])_{\text{ev}}, \langle s, d, [] \rangle_{\text{ev}} \mid (t, e) \approx_{\text{app}} (s, d) \} \\ & \cup \{ \langle (n, \kappa, \rho, (t, e))_{\text{ind}}, \langle n, \kappa, \rho, (s, d) \rangle_{\text{ind}} \mid \\ & \quad (\lambda.t, e) \approx_{\text{app}} (\lambda.s, d) \} \\ & \cup \{ \langle (t', \kappa, \rho, (t, e))_{\text{tm}}, \langle t', \kappa, \rho, (s, d) \rangle_{\text{tm}} \mid \\ & \quad (\lambda.t, e) \approx_{\text{app}} (\lambda.s, d) \} \end{aligned}$$

is a machine bisimulation, which is easy to check. \square

C. Translation into HOCORE

Figure 6 gives the translation of the AB machine in HOCORE. We detail each component, starting with the evaluation mode, i.e., the KAM. We follow the same principles as in Section III: a non-empty stack π or environment e is represented by a pair of messages, respectively on hd_c and c , and hd_e and env . A closure is represented by two messages, one containing the term on η_1 and one containing the environment on η_2 . The process representing the empty environment ϵ should never be executed, because all the closures we manipulate are closed; as a result, we can choose any process to represent it, e.g., $\mathbf{0}$. The empty stack $[[[]]]$ and the process P_{rec} are used to generate an argument and are explained later.

The encoding of $t s$ simulates the rule PUSH: we receive the current stack and environment e to create the new stack with (s, e) on top. Because we receive the current environment to put it on the stack, we have to recreate it on env , unchanged. In the encoding of $\lambda.t$, we capture the stack and environment, and if the stack is non-empty, we fetch its head η to create a new environment with η on top. Finally, a de Bruijn index $n > 0$ go through the current environment, until we reach the correct closure (case $n = 0$). In that case, we receive the head η and tail of the environment, discard the tail as it is no longer useful, and we restore the term and environment stored in η .

If $\lambda.t$ is run in the environment e and the empty stack $[]$, then we obtain $[[[]]] \parallel hd_c(z).(\llbracket t \rrbracket \parallel \overline{env}(hd_e(z) \parallel \overline{env}(\llbracket e \rrbracket)))$, so $[[[]]]$ has to start the argument generating process, and the result has then to be sent on hd_c for the evaluation to restart. We write $\text{Stuck}((t, e))$ for the process in parallel with $[[[]]]$, which remains stuck during the whole generation process. We now explain how $\langle n, \kappa, \rho, \eta \rangle_{\text{ind}}$ and $\langle t, \kappa, \rho, \eta \rangle_{\text{tm}}$ are encoded, starting with κ and ρ .

The machine distinguishes cases based on whether κ is 0 or not, to know if we should apply the transition LAMBDA or LAMBDA0. In the encoding of these rules (see the definition of Lambda), we send on name $zero$ the expected behavior if $\kappa = 0$, and on $succ$ what to do otherwise. The translation of the counter receives both messages, executes the corresponding one (e.g., the one on $zero$ for the encoding of 0), and discards the other. Apart from that, κ is translated as a natural number. Similarly, the translation of ρ combines the regular encodings of pairs and stacks, but also indicates whether ρ is empty or not, to know if we can apply the transitions APP and RESTART.

1) Evaluation mode

$$\begin{aligned} \llbracket (t, e, \pi)_{\text{ev}} \rrbracket &\triangleq \llbracket t \rrbracket \parallel \overline{\text{env}}\langle \llbracket e \rrbracket \rangle \parallel \overline{c}\langle \llbracket \pi \rrbracket \rangle \parallel P_{\text{rec}} & \llbracket (t, e) \rrbracket &\triangleq \overline{\eta}_1\langle \llbracket t \rrbracket \rangle \parallel \overline{\eta}_2\langle \llbracket e \rrbracket \rangle \\ \llbracket [] \rrbracket &\triangleq \odot(_).(\overline{\text{ind}}\langle \llbracket 0 \rrbracket \rangle \parallel \overline{k}\langle \llbracket 1 \rrbracket_c \rangle \parallel \overline{r}\langle \llbracket \odot \rrbracket \rangle \parallel \overline{\text{initInd}}\langle \mathbf{0} \rangle) & \llbracket \epsilon \rrbracket &\triangleq \mathbf{0} \\ \llbracket \eta :: \pi \rrbracket &\triangleq \overline{\text{hd}}_c\langle \llbracket \eta \rrbracket \rangle \parallel \overline{c}\langle \llbracket \pi \rrbracket \rangle & \llbracket \eta :: e \rrbracket &\triangleq \overline{\text{hd}}_e\langle \llbracket \eta \rrbracket \rangle \parallel \overline{\text{env}}\langle \llbracket e \rrbracket \rangle \end{aligned}$$

$$\begin{aligned} \llbracket t s \rrbracket &\triangleq \text{App}_{\text{eval}}(\llbracket t \rrbracket, \llbracket s \rrbracket) & \text{App}_{\text{eval}}(P_t, P_s) &\triangleq c(x).\text{env}(y).(P_t \parallel \overline{c}\langle \overline{\text{hd}}_c\langle \overline{v}_1\langle P_s \rangle \parallel \overline{v}_2\langle y \rangle \rangle \parallel x) \\ \llbracket \lambda.t \rrbracket &\triangleq \text{Lam}_{\text{eval}}(\llbracket t \rrbracket) & \text{Lam}_{\text{eval}}(P_t) &\triangleq c(x).(x \parallel \text{hd}_c(y).\text{env}(z).(P_t \parallel \overline{\text{env}}\langle \overline{\text{hd}}_e\langle y \rangle \parallel \overline{\text{env}}\langle z \rangle \rangle)) \\ \llbracket n + 1 \rrbracket &\triangleq \text{Ind}_{\text{eval}}(\llbracket n \rrbracket) & \text{Ind}_{\text{eval}}(P_n) &\triangleq \text{env}(q).x(x \parallel \text{hd}_e(_).P_n) \\ \llbracket 0 \rrbracket &\triangleq \text{env}(x).(x \parallel \text{hd}_e(y).\text{env}(_).(y \parallel \eta_1(y_1).\eta_2(y_2).(y_1 \parallel \overline{\text{env}}\langle y_2 \rangle))) & & \\ P_{\text{rec}} &\triangleq \text{RecInd} \parallel \overline{\text{recint}}\langle \text{RecInd} \rangle \parallel \text{RecTm} \parallel \overline{\text{rectm}}\langle \text{RecTm} \rangle \parallel \text{RecMax} \parallel \overline{\text{recmax}}\langle \text{RecMax} \rangle \end{aligned}$$

2) Counter κ , stack ρ , and stuck process $\text{Stuck}(\eta)$

$$\begin{aligned} \llbracket 0 \rrbracket_c &\triangleq \text{zero}(x).\text{succ}(_).x & \llbracket \odot \rrbracket &\triangleq \text{mt}(x).\text{cons}(_).x \\ \llbracket \kappa + 1 \rrbracket_c &\triangleq \text{Suk}(\llbracket \kappa \rrbracket_c) & \llbracket (t, \kappa) :: \rho \rrbracket &\triangleq \text{ConsR}(\llbracket (t, \kappa) \rrbracket, \llbracket \rho \rrbracket) \\ \text{Suk}(P_\kappa) &\triangleq \overline{\text{suk}}\langle P_\kappa \rangle \parallel \text{zero}(_).\text{succ}(x).x & \text{ConsR}(P_{\text{hd}}, P_\rho) &\triangleq \overline{\text{hd}}_r\langle P_{\text{hd}} \rangle \parallel \overline{r}\langle P_\rho \rangle \parallel \text{mt}(_).\text{cons}(x).x \\ \text{Stuck}((t, e)) &\triangleq \text{hd}_c(z).(\llbracket t \rrbracket \parallel \overline{\text{env}}\langle \overline{\text{hd}}_e\langle z \rangle \parallel \overline{\text{env}}\langle \llbracket e \rrbracket \rangle \rangle) & \llbracket (t, \kappa) \rrbracket &\triangleq \overline{w}_1\langle \llbracket t \rrbracket \rangle \parallel \overline{w}_2\langle \llbracket \kappa \rrbracket_c \rangle \end{aligned}$$

3) Creation of indices

$$\begin{aligned} \langle n, \kappa, \rho, \eta \rangle_{\text{ind}} &\triangleq \overline{\text{ind}}\langle \llbracket n \rrbracket \rangle \parallel \overline{k}\langle \llbracket \kappa \rrbracket_c \rangle \parallel \overline{r}\langle \llbracket \rho \rrbracket \rangle \parallel P_{\text{rec}} \parallel \overline{\text{initInd}}\langle \mathbf{0} \rangle \parallel \text{Stuck}(\eta) \\ \text{RecInd} &\triangleq \overline{\text{initInd}}(_).\text{recind}(x).(x \parallel \overline{\text{recind}}\langle x \rangle \parallel \text{Succ} + \text{Var}) \\ \text{Succ} &\triangleq \boxplus(_).\text{ind}(x).k(y).(\overline{\text{ind}}\langle \text{Ind}_{\text{eval}}(x) \rangle \parallel \overline{k}\langle \text{Suk}(y) \rangle \parallel \overline{\text{initInd}}\langle \mathbf{0} \rangle) \\ \text{Var} &\triangleq \boxminus(_).\text{ind}(x).(\overline{\text{tm}}\langle x \rangle \parallel \overline{\text{initTm}}\langle \mathbf{0} \rangle) \end{aligned}$$

4) Creation of terms

$$\begin{aligned} \langle t, \kappa, \rho, \eta \rangle_{\text{tm}} &\triangleq \overline{\text{tm}}\langle \llbracket t \rrbracket \rangle \parallel \overline{k}\langle \llbracket \kappa \rrbracket_c \rangle \parallel \overline{r}\langle \llbracket \rho \rrbracket \rangle \parallel P_{\text{rec}} \parallel \overline{\text{initTm}}\langle \mathbf{0} \rangle \parallel \text{Stuck}(\eta) \\ \text{Lambda}(P_\kappa) &\triangleq \lambda(_).\text{tm}(x). \left(\overline{\text{tm}}\langle \text{Lam}_{\text{eval}}(x) \rangle \parallel P_\kappa \parallel \overline{\text{zero}}\langle \overline{k}\langle P_\kappa \rangle \parallel \overline{\text{initTm}}\langle \mathbf{0} \rangle \rangle \right. \\ &\quad \left. \parallel \overline{\text{succ}}\langle \text{suk}(y).(\overline{k}\langle y \rangle \parallel \overline{\text{initTm}}\langle \mathbf{0} \rangle) \rangle \right) \\ \text{AppFun}(P_\kappa, P_\rho) &\triangleq \text{=}\!|\!|(_).\text{tm}(x).(\overline{r}\langle \text{ConsR}(\overline{\text{hd}}_r\langle \overline{w}_1\langle x \rangle \parallel \overline{w}_2\langle P_\kappa \rangle \rangle, P_\rho) \rangle \parallel \overline{\text{ind}}\langle \llbracket 0 \rrbracket \rangle \parallel \overline{k}\langle \llbracket 1 \rrbracket_c \rangle \parallel \overline{\text{initInd}}\langle \mathbf{0} \rangle) \\ \text{Done} &\triangleq \star(_).\text{tm}(x).(\overline{\text{hd}}_c\langle \overline{\eta}_1\langle x \rangle \parallel \overline{\eta}_2\langle \llbracket \epsilon \rrbracket \rangle \rangle \parallel \overline{c}\langle \llbracket [] \rrbracket \rangle) \\ \text{App}(P_\kappa, P_{\text{hd}}, P_\rho) &\triangleq \text{@}(_).\text{tm}(x_2). \left(P_{\text{hd}} \parallel \overline{w}_2(y).\overline{w}_1(x_1). \left(\overline{\text{max1}}\langle y \rangle \parallel \overline{\text{max2}}\langle P_\kappa \rangle \parallel \overline{\text{init1}}\langle y \rangle \parallel \overline{\text{init2}}\langle P_\kappa \rangle \parallel \right. \right. \\ &\quad \left. \left. \overline{\text{resu}}\langle z \rangle.(\overline{\text{tm}}\langle \text{App}_{\text{eval}}(x_1, x_2) \rangle \parallel P_\rho \parallel \overline{k}\langle z \rangle \parallel \overline{\text{initTm}}\langle \mathbf{0} \rangle) \right) \right) \\ \text{RecMax} &\triangleq \overline{\text{init1}}\langle x_1 \rangle.\overline{\text{init2}}\langle x_2 \rangle.\overline{\text{recmax}}\langle y \rangle.(y \parallel \overline{\text{recmax}}\langle y \rangle \parallel \text{Max}(x_1, x_2)) \\ \text{Max}(P_1, P_2) &\triangleq P_1 \parallel \overline{\text{zero}}\langle \overline{\text{max2}}\langle x \rangle.\overline{\text{resu}}\langle x \rangle \rangle \parallel \overline{\text{succ}}\left\langle \text{suk}(x_1). \left(P_2 \parallel \overline{\text{zero}}\langle \overline{\text{max1}}\langle x \rangle.\overline{\text{resu}}\langle x \rangle \rangle \right. \right. \\ &\quad \left. \left. \parallel \overline{\text{succ}}\langle \text{suk}(x_2).(\overline{\text{init1}}\langle x_1 \rangle \parallel \overline{\text{init2}}\langle x_2 \rangle) \rangle \right) \right\rangle \\ \text{RecTm} &\triangleq \overline{\text{initTm}}(_).\overline{\text{rectm}}(x). \left(x \parallel \overline{\text{rectm}}\langle x \rangle \parallel r(y).y \parallel \overline{\text{cons}}\left\langle k(z).\text{hd}_r(y_1).r(y_2). \right. \right. \\ &\quad \left. \left. (\text{Lambda}(z) \parallel y) + \text{AppFun}(z, y) + \text{App}(z, y_1, y_2) \right\rangle \right. \\ &\quad \left. \parallel \overline{\text{mt}}\left\langle k(z). \left(z \parallel \overline{\text{zero}}\left\langle \begin{array}{l} (\text{Lambda}(z) \parallel \llbracket \odot \rrbracket) \\ + \text{AppFun}(z, \llbracket \odot \rrbracket) + \text{Done} \end{array} \right\rangle \right. \right. \right. \\ &\quad \left. \left. \parallel \overline{\text{succ}}\left\langle \text{suk}(_). \begin{array}{l} (\text{Lambda}(z) \parallel \llbracket \odot \rrbracket) \\ + \text{AppFun}(z, \llbracket \odot \rrbracket) \end{array} \right\rangle \right) \right\rangle \right) \end{aligned}$$

Fig. 6: Translation of the AB machine

After flagging \odot , the process $\llbracket [] \rrbracket$ starts the argument generation process in index mode: the index being built is sent on ind (here, initialized with $\llbracket 0 \rrbracket$), the counter on k , and the stack on r . The message on $initInd$ triggers the recursive process $RecInd$, which non-deterministically chooses between Succ and Var. Executing Succ flags \boxplus , increases the values of the index (thanks to Ind_{eval}) and the counter (with Suk), and relaunches the $RecInd$ process with a message on $initInt$. Executing Var flags \boxminus , moves the index from ind to tm , and initiates the term mode by sending a message on $initTm$, which triggers the process $RecTm$.

The goal of $RecTm$ is to non-deterministically choose between the four transitions available in term mode, namely $\xrightarrow{\lambda}$, $\xrightarrow{-\parallel}$, $\xrightarrow{\textcircled{A}}$, and $\xrightarrow{\star}$. However, some of these transitions have requirements: $\xrightarrow{\textcircled{A}}$ needs $\rho \neq \odot$ and $\xrightarrow{\star}$ needs $\rho = \odot$ and $\kappa = 0$. The process $RecTm$ is therefore doing a case analysis to check these conditions. First, it captures $\llbracket \rho \rrbracket$ on r : if $\rho \neq \odot$, it executes the message on $cons$, which makes a choice between λ , $-\parallel$, and \textcircled{A} , which are represented by respectively Lambda, AppFun, and App. If $\rho = \odot$, then we do a case analysis on κ . If $\kappa = 0$, then we can do either λ , $-\parallel$, or \star (represented by Done), otherwise, only λ or $-\parallel$ are possible.

The process Lambda adds a λ -abstraction to the term in tm , updating κ (represented by P_κ) accordingly: if $\kappa = 0$, then it is restored unchanged on k , otherwise, it is decreased by 1 by releasing the message in suk . The process AppFun pushes on the stack P_ρ the current term t_1 on tm and its counter κ_1 (represented by P_κ), which is the term in function position of an application. It then relaunches the index mode to build the argument t_2 with its counter κ_2 . The process App can then build the application itself, by computing the maximum between κ_1 and κ_2 with the processes $RecMax$ and Max .

We compute the maximum between κ_1 and κ_2 by removing the layers of successors common to κ_1 and κ_2 , until we reach 0 for one of them. If we reach 0 for κ_1 first, then κ_2 is the \max , otherwise it is κ_1 . We store the initial values of κ_1 and κ_2 in respectively $max1$ and $max2$, and the decomposition occurs in Max , where $init1$ is initialized with κ_1 and $init2$ with κ_2 . If $P_1 = \llbracket \kappa_1 \rrbracket_c = \llbracket 0 \rrbracket_c$, then we send κ_2 (stored in $max2$) on $resu$. Otherwise, $P_1 = suk\langle P'_1 \rangle$, and we do a case analysis on the process $P_2 = \llbracket \kappa_2 \rrbracket_c$. If $P_2 = \llbracket 0 \rrbracket_c$, then we send κ_1 on $resu$, otherwise $P_2 = suk\langle P'_2 \rangle$, and we restart $RecMax$ by sending P'_1 and P'_2 on $init1$ and $init2$, respectively. Once the \max is known on $resu$, then App builds the application $t_1 t_2$ and relaunches $RecTm$.

Finally, the process Done ends the argument generation phase, and restarts the computation by restoring the empty continuation and by passing the term in tm to $Stuck(\eta)$. The process P_{rec} contains all the processes necessary to encode the different recursive mechanisms.

Full abstraction: We use the notions of Section IV-D to prove the correspondence between the AB machine and its translation. The definition of next carries over to the translation of the AB machine, and Lemmas 12 and 13 still hold (but without terminating transitions). Recall that \approx_{HO} stands for $\approx_{HO}^{\mathbb{H}}$, where \mathbb{H} contains the names of the translation




Fig. 7: Contextual equivalence machine

that are not flags. The proof of the following theorem is then the same as for Theorem 14.

Theorem 22. $C \approx_m C'$ iff $\llbracket C \rrbracket \approx_{HO} \llbracket C' \rrbracket$.

We then deduce a full abstraction result between λ -calculus with applicative bisimilarity and HOcore.

Corollary 23. If (t, e) and (s, d) are closed closures, then $(t, e) \approx_{app} (s, d)$ iff $\llbracket \langle t, e, [] \rangle_{ev} \rrbracket \approx_{HO} \llbracket \langle s, d, [] \rangle_{ev} \rrbracket$.

Remark 24. As explained in the introduction, the encoding of [22] is not complete w.r.t. applicative bisimilarity, while ours is. The difference is that our translation is tailored to protect itself from bad behaviors of the environment, by limiting its interactions with the outside to synchronizations on flags. A translated λ -abstraction in [22] is waiting for a channel name which gives access to the argument, but the environment may use this name to give access to a process which is not an encoded λ -term, thus breaking the encoding.

D. Internalizing Contextual Equivalence

Corollary 23 is enough to deduce full abstraction w.r.t. contextual equivalence, since $t \approx_{ctx} s \iff (t, \epsilon) \approx_{app} (s, \epsilon)$. However, it is possible to prove this result directly, by internalizing contextual equivalence in an abstract machine.

Figure 7 gives the transitions of this machine, except for the $\xrightarrow{\tau}$ transitions, which are the same as in Section V-A. In contrast with the AB machine, the contextual equivalence machine produces a context first, and then reduces the resulting term; consequently, the starting point is a state $\langle t, 0, \odot \rangle_{tm}$, where t is the closed term we want to plug in the context. When the context is finished, the transition $\xrightarrow{\textcircled{S}}$ switches to the evaluation mode. Also, the evaluation part of the machine is not executed several times, since \approx_{ctx} is not coinductive. We flag \star when the evaluation terminates, to distinguish a terminating term from a diverging one.

Creating a context \mathcal{C} is almost the same as generating an argument in the AB machine, except that we want to plug a closed term t inside. We build $\mathcal{C}[t]$ by starting the generation process from t ; t can be anywhere in $\mathcal{C}[t]$, not necessarily at

the leftmost position, so we cannot do the generation process going left to right in an application, as with the AB machine (Section V-B). Instead, after producing a term t and with a term s on the stack, we can do either the transition APPLEFT to build $t s$, or APPRIGHT to build $s t$.

Example 25. We show how to generate the context $\lambda.(0\ 0) (\square\ 0)$ around t .

$$\begin{aligned}
\langle t, 0, \odot \rangle_{\text{tm}} &\xrightarrow{-i} \langle 0, 1, (t, 0) :: \odot \rangle_{\text{tm}} \\
&\xrightarrow{\overline{\odot}} \langle t\ 0, 1, \odot \rangle_{\text{tm}} \\
&\xrightarrow{-i} \langle 0, 1, (t\ 0, 1) :: \odot \rangle_{\text{tm}} \\
&\xrightarrow{-i} \langle 0, 1, (0, 1) :: (t\ 0, 1) :: \odot \rangle_{\text{tm}} \\
&\xrightarrow{\overline{\odot}} \langle 0\ 0, 1, (t\ 0, 1) :: \odot \rangle_{\text{tm}} \\
&\xrightarrow{\overline{\odot}} \langle (0\ 0) (t\ 0), 1, \odot \rangle_{\text{tm}} \\
&\xrightarrow{\lambda} \langle \lambda.(0\ 0) (t\ 0), 0, \odot \rangle_{\text{tm}}
\end{aligned}$$

The translation of the contextual equivalence machine into HOcore and the full abstraction proofs are then the same as with the AB machine.

Theorem 26. If t and s are closed terms, then $t \approx_{\text{ctx}} s$ iff $\llbracket \langle t, 0, \odot \rangle_{\text{tm}} \rrbracket \approx_{\text{HO}} \llbracket \langle s, 0, \odot \rangle_{\text{tm}} \rrbracket$.

VI. CONCLUSION AND FUTURE WORK

We propose encodings of the call-by-name λ -calculus into HOcore, fully abstract w.r.t. different equivalences of the λ -calculus, namely normal-form and applicative bisimilarities, and contextual equivalence. This shows that a minimal higher-order calculus with a fixed number of hidden names, which is much less expressive than the name-passing π -calculus, still has enough expressive power to faithfully encode the call-by-name λ -calculus. Note that our encodings can immediately be ported to $\text{HO}\pi$ by adding a few top-level restrictions. Our full-abstraction results then hold with the usual barbed equivalence of $\text{HO}\pi$, no longer requiring hidden names. This, however, gives no intuition on the expressiveness of HOcore.

We use abstract machines not only to fix the reduction strategy, but also as an intermediary step between the equivalences of the λ -calculus and HOcore. We turn the equivalences of the λ -calculus, and their potentially complex testing conditions, into a bisimilarity over a labeled transition system (a flags-generating machine), which is closer to the HOcore equivalence. We believe this internalization technique can be applied to other languages for which an abstract machine has been defined, like, e.g., the call-by-value calculus and its CK machine [11] (see Appendix C), or a calculus with control operators [7]. Even though the bisimilarities for these calculi can be quite intricate (see, e.g., [8]), it should be always possible to generate a context as in Section V-D to internalize contextual equivalence. We also think we can internalize more complex equivalences, like environmental bisimilarities, using Madiot's framework [18], which expresses these bisimilarities as a labeled transition system.

Finally, the encodings of the extended abstract machines into HOcore rely on the same principles, e.g., to represent

stacks, non-deterministic choice, case analyses on terms, etc. We believe it is possible to automatically derive the encoding from an abstract machine, so that the generated translation is deterministic (up to flags for choice processes, as in Lemma 12) and with an operational correspondence result similar to Lemma 13. As these two ingredients are almost sufficient to get full abstraction between machines and HOcore, it would give us Theorem 14 or Theorem 22 for free.

REFERENCES

- [1] S. Abramsky and C.-H. L. Ong. Full abstraction in the lazy lambda calculus. *Information and Computation*, 105:159–267, 1993.
- [2] B. Accattoli. Evaluating functions as processes. In *TERMGRAPH 2013*, volume 110 of *EPTCS*, pages 41–55, 2013.
- [3] M. S. Ager, D. Biernacki, O. Danvy, and J. Midtgaard. A functional correspondence between evaluators and abstract machines. In *PPDP'03*, pages 8–19, 2003. ACM Press.
- [4] R. M. Amadio. A decompilation of the pi-calculus and its application to termination. *CoRR*, abs/1102.2339, 2011.
- [5] E. Beffara. Functions as proofs as processes. *CoRR*, abs/1107.4160, 2011.
- [6] M. Berger, K. Honda, and N. Yoshida. Sequentiality and the pi-calculus. In *TLCA 2001*, number 2044 in *LNCS*, pages 29–45, 2001. Springer-Verlag.
- [7] M. Biernacka, D. Biernacki, and O. Danvy. An operational foundation for delimited continuations in the CPS hierarchy. *LMCS*, 1(2:5):1–39, 2005.
- [8] D. Biernacki, S. Lenglet, and P. Polesiuk. Bisimulations for delimited-control operators. available at <https://hal.inria.fr/hal-01207112v1>, 2015. Accepted for publication at *Information and Computation*.
- [9] M. Cimini, C. S. Coen, and D. Sangiorgi. Functions as processes: Termination and the $\lambda\mu\tilde{\mu}$ -calculus. In *TGC 2010*, volume 6084 of *LNCS*, pages 73–86, 2010. Springer.
- [10] P. Downen, L. Maurer, Z. M. Ariola, and D. Varacca. Continuations, processes, and sharing. In *PPDP 2014*, pages 69–80, 2014. ACM.
- [11] M. Felleisen, R. B. Findler, and M. Flatt. *Semantics Engineering with PLT Redex*. The MIT Press, 2009.
- [12] K. Honda, N. Yoshida, and M. Berger. Process types as a descriptive tool for interaction - control and the pi-calculus. In *RTA-TLCA 2014*, volume 8560 of *LNCS*, pages 1–20, 2014. Springer.
- [13] J.-L. Krivine. A call-by-name lambda-calculus machine. *HOSC*, 20(3):199–207, 2007.
- [14] I. Lanese, J. A. Pérez, D. Sangiorgi, and A. Schmitt. On the expressiveness of polyadic and synchronous communication in higher-order process calculi. In *ICALP 2010*, volume 6199 of *LNCS*, pages 442–453, 2010. Springer.
- [15] I. Lanese, J. A. Pérez, D. Sangiorgi, and A. Schmitt. On the expressiveness and decidability of higher-order process calculi. *Information and Computation*, 209(2):198–226, 2011.
- [16] S. B. Lassen. Bisimulation in untyped lambda calculus: Böhm trees and bisimulation up to context. In *MFPS 1999*, volume 20 of *ENTCS*, pages 346–374, New Orleans, LA, Apr. 1999.
- [17] S. B. Lassen. Eager normal form bisimulation. In *LICS 2005*, pages 345–354, 2005.
- [18] J.-M. Madiot. *Higher-order languages: dualities and bisimulation enhancements*. PhD thesis, Université de Lyon and Università di Bologna, 2015.
- [19] R. Milner. Functions as processes. *Mathematical Structures in Computer Science*, 2(2):119–141, 1992.
- [20] D. Sangiorgi. Bisimulation for higher-order process calculi. *Information and Computation*, 131(2):141–178, 1996.
- [21] D. Sangiorgi. From lambda to pi; or, rediscovering continuations. *Mathematical Structures in Computer Science*, 9(4):367–401, 1999.
- [22] D. Sangiorgi and D. Walker. *The π -Calculus: A Theory of Mobile Processes*. Cambridge University Press, 2001.
- [23] B. Toninho, L. Caires, and F. Pfenning. Functions as session-typed processes. In *FOSSACS'12*, number 7213 in *LNCS*, pages 346–360, 2012.
- [24] S. van Bakel and M. G. Vigliotti. A fully-abstract semantics of lambda-mu in the pi-calculus. In *CL&C 2014*, volume 164 of *EPTCS*, pages 33–47, 2014.
- [25] N. Yoshida, M. Berger, and K. Honda. Strong normalisation in the pi-calculus. *Information and Computation*, 191(2):145–202, 2004.

APPENDIX

A. Proofs for Normal-Form Bisimilarity

Theorem 9. We have $t \approx_{\text{nf}} s$ iff there exists $n > \max(\text{fv}(t) \cup \text{fv}(s))$ such that $\langle t, [], n \rangle_{\text{ev}} \approx_m \langle s, [], n \rangle_{\text{ev}}$.

Proof. To prove that machine bisimilarity implies normal-form bisimilarity, we show that $\mathcal{R} \triangleq \{ \langle t, s \rangle \mid \langle t, [], n \rangle_{\text{ev}} \approx_m \langle s, [], n \rangle_{\text{ev}}, n > \max(\text{fv}(t) \cup \text{fv}(s)) \}$ is a normal-form bisimulation. If $t \star [] \mapsto^* \lambda x.t' \star []$, then $\langle t, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle \lambda x.t', [], n \rangle_{\text{ev}}$ by Lemma 7, and then $\langle \lambda x.t', [], n \rangle_{\text{ev}} \xrightarrow{\circ} \langle [n/x]t', [], n+1 \rangle_{\text{ev}}$. Since $\langle t, [], n \rangle_{\text{ev}} \approx_m \langle s, [], n \rangle_{\text{ev}}$, there exists C' such that $\langle s, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle C' \rangle$ and $\langle [n/x]t', [], n+1 \rangle_{\text{ev}} \approx_m C'$. The $\xrightarrow{\tau}$ transitions cannot change the value of n , and the $\xrightarrow{\circ}$ transition tells us that $\langle s, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle \lambda x.s', [], n \rangle_{\text{ev}} \xrightarrow{\circ} \langle [n/x]s', [], n+1 \rangle_{\text{ev}}$ for some s' . Consequently, we have $s \star [] \mapsto^* \lambda x.s' \star []$ (by Lemma 7) with $\langle [n/x]t', [], n+1 \rangle_{\text{ev}} \mathcal{R} \langle [n/x]s', [], n+1 \rangle_{\text{ev}}$, as wished.

Suppose $t \star [] \mapsto^* \alpha \star \pi$ with $\pi = t_1 :: \dots :: t_m :: []$. Then $\langle t, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle \alpha, \pi, n \rangle_{\text{ev}}$ by Lemma 7, and then $\langle \alpha, \pi, n \rangle_{\text{ev}} \xrightarrow{\alpha} \xrightarrow{m} \xrightarrow{\star} \xrightarrow{\circ}$. Since $\langle t, [], n \rangle_{\text{ev}} \approx_m \langle s, [], n \rangle_{\text{ev}}$, we have $\langle s, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle \alpha, \pi', n \rangle_{\text{ev}}$ as well, which is possible only if $\langle s, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle \alpha, \pi', n \rangle_{\text{ev}}$, where $\pi' = s_1 :: \dots :: s_m :: []$. Let $1 \leq i \leq m$; then $\langle \alpha, \pi, n \rangle_{\text{ev}} \xrightarrow{\alpha} \xrightarrow{i-1} \xrightarrow{\blacktriangleright} \langle t_i, [], n \rangle_{\text{ev}}$, which can only be matched with $\langle \alpha, \pi', n \rangle_{\text{ev}} \xrightarrow{\alpha} \xrightarrow{i-1} \xrightarrow{\blacktriangleright} \langle s_i, [], n \rangle_{\text{ev}}$, therefore we have $\langle t_i, [], n \rangle_{\text{ev}} \approx_m \langle s_i, [], n \rangle_{\text{ev}}$. We have $s \star [] \mapsto^* \alpha \star \pi'$ (by Lemma 7) and $\pi \mathcal{R} \pi'$, as required.

For the reverse implication, we show that

$$\mathcal{R} = \{ \langle (t, [], n)_{\text{ev}}, (s, [], n)_{\text{ev}} \rangle \mid t \approx_{\text{nf}} s, n > \max(\text{fv}(t) \cup \text{fv}(s)) \} \cup \{ \langle (\pi, n)_{\text{ct}}, (\pi', n)_{\text{ct}} \rangle \mid \pi \approx_{\text{nf}} \pi', n > \max(\text{fv}(\pi) \cup \text{fv}(\pi')) \}$$

is a machine bisimulation. If $\langle t, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle [n/x]t', [], n+1 \rangle_{\text{ev}}$, then $t \star [] \mapsto^* \lambda x.t' \star []$ by Lemma 7, so there exists s' such that $s \star [] \mapsto^* \lambda x.s' \star []$ and $[n/x]t' \approx_{\text{nf}} [n/x]s'$. Hence we have $\langle s, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle [n/x]s', [], n+1 \rangle_{\text{ev}}$ by Lemma 7, with $\langle [n/x]t', [], n+1 \rangle_{\text{ev}} \mathcal{R} \langle [n/x]s', [], n+1 \rangle_{\text{ev}}$, as wished. If $\langle t, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle \pi, n \rangle_{\text{ct}}$, then $t \star [] \mapsto^* \alpha \star \pi$ by Lemma 7, so there exists π' such that $s \star [] \mapsto^* \alpha \star \pi'$ and $\pi \approx_{\text{nf}} \pi'$. Hence we have $\langle s, [], n \rangle_{\text{ev}} \xrightarrow{\tau} \langle \pi', n \rangle_{\text{ct}}$ by Lemma 7, with $\langle \pi, n \rangle_{\text{ct}} \mathcal{R} \langle \pi', n \rangle_{\text{ct}}$, as wished.

If $\langle \pi, n \rangle_{\text{ct}} \xrightarrow{\star} \langle \pi', n \rangle_{\text{ct}}$, then $\pi = []$, which implies $\pi' = []$, which gives $\langle \pi', n \rangle_{\text{ct}} \xrightarrow{\star} \langle \pi, n \rangle_{\text{ct}}$. If $\langle \pi, n \rangle_{\text{ct}} \xrightarrow{\blacktriangleright} \langle \pi_2, n \rangle_{\text{ct}}$ or $\langle \pi, n \rangle_{\text{ct}} \xrightarrow{\blacktriangledown} \langle t, [], n \rangle_{\text{ev}}$, then $\pi = t :: \pi_2$, so $\pi' = s :: \pi'_2$ with $t \approx_{\text{nf}} s$ and $\pi_2 \approx_{\text{nf}} \pi'_2$. Hence $\langle \pi', n \rangle_{\text{ct}} \xrightarrow{\blacktriangleright} \langle \pi'_2, n \rangle_{\text{ct}}$ with $\langle \pi_2, n \rangle_{\text{ct}} \mathcal{R} \langle \pi'_2, n \rangle_{\text{ct}}$ or $\langle \pi, n \rangle_{\text{ct}} \xrightarrow{\blacktriangledown} \langle t, [], n \rangle_{\text{ev}}$ with $\langle t, [], n \rangle_{\text{ev}} \mathcal{R} \langle s, [], n \rangle_{\text{ev}}$, as wished. \square

Lemma 12. Let P be a machine process which is not a choice process. If $P \xrightarrow{\hat{f}} P'$ and $P \xrightarrow{\hat{f}} P''$, then $P' = P''$. Let P be a choice process. If $P \xrightarrow{\tau} P'$, $P \xrightarrow{\tau} P''$, and $\text{WkObs}(P') = \text{WkObs}(P'')$, then $P' = P''$.

Proof. A machine process which is not a choice process has at most one possible communication. For choice processes, the flags tell us which branch has been selected. \square

Lemma 13. The following assertions hold:

- $C \xrightarrow{\hat{f}} C'$ iff $\llbracket C \rrbracket \xrightarrow{\hat{f}} \llbracket C' \rrbracket$ and $\text{next}(\hat{f}, \llbracket C \rrbracket) = \llbracket C' \rrbracket$.
- $C \xrightarrow{\star} P$ iff $\llbracket C \rrbracket \xrightarrow{\tau} \xrightarrow{\star} P$ and $P \approx_{\text{HO}} \mathbf{0}$.

Proof. By case analyses on $C \xrightarrow{\tau} C'$, on $C \xrightarrow{\hat{f}} C'$, and $\llbracket C \rrbracket$. \square

The following lemma allows us to reason up to τ -transitions on the HOcore side.

Lemma 27. Let P, P' be machine processes. If $P \xrightarrow{\tau} P'$ and $\text{WkObs}(P) = \text{WkObs}(P')$, then $P \approx_{\text{HO}} P'$.

Proof. One can show that $\{ (P \parallel R, P' \parallel R) \mid P \xrightarrow{\tau} P', \text{WkObs}(P) = \text{WkObs}(P') \} \cup \{ (R, R) \}$ is a barbed bisimulation, using Lemma 12. \square

Lemma 28. The relation $\mathcal{R} \triangleq \{ (C, C') \mid \llbracket C \rrbracket \approx_{\text{HO}} \llbracket C' \rrbracket \}$ is a machine bisimulation.

Proof. Let $C_1 \mathcal{R} C'_1$. Suppose first that $C_1 \xrightarrow{\tau} C_2 \xrightarrow{\hat{f}} C_3$; then $\llbracket C_1 \rrbracket \parallel \bar{f} \xrightarrow{\tau} \llbracket C_2 \rrbracket \parallel \bar{f} \xrightarrow{\tau} \llbracket C_3 \rrbracket$ by Lemma 13, which in turn implies that there exists P such that $\llbracket C'_1 \rrbracket \parallel \bar{f} \xrightarrow{\tau} P$ and $\llbracket C_3 \rrbracket \approx_{\text{HO}} P$. In particular, we have $\text{WkObs}(\llbracket C_3 \rrbracket) = \text{WkObs}(P)$, meaning that $\neg(P \downarrow_{\bar{f}})$. Consequently, there exists P' such that $\llbracket C'_1 \rrbracket \xrightarrow{\tau} P'$ and $P' \downarrow_{\bar{f}}$. Then we have $P' \parallel \bar{f} \xrightarrow{\tau} \text{next}(f, P')$ and $P' \parallel \bar{f} \xrightarrow{\tau} P$. By Lemma 12, we have either $P \xrightarrow{\tau} \text{next}(f, P')$ or $\text{next}(f, P') \xrightarrow{\tau} P$. Suppose $\text{WkObs}(P) \neq \text{WkObs}(\text{next}(f, P'))$. It is possible only if there is a choice process between the two. If we have $\text{next}(f, P') \xrightarrow{\tau} P_c \xrightarrow{\tau} P$, with P_c a choice process, then $\text{WkObs}(P) = \{ \blacktriangleright \}$ or $\text{WkObs}(P) = \{ \blacktriangledown \}$, which contradicts $\text{WkObs}(P) = \text{WkObs}(\llbracket C_3 \rrbracket)$ (the translation of a configuration cannot have only \blacktriangleright or only \blacktriangledown as weak observable actions). If $P \xrightarrow{\tau} P_c \xrightarrow{\tau} \text{next}(f, P')$, with P_c a choice process, then a communication on a flag \blacktriangleright or \blacktriangledown happens between P_c and $\text{next}(f, P')$, as it is the only way to reach a configuration from a choice process. It means that $P \downarrow_{\bar{f}}$, which is again in contradiction with $\text{WkObs}(P) = \text{WkObs}(\llbracket C_3 \rrbracket)$. Consequently, we have $\text{WkObs}(P) = \text{WkObs}(\text{next}(f, P'))$ so by Lemma 27, we have $\text{next}(f, P') \approx_{\text{HO}} P \approx_{\text{HO}} \llbracket C_3 \rrbracket$, and by Lemma 13, we have $C'_1 \xrightarrow{\tau} \xrightarrow{\hat{f}} C'_3$ where $\llbracket C'_3 \rrbracket = \text{next}(f, P')$, so we have $C_3 \mathcal{R} C'_3$, as wished.

Suppose $C_1 \xrightarrow{\tau} C_2 \xrightarrow{\star} P$; then $\llbracket C_1 \rrbracket \parallel \bar{\star} \xrightarrow{\tau} \llbracket C_2 \rrbracket \parallel \bar{\star} \xrightarrow{\tau} P$ with $P \approx_{\text{HO}} \mathbf{0}$ by Lemma 13, which in turn implies that there exists P' such that $\llbracket C'_1 \rrbracket \parallel \bar{\star} \xrightarrow{\tau} P'$ and $P \approx_{\text{HO}} P' \approx_{\text{HO}} \mathbf{0}$, so by Lemma 13, we have $C'_1 \xrightarrow{\tau} \xrightarrow{\star} P'$, as wished. \square

Lemma 29. If $C \approx_m C'$, then $\llbracket C \rrbracket \approx_{\text{HO}} \llbracket C' \rrbracket$.

Proof. We prove that

$$\mathcal{R} \triangleq \left\{ (P_C \parallel R, P_{C'} \parallel R) \mid C \approx_m C', P_C \xrightarrow{\tau} \llbracket C \rrbracket, P_{C'} \xrightarrow{\tau} \llbracket C' \rrbracket \text{ or } \llbracket C \rrbracket \xrightarrow{\tau} P_C, \llbracket C' \rrbracket \xrightarrow{\tau} P_{C'}, \text{WkObs}(P_C) = \text{WkObs}(P_{C'}) \right\} \cup \{ (P \parallel R, P' \parallel R) \mid P \approx_{\text{HO}} \mathbf{0} \approx_{\text{HO}} P' \}$$

is a barbed bisimulation. The congruence condition is an immediate consequence of the definition. The transitions from

R are easy to match, and the observable actions are the same because of the condition $\text{WkObs}(P_C) = \text{WkObs}(P_{C'})$. What is left to check are the reductions.

If $P_C \xrightarrow{\tau} P'_C$, then we show that there exists $P'_{C'} \parallel R$ such that $P_{C'} \parallel R \xrightarrow{\tau} P'_{C'} \parallel R$ and $P'_C \parallel R \mathcal{R} P'_{C'} \parallel R$. If $P_C \xrightarrow{\tau} \llbracket C \rrbracket$ and $P_{C'} \xrightarrow{\tau} \llbracket C' \rrbracket$, then taking $P'_{C'} = \llbracket C' \rrbracket$ works. If $\llbracket C \rrbracket \xrightarrow{\tau} P_C$ and $\llbracket C' \rrbracket \xrightarrow{\tau} P_{C'}$, then we distinguish two cases. If P_C is not a choice process, then $\text{WkObs}(P_C) = \text{WkObs}(P'_C)$, and we can simply choose $P'_{C'} = P_{C'}$. Otherwise, the transition $P_C \xrightarrow{\tau} P'_C$ is making a choice, and we have $\text{WkObs}(P_C) = \{\blacktriangleright, \blacktriangledown\}$, and either $\text{WkObs}(P'_C) = \{\blacktriangleright\}$ or $\text{WkObs}(P'_C) = \{\blacktriangledown\}$. But $\text{WkObs}(P_C) = \text{WkObs}(P_{C'})$, which is possible only if $P_{C'}$ also reduces to a choice process. Then it is possible to choose the corresponding branch, and define $P'_{C'}$ accordingly.

Suppose $P_C \parallel R \xrightarrow{\tau} P'$ with a communication on a flag f — a communication on another name is not possible; then $P_C \downarrow_f$, which is possible only in the case $\llbracket C \rrbracket \xrightarrow{\tau} P_C$, and $R \downarrow_{\bar{f}}$. Suppose $f \neq \star$. Then we also have $P_C \xrightarrow{f} P_{C_2} \xrightarrow{\tau} \llbracket C_2 \rrbracket$ for some P_{C_2} and with $\llbracket C_2 \rrbracket = \text{next}(f, P_C)$, so by Lemma 12, there exist R' such that $P' = P_{C_2} \parallel R'$. We have $\llbracket C \rrbracket \xrightarrow{f} \llbracket C_2 \rrbracket$, so $C \xrightarrow{\tau} \xrightarrow{f} C_2$ holds by Lemma 13. Because $C \approx_m C'$, there exists C'_2 such that $C' \xrightarrow{\tau} \xrightarrow{f} C'_2$ and $C_2 \approx_m C'_2$. By Lemma 13, this implies $\llbracket C' \rrbracket \xrightarrow{f} \llbracket C'_2 \rrbracket$; in particular, there exists $P_{C'_2}$ such that $\llbracket C' \rrbracket \xrightarrow{\tau} \xrightarrow{f} P_{C'_2} \xrightarrow{\tau} \llbracket C'_2 \rrbracket$. By Lemma 12, we also have $P_{C'} \xrightarrow{\tau} \xrightarrow{f} P_{C'_2}$, therefore we have $P_{C'} \parallel R \xrightarrow{\tau} P_{C'_2} \parallel R'$, with $P_{C_2} \parallel R' \mathcal{R} P_{C'_2} \parallel R'$, as wished.

If $f = \star$, then $P_C \xrightarrow{\star} P$ for some P'' such that $P'' \approx_{\text{HO}} \mathbf{0}$, and by Lemma 12, we have $P' = P'' \parallel R'$. We have $\llbracket C \rrbracket \xrightarrow{\tau} \xrightarrow{\star} P''$, so $C \xrightarrow{\tau} \xrightarrow{\star} P''$ holds by Lemma 13. Because $C \approx_m C'$, we have $C' \xrightarrow{\tau} \xrightarrow{\star} P''$. By Lemma 13, this implies $\llbracket C' \rrbracket \xrightarrow{\tau} \xrightarrow{\star} P'''$ for some P''' such that $P''' \approx_{\text{HO}} \mathbf{0}$. By Lemma 12, we also have $P_{C'} \xrightarrow{\tau} \xrightarrow{f} P'''$, therefore we have $P_{C'} \parallel R \xrightarrow{\tau} P''' \parallel R'$, so we obtain processes in the second set of \mathcal{R} . \square

Theorem 14. $C \approx_m C'$ iff $\llbracket C \rrbracket \approx_{\text{HO}} \llbracket C' \rrbracket$.

Proof. By Lemmas 28 and 29. \square

B. Proofs for Applicative Bisimilarity

Lemma 20. If t' is closed, then $\langle 0, 1, \odot, (t, e) \rangle_{\text{ind}} \xrightarrow{\text{Seq}(t')} \langle t, (t', \epsilon) :: e, [] \rangle_{\text{ev}}$.

Proof. Let $\eta = (t, e)$. First, for all n , we have $\langle 0, 1, \rho, \eta \rangle_{\text{ind}} \xrightarrow{n} \langle n, n+1, \rho, \eta \rangle_{\text{tm}}$. Then we show by induction on t' that $\langle 0, 1, \rho, \eta \rangle_{\text{ind}} \xrightarrow{\text{SeqTm}(t')} \langle t', \kappa, \rho, \eta \rangle_{\text{tm}}$ where $\kappa = \max(\text{fv}(t')) + 1$ if t' is not closed, and $\kappa = 0$ otherwise. The case $t = n$ is concluded with the previous observation. If

$t' = t'_1 t'_2$, then

$$\begin{aligned} \langle 0, 1, \rho, \eta \rangle_{\text{ind}} &\xrightarrow{\text{SeqTm}(t'_1)} \langle t'_1, \kappa_1, \rho, \eta \rangle_{\text{tm}} \text{ (by induction)} \\ &\xrightarrow{-\text{!}} \langle 0, 1, (t'_1, \kappa_1) :: \rho, \eta \rangle_{\text{ind}} \\ &\xrightarrow{\text{SeqTm}(t'_2)} \langle t'_2, \kappa_2, (t'_1, \kappa_1) :: \rho, \eta \rangle_{\text{tm}} \text{ (by induction)} \\ &\xrightarrow{\text{!}} \langle t'_1 t'_2, \max(\kappa_1, \kappa_2), \rho, \eta \rangle_{\text{tm}} \end{aligned}$$

And by case analysis on (κ_1, κ_2) , one can check that $\max(\kappa_1, \kappa_2)$ is the desired value. If $t' = \lambda.t''$, then by induction, we have $\langle 0, 1, \rho, \eta \rangle_{\text{ind}} \xrightarrow{\text{SeqTm}(t'')} \langle t'', \kappa, \rho, \eta \rangle_{\text{tm}}$, and then $\langle t'', \kappa, \rho, \eta \rangle_{\text{tm}} \xrightarrow{\lambda} \langle t', \kappa', \rho, \eta \rangle_{\text{tm}}$ where κ' is as wished depending on κ .

This implies that for a closed term t' , we have $\langle 0, 1, \odot, \eta \rangle_{\text{ind}} \xrightarrow{\text{SeqTm}(t')} \langle t', 0, \odot, \eta \rangle_{\text{tm}}$, and the last transition $\xrightarrow{\star}$ gives what we want. \square

Theorem 21. $(t, e) \approx_{\text{app}} (s, d)$ iff $\langle t, e, [] \rangle_{\text{ev}} \approx_m \langle s, d, [] \rangle_{\text{ev}}$.

Proof. To prove that machine bisimilarity implies applicative bisimilarity, we show that $\mathcal{R} \stackrel{\Delta}{=} \{ \langle (t, e), (s, d) \rangle \mid \langle t, e, [] \rangle_{\text{ev}} \approx_m \langle s, d, [] \rangle_{\text{ev}} \}$ is an applicative bisimulation. If $\langle t, e, [] \rangle_{\text{ev}} \xrightarrow{\tau} \langle \lambda.t', e', [] \rangle_{\text{ev}}$, then because we have also $\langle \lambda.t', e', [] \rangle_{\text{ev}} \xrightarrow{\circ} \langle 0, 1, \odot, (t', e') \rangle_{\text{ind}}$, there exist s' and d' such that

$$\langle s, d, [] \rangle_{\text{ev}} \xrightarrow{\tau} \langle \lambda.s', d', [] \rangle_{\text{ev}} \xrightarrow{\circ} \langle 0, 1, \odot, (s', d') \rangle_{\text{ind}}$$

and $\langle 0, 1, \odot, (t', e') \rangle_{\text{ind}} \approx_m \langle 0, 1, \odot, (s', d') \rangle_{\text{ind}}$. For all closed t'' , we have then $\langle 0, 1, \odot, (t', e') \rangle_{\text{ind}} \xrightarrow{\text{Seq}(t)} \langle t', (t'', \epsilon) :: e', [] \rangle_{\text{ev}}$ by Lemma 20, which can only be matched by $\langle 0, 1, \odot, (s', d') \rangle_{\text{ind}} \xrightarrow{\text{Seq}(t)} \langle s', (t'', \epsilon) :: d', [] \rangle_{\text{ev}}$ meaning that $\langle t', (t'', \epsilon) :: e', [] \rangle_{\text{ev}} \approx_m \langle s', (t'', \epsilon) :: d', [] \rangle_{\text{ev}}$. We can then easily conclude.

For the reverse implication, we show that

$$\begin{aligned} \mathcal{R} \stackrel{\Delta}{=} &\{ \langle (t, e, []), (s, d, []), (t, e) \rangle_{\text{ev}} \mid (t, e) \approx_{\text{app}} (s, d) \} \\ &\cup \{ \langle (n, \kappa, \rho, (t, e))_{\text{ind}}, (n, \kappa, \rho, (s, d))_{\text{ind}} \mid (\lambda.t, e) \approx_{\text{app}} (\lambda.s, d) \} \\ &\cup \{ \langle (t', \kappa, \rho, (t, e))_{\text{tm}}, (t', \kappa, \rho, (s, d))_{\text{tm}} \mid (\lambda.t, e) \approx_{\text{app}} (\lambda.s, d) \} \end{aligned}$$

is a machine bisimulation. If $\langle t, e, [] \rangle_{\text{ev}} \xrightarrow{\tau} \langle \lambda.t', e', [] \rangle_{\text{ev}} \xrightarrow{\circ} \langle 0, 1, \odot, (t', e') \rangle_{\text{ind}}$, then we have $\langle t, e, [] \rangle_{\text{ev}} \xrightarrow{\tau} \langle \lambda.t', e', [] \rangle_{\text{ev}}$, so there exists (s', d') such that $\langle s, d, [] \rangle_{\text{ev}} \xrightarrow{\tau} \langle \lambda.s', d', [] \rangle_{\text{ev}}$ so that $(\lambda.t', e') \approx_{\text{app}} (\lambda.s', d')$. But then also $\langle \lambda.t', e', [] \rangle_{\text{ev}} \xrightarrow{\circ} \langle 0, 1, \odot, (s', e') \rangle_{\text{ind}}$ and we obtain configurations in \mathcal{R} . If $\langle t', \kappa, \rho, (t, e) \rangle_{\text{arg}} \xrightarrow{f} C$ with $f \neq \star$ and $\text{arg} \in \{\text{ind}, \text{tm}\}$, then $\langle t', \kappa, \rho, (t, e) \rangle_{\text{arg}} \xrightarrow{f} C'$ where C' is the same as C except for the stored closures, which are not changed and therefore still applicative bisimilar; we have therefore $C \mathcal{R} C'$. If $\langle t', \kappa, \rho, (t, e) \rangle_{\text{tm}} \xrightarrow{\star} \langle t, (t', \epsilon) :: e, [] \rangle_{\text{ev}}$, then $\langle t', \kappa, \rho, (s, d) \rangle_{\text{tm}} \xrightarrow{\star} \langle s, (t', \epsilon) :: d, [] \rangle_{\text{ev}}$, and the resulting terms are in \mathcal{R} because $(\lambda.t, e) \approx_{\text{app}} (\lambda.s, d)$ implies $\langle t, (t', \epsilon) :: e \rangle_{\text{ev}} \approx_{\text{app}} \langle s, (t', \epsilon) :: d \rangle_{\text{ev}}$ for all closed t' . \square

| | |
|--|-----------|
| $\langle t s, E, n \rangle_{\text{ev}} \xrightarrow{\tau} \langle t, E s, n \rangle_{\text{ev}}$ | (FUN) |
| $\langle v, E, n \rangle_{\text{ev}} \xrightarrow{\tau} \langle E, v, n \rangle_{\text{ct}}$ | (SWITCH) |
| $\langle E t, v, n \rangle_{\text{ct}} \xrightarrow{\tau} \langle t, v E, n \rangle_{\text{ev}}$ | (ARG) |
| $\langle (\lambda x.t) E, v, n \rangle_{\text{ct}} \xrightarrow{\tau} \langle [v/x]t, E, n \rangle_{\text{ev}}$ | (BETA) |
| $\langle \square, \alpha, n \rangle_{\text{ct}} \xrightarrow{\alpha \star} \square$ | (DONE) |
| $\langle \square, \lambda x.t, n \rangle_{\text{ct}} \xrightarrow{\circledast} \langle [n/x]t, \square, n+1 \rangle_{\text{ev}}$ | (LAMBDA) |
| $\langle \alpha E, v, n \rangle_{\text{ct}} \xrightarrow{\alpha \blacktriangledown} \langle \square, v, n \rangle_{\text{ct}}$ | (VAL) |
| $\langle \alpha E, v, n \rangle_{\text{ct}} \xrightarrow{\alpha \blacktriangleright} \langle E, n, n+1 \rangle_{\text{ct}}$ | (CONTEXT) |

Fig. 8: NFB machine for call by value

C. Call-by-Value

We informally present the machines for the call-by-value λ -calculus, without the proofs, to see how our techniques can be adapted in that case. The syntax of values and call-by-value contexts are as follows.

$$v ::= \alpha \mid \lambda x.t$$

$$E ::= \square \mid E t \mid v E$$

Contexts are represented inside-out: plugging is defined as $\square[t] \triangleq t$, $(E t_2)[t_1] \triangleq E[t_1 t_2]$, and $(v E)[t] \triangleq E[v t]$. Normal forms are either values or terms of the form $E[\alpha v]$.

We present a variant of the CK machine [11] with evaluation contexts below.

$$t s \star E \xrightarrow{\tau} t \star E s \quad (\text{FUN})$$

$$v \star E \xrightarrow{\tau} E \star_{\text{ct}} v \quad (\text{SWITCH})$$

$$E t \star_{\text{ct}} v \xrightarrow{\tau} t \star v E \quad (\text{ARG})$$

$$(\lambda x.t) E \star_{\text{ct}} v \xrightarrow{\tau} [v/x]t \star E \quad (\text{BETA})$$

When we get a value in an evaluation context, we switch into the continuation mode (rule SWITCH) to evaluate the argument of that value (rule ARG), or do the β -reduction if the argument itself is a value (rule BETA).

Normal-form bisimilarity for call-by-value is defined as follows.

Definition 30. A symmetric relation \mathcal{R} is a normal-form bisimulation if $t \mathcal{R} s$ implies:

- if $t \star \square \xrightarrow{\tau}^* \square \star_{\text{ct}} \alpha$, then $s \star \square \xrightarrow{\tau}^* \square \star_{\text{ct}} \alpha$;
- if $t \star \square \xrightarrow{\tau}^* \square \star_{\text{ct}} \lambda x.t'$, then there exists s' such that $s \star \square \xrightarrow{\tau}^* \square \star_{\text{ct}} \lambda x.s'$ and $[\alpha/x]t' \mathcal{R} [\alpha/x]s'$ for a fresh α ;
- if $t \star \square \xrightarrow{\tau}^* \alpha E \star_{\text{ct}} v$, then there exist E' and v' such that $s \star \square \xrightarrow{\tau}^* \alpha E' \star_{\text{ct}} v'$, $v \mathcal{R} v'$, and $E[\alpha'] \mathcal{R} E'[\alpha']$ for a fresh α' .

Normal-form bisimilarity \approx_{nf} , is the largest normal-form bisimulation.

We extend the CK machine into a NFB machine in Figure 8. We have a counter n to generate fresh variables, as in call-by-name. The four last rules tell us what to do to compare normal forms. If we get a variable α , then we signal its value and we are done (transition DONE). If we have a λ -abstraction, we restart the machine (transition LAMBDA). If we have a

context of the form αE , it means that we have a normal form $E[\alpha v]$: we have to choose if we either test v (transition VAL) or E (transition CONTEXT). To test v , we forget about E , and enter the context mode with $\langle \square, v, n \rangle_{\text{ct}}$ after flagging α and \blacktriangledown . Depending on v , we will then apply the transition DONE or LAMBDA. To test E , we generate a fresh variable to plug into E , meaning that we continue with $\langle E, n, n+1 \rangle_{\text{ct}}$ (after flagging α and \blacktriangleright).

We can translate this machine in HOCORE using the same techniques and prove full abstraction as in call-by-name. Internalizing applicative bisimilarity in call-by-value is straightforward: its definition differs from call-by-value only in the fact that the testing argument must be a λ -abstraction. Consequently, we just use the argument-generating rules from Figure 4 on top of the CK machine, and change the rule RESTART so that it can trigger only if the constructed term is a λ -abstraction.