



Processes, Systems & Tests: Defining Contextual Equivalences

Clément Aubert, Daniele Varacca

► To cite this version:

Clément Aubert, Daniele Varacca. Processes, Systems & Tests: Defining Contextual Equivalences. Proceedings 14th Interaction and Concurrency Experience (ICE 2021), Jun 2021, Virtual Event, Malta. pp.1-21, 10.4204/EPTCS.347.1 . hal-02895417v4

HAL Id: hal-02895417

<https://hal.science/hal-02895417v4>

Submitted on 1 Oct 2021

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.



Distributed under a Creative Commons Attribution 4.0 International License

Processes, Systems & Tests: Defining Contextual Equivalences

Clément Aubert

School of Computer and Cyber Sciences, Augusta University, Georgia, USA

caubert@augusta.edu

Daniele Varacca

LACL, Université Paris-Est Créteil, France

daniele.varacca@u-pec.fr

In this position paper, we would like to offer and defend a template to study equivalences between programs—in the particular framework of process algebras for concurrent computation. We believe that our layered model of development will clarify the distinction that is too often left implicit between the tasks and duties of the programmer and of the tester. It will also enlighten pre-existing issues that have been running across process algebras such as the calculus of communicating systems, the π -calculus—also in its distributed version—or mobile ambients. Our distinction starts by subdividing the notion of process in three conceptually separated entities, that we call *Processes*, *Systems* and *Tests*. While the role of what can be observed and the subtleties in the definitions of congruences have been intensively studied, the fact that *not every process can be tested*, and that *the tester should have access to a different set of tools than the programmer* is curiously left out, or at least not often formally discussed. We argue that this blind spot comes from the under-specification of contexts—environments in which comparisons occur—that play multiple distinct roles but supposedly always “stay the same”.

1 Introduction

In the study of programming languages, contextual equivalences play a central role: to study the behavior of a program, or a process, one needs to observe its interactions with different environments, e.g. what outcomes it produces. If the program is represented by a term in a given syntax, environments are often represented as contexts surrounding the terms. But contexts play multiple roles that serve different actors with different purposes. The programmer uses them to construct larger programs, the user employs them to provide input and obtain an output, and the tester or attacker uses them to debug and compare the program or to try to disrupt its intended behavior.

We believe that representing those different purposes with the same “monolithic” syntactical notion of context forced numerous authors to repeatedly “adjust” their definition of context without always acknowledging it. We also argue that collapsing multiple notions of contexts into one prevented further progress. In this article, we propose a way of clarifying how to define contextual equivalences, and show that having co-existing notions of equivalences legitimates and explains recurring choices, and supports a rigorous guideline to separate the development of a program from its usage and testing.

Maybe in the mind of most of the experts in the formal study of programming language is our proposal too obvious to discuss. However, if this is the case, we believe that this “folklore” remains unwritten, and that since we were not at *that* “seminar at Columbia in 1976”¹, we are to remain in darkness.

We believe the interactive and friendly community of ICE to be the ideal place to reach curious and open-minded actors in the field, and to either be proven wrong or—hopefully—impact some researchers.

¹To re-use in our setting Paul Taylor’s witty comment [83].

Non-technical articles can at times have a tremendous impact [27], and even if we do not claim to have Dijkstra’s influence or talent, we believe that our precise, documented proposition can shed a new light on past results, and frame current reflections and future developments. We also believe that ignoring or downplaying the distinctions we stress have repeatedly caused confusions.

2 The Flow of Testing

We begin by illustrating with a simple Java example the three syntactic notions—*process*, *system* and *test*—we will be using. Imagine giving a user the code `while(i < 10){x *= x; i++;}`. A user cannot execute or use this “snippet” unless it is *wrapped* into a method, with adequate header, and possibly variable declaration(s) and `return` statement. Once the programmer performed this operation, the user can *use* the obtained program, and the tester can *interact* with it further, e.g. by calling it from a `main` method.

All in all, a programmer would build on the snippet, then the tester would build an environment to interact with the resulting program, and we could obtain the code below. Other situations could arise—e.g., if the snippet was already wrapped—, but we believe this is a fair rendering of “the life of a snippet”.

```
public class Main{
    public static int foo(int x){
        int i = 0;
        while(i < 10){
            x *= x;
            i++;
        } // Snippet
        return x;
    } // Wrapping
    public static void main(){
        System.out.print(foo(2));
    }
} // Interaction
```

In this example, the snippet is what we will call a *process*, the snippet once wrapped is what we will call a *system* and the “Interaction” part without the system in it, but with additional “observations”—i.e. measures on the execution, terminal output—, is what we will call a *test*. Our terminology comes from the study of concurrent process algebras, where most of our intuitions and references are located, but let us first make a brief detour to examine how our lens applies to λ -calculus.

3 A Foreword on λ -Calculus

Theoretical languages often take λ -calculus as a model or a comparison basis. It is often said that the λ -calculus is to sequential programs what the π -calculus is to concurrent programs [76,85]. Indeed, pure λ -calculus (i.e. without types or additional features like probabilistic sum [29] or quantum capacities [81,84]) is a reasonable [6], Turing-complete and elegant language, that requires only a couple of operators, one reduction rule and one equivalence relation to produce a rich and meaningful theory, sometimes seen as an idealized target language for functional programming languages.

Since most terms² do not reduce as they are, to study their behavior, one needs first to make them interact with an environment, represented by a context. Contexts are generally defined as “term[s] with some holes” [13, p. 29, 2.1.18], that we prefer to call *slots* and denote $[\square]$. Under this apparent simplicity, they should not be manipulated carelessly, as having multiple slots or not being careful when defining what it means to “fill a slot” can lead to e.g. lose confluence [18, pp. 40–41, Example 2.2.1], and as those issues persist even in the presence of a typing system [35]. Furthermore, definitions and theorems that use contexts frequently impose some restrictions on the contexts considered, to exclude e.g. $(\lambda x.y)[\square]$ that

²Actually, if application, abstraction and variables all count as one, the ratio between normal term and term with redexes is unknown [17]. We imply here “since *most interesting* terms”, in the sense of terms that represent programs.

simply “throw away” the term put in the slot in one step of reduction. Following the above observations, we conclude that contexts often come in two flavors, depending on the nature of the term considered:

For closed terms (i.e. without free variables), a context is essentially a series of arguments to feed the term. This observation allows to define e.g. *solvable terms* [13, p. 171, 8.3.1 and p. 416, 16.2.1].

For open terms (i.e. with free variables), a context is a *Böhm transformation* [13, p. 246, 10.3.3], which is equivalent [13, p. 246, 10.3.4] to a series of abstractions followed by a series of applications, and sometimes called “head context” [7, p. 25].

Being closed corresponds to being “wrapped”—ready to use—, and feeding arguments to a term corresponds to interacting with it from a `main` method: the Böhm transformation actually encapsulates two operations at once. In this case, the interaction can observe different aspects: whether the term terminates, whether it grows in size, etc., but it is generally agreed upon that no additional operator or reduction rule should be used. Actually, the syntax is restricted when testing, as only application is allowed: the tested term should not be wrapped in additional layers of abstraction if it is already closed.

Adding features to the λ -calculus certainly does not restore the supposed purity or unicity of the concept of context, but actually distances it even further from being simply “a term with a slot”. For instance, contexts are narrowed down to term context [84, p. 1126] and surface context [29, pp. 4, 10] for respectively quantum and probabilistic λ -calculus, to “tame” the power of contexts. In resource sensitive extensions of the λ -calculus, the quest for full abstraction even led to a drastic separation of λ -terms between terms and tests [22], a separation naturally extended to contexts [21, p. 73, Figure 2.4].

This variety happened after the 2000’s formal studies of contexts was undertaken [18, 19, 35], which led to the observation that treating contexts “merely as a notation [...] hinders any formal reasoning[, while treating them] as first-class objects [allows] to gain control over variable capturing and, more generally, ‘communication’ between a context and expressions to be put into its holes” [19, p. 29]. It is ironic that λ -calculus took inspiration from a concurrent language to split their syntax in two right at its core [22, p. 97], or to study formally the *communication* between a context and the term in its slot, while concurrent languages sometimes tried to keep the “purity” and indistinguishability of their contexts. It can be noted that this re-definition of contexts had impacts on other fields, e.g. on modal type theory [63].

In the case of concurrent calculi like the calculus of communicating systems (CCS) or the π -calculus, interactions with environments are also represented using a notion of context. But the status of contexts in concurrent calculi is even more unsettling when one notes that, while “wrapping” contexts are of interest mainly for open terms in lambda calculus, *all* terms need a pertinent notion of context in concurrent systems to be tested and observed. Indeed, as the notion of “feeding arguments to a concurrent process” concurs with the idea of wrapping it into a larger process, it seems that the distinction we just made between two kinds of contexts in λ -calculus cannot be ported to concurrent calculi. Our contribution starts by questioning whenever, indeed, process calculi have treated contexts as a uniform notion independently from the nature of the term or what it was used for.

4 Contextual Relations

Comparing terms is at the core of the study of programming languages, and process algebra is no exception. Generally, and similarly to what is done in λ -calculus, a comparison is deemed of interest only if it is valid in every possible context³ an idea formally captured by the notion of (pre-)congruence. An

³As a reviewer noted, “Proving a behavioural equivalence is a congruence has a reading that is somewhat left out of the discussion in the paper. Namely, the behavioural equivalence provides a mathematical notion of behaviour and showing such

equivalence relation \mathcal{R} is usually said to be a congruence if it is closed by context, i.e. if for all P, Q (open or closed) terms, $(P, Q) \in \mathcal{R}$ implies that for all context $C[\square]$, $(C[P], C[Q]) \in \mathcal{R}$. (Sometimes, the additional requirement that terms in the relation needs to be similar up to uniform substitution is added [40], and sometimes [68, p. 516, Definition 2], only the closure by substitution—seen as a particular kind of context—is required.)

A notable example of congruence is *barbed congruence* [50, Definition 2.1.4, 61, Definition 8], which closes by context a reduction-closed relation used to observe “barbs”—the channel(s) on which a process can emit or receive. It is often taken to be *the* “reference behavioural equivalence” [50, p. 4], as it observes the interface of processes, i.e. on which channels they can interact over the time and in parallel.

But behind this apparent uniformity in the definition of congruences, the definition of contextual relations itself have often been tweaked by altering the definition of context, with no clear explanation nor justification, as we illustrate below.

In the calculus of communicating systems, notions as central as contextual bisimulation [8, pp. 223–224, Definition 421] and barbed equivalence [8, p. 224, Definition 424] considers only *static* contexts [8, p. 223, Definition 420], which are composed only of parallel composition with arbitrary term and restriction. As the author of those notes puts it himself, “the rules of the bisimulation game may be hard to justify [and] contextual bisimulation [...] is more natural” [8, p. 227]. But there is no justification—other than technical, i.e. because they “they persist after a transition” [8, p. 223]—as to *why* one should consider only some contexts in defining contextual equivalences.

In the π -calculus, contexts are defined liberally [79, p. 19, Definition 1.2.1], but still exclude contexts like e.g. $[\square] + 0$ right from the beginning. Congruences are then defined using this notion of context [79, p. 19, Definition 1.2.2], and strong barbed congruence is no exception [79, p. 59, Definition 2.1.17]. Other notions, like strong barbed equivalence [79, p. 62, Definition 2.1.20], are shown to be a non-input congruence [79, p. 63, Lemma 2.1.24], which is a notion relying on contexts that forbids the slot to occur under an input prefix [79, p. 62, Definition 2.1.22]. In other words, two notions of contexts and of congruences co-exist generally in π -calculus, but “[i]t is difficult to give rational arguments as to why one of these relations is more reasonable than the other.” [36, p. 245]

In the distributed π -calculus, contexts are restricted right from the beginning to particular operators [36, Definition 2.6]. Then, relations are defined to be contextual if they are preserved by static contexts [36, Definition 2.6], which contains only parallel composition with arbitrary terms and name binding. These contexts also appear as “configuration context” [45, p. 375] or “harness” in the ambient calculus [34, p. 372]. Static operators are deemed “sufficient for our purpose” [36, p. 37] and static contexts only are considered “[t]o keep life simple” [36, p. 38], but no further justification is given.

In the semantic theory for processes, at least in the foundational theory we would like to discuss below, one difficulty is that the class of formal theories restricted to “reduction contexts” [40, p. 448] still fall short on providing a satisfactory “formulation of semantic theories for processes which does not rely on the notion of observables or convergence”. Hence, the authors have to furthermore restrict the class of terms to *insensitive* terms [40, p. 450] to obtain a notion of *generic reduction* [40, p. 451] that allows a satisfactory definition of sound theories [40, p. 452]. Insensitive terms are essentially the collection of terms that do not interact with contexts [40, p. 451, Proposition 3.15], an analogue to λ -calculus’ genericity Lemma [13, p. 374, Proposition 14.3.24]. Here, contexts are restricted by duality: insensitive

behaviour is preserved by language contexts attests that the latter are proper functions of behaviour (i.e., if $C[\cdot]$ is provided with ‘behaviour’ then the result is still ‘behaviour’). In this sense, the congruence result can be viewed as a sanity check on the language constructors.”

terms are terms that will *not* interact with the context in which they are placed, and that need to be equated by sound theories.

Across calculi, a notion of “closing context”—that emerged from λ -calculus [8, p. 85], and matches the “wrapping” of a snippet—can be found in e.g. typed versions of the π -calculus [79, p. 479], in mobile ambient [85, p. 134], in the applied π -calculus [1, p. 7], and in the fusion calculus [51, p. 6]. Also known as “completing context” [74, p. 466], those contexts are parametric in a term, the idea being that such a context will “close” the term under study, making it amenable to tests and comparisons.

Let us try to extract some general principles from this short survey. It seems that contexts are 1. *in appearance* given access to the same operators than terms, 2. sometimes deemed to be “un-reasonable”, without always a clear justification, 3. shrunk by need, to bypass some of the difficulties they raise, or to preserve some notions, 4. sometimes picked by the term itself—typically because the same “wrapping” cannot be applied to all processes. Additionally, in all those cases, contexts are given access to a subset of operators, or restricted to contexts with particular behavior, *but never extended*. If we consider that contexts are the main tool to test the equivalence of processes, then why should the testers—or the attacker—always have access to *fewer* tools than the programmer? What reason is there not to *extend* the set of tools, of contexts, or simply take it to be orthogonal? The method we sketch below allows and actually encourages such nuances, would justify and acknowledge the restrictions we just discussed instead of adding them *passing-by*, and actually corresponds to common usage.

5 Processes, Systems and Tests

As in the λ -calculus, most concurrent calculi make a distinction between open and closed terms. For instance, the distributed π -calculus [36] implements a distinction between closed terms (called processes [36, p. 14]) and open terms, based on binding operators (input and recursion).

Most of the time, and since the origin of the calculus of communicating systems, the theory starts by considering only programs—“closed behaviour expression[s], i.e. ones with no free variable” [54, p. 73]—when comparing terms, as—exactly like in λ -calculus—they correspond to self-sufficient, well-rounded programs: it is generally agreed upon that open terms should not be released “into the wild”, as they are not able to remain in control of their internal variables, and can be subject to undesirable or uncontrolled interferences. Additionally, closed terms are also the only ones to have a *reduction semantics*, which means that they can evolve without interacting with the environment—this would corresponds, in Java, to being wrapped, i.e. inserted into a proper header and ready to be used or tested.

However, in concurrent calculi, the central notions of binders and of variables have been changing, and still seem today sometimes “up in the air”. For instance, in the original CCS, restriction was not a binder [54, p. 68], and by “refusing to admit channels as entities distinct from agents” [57, p. 16] and defining two different notions of scopes [57, p. 18], everything was set-up to produce a long and recurring confusion as to what a “closed” term meant in CCS. In the original definition of π -calculus [59, 60], there is no notion of closed terms, as every (input) binding on a channel introduces a new and free occurrence of a variable. However, the language they build upon—ECCS [28]—made this distinction clear, by separating channel constants and variables.

Once again in an attempt to mimic the “economy” [58, p. 86] of λ -calculus, but also taking inspiration from the claimed “monotheism” of the actor model [38], different notions such as values, variables, or channels have been united under the common terminology of “names”. This is at times identified as a strength, to obtain a “richer calculus in which values of many kinds may be communicated, and in which

value computations may be freely mixed with communications.” [59, p. 20] However, it seems that a distinction between those notions always needs to be carefully re-introduced when discussing technically the language [8, p. 258, Remark 493], extensions to it [1, p. 4] or possible implementations [16, p. 13, 31]. Finally, let us note that extensions of π -calculus can sometimes have different binders, as e.g. output binders are binding in the private π -calculus [67, p. 113].

In the λ -calculus, being closed is what makes a term “ready to be executed in an external environment”. But in concurrent calculi, being a closed term—no matter how it is defined—is often not enough, as it is routine to exclude e.g. terms with un-guarded operators like sum [79, p. 416] or recursion [57, p. 166]. However, these operators are sometimes not excluded from the start, even if they can never be parts of tested terms. The usual strategy [8, Remark 414, 57] is often to keep them “as long as possible”, and to exclude them only when their power cannot be tamed any more to fit the framework or prove the desired result, such as the preservation of weak bisimulation by all contexts.

In our opinion, the right distinction is not about binders of free variables, but about the role played by the syntactic objects in the theory. As “being closed” is 1. not always well-defined, or at least changing, 2. sometimes not the only condition, we would like to use the slightly more generic adjectives *complete* and *incomplete*—wrapped or not, in our Java terminology. Process algebras generally study terms by 1. completing them if needed, 2. inserting them in an environment, 3. executing them, 4. observing them thanks to predicates on the execution (“terminates”, “emitted the barb a ”, etc.), hence constructing equivalences, preorders or metrics [41] on them. Often, the environment is essentially made of another process composed in parallel with the one studied, and tweaked to improve the likeliness of observing a particular behavior: hence, we would like to think of them as tests that the observed systems has to pass, justifying the terminology we will be using.

Processes are “partial” programs, still under development; sometimes called “open terms”, they correspond to *incomplete terms*. They would be called code fragments, or snippets, in standard programming.

Systems are “configured processes”, ready to be executed in any external environment: sometimes called “closed terms”, they correspond to *complete terms*. They would be functions shipped with a library in standard programming, and ready to be executed.

Tests are defined using contexts and observations, and aims at executing and testing systems. They would be `main` methods calling a library or an API in standard programming, along with a set of observables.

Our terminology is close to the one used e.g. in ADPI [36, Chapter 5] or mobile ambients [52, Table 1], which distinguish processes and systems. In the literature of process algebra, the term “process” is commonly used to denote these three layers, possibly generating confusion. We believe this usage comes from a strong desire to keep the three layers uniform, using the same name, operators and rules, but this principle is actually constantly dented (as discussed in Sect. 4), for reasons we expose below.

6 Designing Layered Concurrent Languages

Concurrent languages could benefit from this organization from their conception:

Define processes The first step is to select a set of operators called *construction operators*, used by the programmer to write processes. Those operators should be expressive, easy to combine, with constraints as light as possible, and selected with the situation that is being modeled in mind—and not depending on whenever they fare well with not-yet-defined relations, as it is often done to privilege the guarded sum over the internal choice. To ease their usage, a “meta-syntax” can be used, something that is generally

represented by the structural equivalence. (Another interesting approach is proposed in “the π -calculus, at a distance” [4, p. 45], that bypasses the need for a structural equivalence without losing the flexibility it usually provides.)

Define deployment criteria How a process can become a system ready to be executed and tested should then be defined as a series of conditions on the binding of variables, the presence or absence of some construction operators at top-level, and even the addition of *deployment operators*, marking the process as ready to be deployed in an external environment⁴. Having a set of deployment operators that restricts, expands or intersects with the set of construction operators is perfectly acceptable, and it should enable the transformation of processes into systems and their composition.

Define tests The last step requires to define 1. a set of *testing operators*, 2. a notion of environment constructed from those operators, along with instructions on how to place a system in it, 3. a system of reduction rules regimenting how a system can execute in an environment, 4. a set of observables, i.e. a function from systems in environments to a subset of a set of atomic proposition (like “emits barb a ”, “terminates”, “contains recursion operator”, etc.).

Observe that each step uses its own set of operators and generates its own notion of context—to construct, to deploy, or to test. Tests would be key in defining notions of congruence, that would likely be reduction-closed, observational contextually-closed relations. Determining if tests should wrap processes into systems or if that should be done ahead of the test itself resonates with a long-standing debate in process algebra, and is discussed in Sect. 9.2. Note that compared to how concurrent languages are generally designed, our approach is refined along two axis: 1. every step previously exposed allows the introduction of novel operators, 2. multiple notions of systems or tests can and should co-exist in the same process algebra.

7 Addressing Existing Issues

In the process algebras literature, processes and systems often have the same structure as tests and all use the same operators and contexts, to preserve and nurture a supposedly required simplicity—at least on the surface of it. But actually, we believe that the distinction we offer is constantly used “under the hood”, without always a clear discussion, but that it captures and clarifies some of the choices, debates, improvements and explanations that have been proposed.

Co-defining observations and contexts Originally, the barb was a predicate [61, p. 690], whose definition was purely syntactic. Probably inspired by the notion of observer for testing equivalences [25, p. 91], an alternative definition was made in terms of parallel composition with a tester process [50, p. 10, Definition 2.1.3]. This illustrates perfectly how the set of observables and the notion of context are inter-dependent, and that tests should always come with a definition of observable *and* a notion of context: we believe our proposal could help in clarifying the interplay between observations and contexts. One could even imagine having a series of “contexts and observations lemmas” illustrating how certain observations can be simulated by some operators, or reciprocally.

Justifying the “silent” transition’s treatment It is routine to define relations (often called “weak”) that ignore silent (a.k.a. τ -) transitions, seen as “internal”. This sort of transitions was dubbed “unobservable internal activity” [36, p. 6] and sometimes opposed to “externally observable actions” [77, p. 230]. While we agree that “[t]his abstraction from internal differences is essential for any tractable theory

⁴Exactly like a Java method header can use keywords—*extends*, *implements*, etc.—that cannot be used in a method body.

of processes” [57, p. 3], we would also like to stress that both can and should be accommodated, and that “internal” transition should be treated as invisible *to the user*, but should still be accessible *to the programmer* when they are running their own tests.

The question “to what extent should one identify processes differing only in their internal or silent actions?” [15, p. 6] is sometimes asked, and discussed as if it was a property of the process algebra and not something that can be *internally* tuned when needed. We argue that the answer to that question is “*it depends who is asking!*”: from a user perspective, internal actions should *not* be observed, but it makes sense to let a programmer observe them when testing to help in deciding which process to prefer based on information not available to users.

Letting multiple comparisons co-exist The discussion on τ -transitions resonates with a long debate on which notion of behavioral relation is the most “reasonable”, and—still recently—a textbook can conclude a brief overview of this issue by “hop[ing] that [they] have provided enough information to [their] readers so that they can draw their own conclusions on this long-standing debate” [77, p. 160]. Sometimes, a similar question is phrased in terms of choosing the right level of abstraction to obtain meaningful language comparisons [48, Section 3]. We firmly believe that the best answer to both questions is to acknowledge that different relations and comparisons tools match different needs, and that there is no “one size fits all” answer for the needs of all the variety of testers. Of course, comparing multiple relations is an interesting and needed task [30, 33], but one should also state that multiple comparison tools can and should co-exist, and such vision will be encapsulated by the division we are proposing.

Embracing a feared distinction The distinction between our notions of processes and systems is rampant in the literature, but too often feared, as if it was a parenthesis that needed to be closed to restore some supposedly required purity and uniformity of the syntax. A good example is probably given by mobile ambients [52]. The authors start with a two-level syntax that distinguishes between processes and systems [52, p. 966]. Processes have access to strictly more constructors than systems [52, p. 967, Table 1], that are supposed to hide the threads of computation [52, p. 965]. A notion of *system context* is then introduced—as a restriction of arbitrary contexts—and discussed, and two different ways for relations to be preserved by context are defined [52, p. 969, Definition 2.2].

The authors even extend further the syntax for processes with a special \circ operator [52, p. 971, Definition 3.1], and note that the equivalences studied will not consider this additional constructor: we can see at work the distinction we sketched, where operators are added and removed based on different needs, and where the language needs not to be monolithic. The authors furthermore introduce two different reduction barbed congruences [52, p. 969, Definition 2.4]—one for systems, and one for processes, with different notions of contexts—but later on prove that they coincide on systems [52, p. 989, Theorem 6.10]. It seems to us that the distinction between processes and systems was essentially introduced for technical reasons, but that re-unifying the syntax—or at least prove that systems do not do more than processes—was a clear goal right from the start. We believe it would have been fruitful to embrace this distinction in a framework similar to the one we sketched: while retaining the interesting results already proven, maintaining this two-level syntax would allow to make a clearer distinction between the user’s and the programmer’s roles and interests, and assert that, sometimes, systems can and *should* do more than processes—for instance, interacting with users!—, and can be compared using different tools.

Keeping on extending contexts We are not the first to argue that constructors can and should be added to calculi to access better discriminatory power, but without necessarily changing the “original” language. The mismatch operator, for instance, has a similar feeling: “reasonable” testing equivalences [20, p. 280]

require it, and multiple languages [2, p. 24] use it to provide finer-grained equivalences. For technical reasons [79, p. 13], this operator is generally not part of the “core” of π -calculus, but resurfaces *by need* to obtain better equivalences: we defend a liberal use of this fruitful technique, by making a clear separation between the construction operators—added for their expressivity—and the testing operators—that improve the testing capacities.

Treating extensions as different completions It would benefit their study and usage to consider different extensions of processes algebras as different completion strategies for the same construction operators. For instance, reversible [46] or timed [87] extensions of CCS could be seen as two completion strategies—different conditions for a process to become a system—for the same class of processes, inspired from the usual CCS syntax [8, Chapter 28.1]. Those completion strategies would be suited for different needs, as one could e.g. complete a CSS process as a RCCS [24] system to test for relations such as hereditary history-preserving bisimulation [9], and then complete it with time markers as a safety-critical system. This would correspond to having multiple compilation, or deployment, strategies, based on the need, similar to “debug” and “real-time”, versions of the same piece of software. We think also of Debian’s DebugPackage, enabling generation of stack traces for any package, or of the CONFIG_PREEMPT_RT patch that converts a kernel into a real-time micro-kernel: both uses the same source code as their “casual” versions.

Obtaining fine-grained typing systems The development of typing systems for concurrent programming languages is a notoriously difficult topic. Some results in π -calculus have been solidified [79, Part III], but diverse difficulties remain. Among them, the co-existence of multiple systems for e.g. session types [37], the difficulty to tie them precisely to other type systems as Linear Logic [23], and the doubts about the adaptation of the “proof-as-program” paradigm in a concurrent setting [14], make this problem active and diverse. The ultimate goal seems to find a typing system that would accommodate different uses and scenarios that are not necessarily comparable.

Using our proposal, one could imagine easing this process by developing two different typing systems, one aimed at programmers—to track bugs and produce meaningful error messages—and one aimed at users—to track security leaks or perform user-input validation. Once again, having a system developed along the layers we recommend would allow to have e.g. a type system for processes only, and to erase the information when completing the process, so that the typing discipline would be enforced only when the program is being developed, but not executed. This is similar to arrays of parameterized types in Java [66, pp. 253–258], that checks the typing discipline at compilation time, but not at run-time.

While this series of examples and references illustrates how our proposal could clarify pre-existing distinctions, we would like to stress that 1. nothing prevents from collapsing our distinction when it is not needed, 2. additional progresses could be made using it, as we sketch in the next section.


8 Exploiting Context Awareness

We would like to sketch below some possible exploitations of our frame that we believe could benefit the study and expressivity of some popular concurrent languages.

For CCS, we sketch below two possible improvements, the second being related to security.

Testing for auto-concurrency Auto-concurrency (a.k.a. auto-parallelism) is when a system has two different transitions—leading to different states—labeled with the same action [64, p. 391, Definition 5]. Systems with auto-concurrency are sometimes excluded as non-valid terms [26, p. 155] or simply not considered in particular models [65, p. 531], as the definition of bisimulation is problematic for them.

Consider e.g. the labeled configuration structures (a.k.a. stable family [86, Section 3.1]) on the right, where the label of the event executed is on the edge and configurations are represented with \circ . Non-interleaving models of concurrency [80] distinguishes between them, as “true concurrency models” would.



Some forms of “back-and-forth-bisimulations” cannot discriminate between them if $a = b$ [71]. While not being able to distinguish between those two terms may make sense from an external—user’s—point of view, we argue that a programmer should have access to an internal mechanism that could answer the question “*Can this process perform two barbs with the same label at the same time?*”. Such an observation—possibly coupled with a testing operator—would allow to distinguish between e.g. $!a.P \mid !a.P$ and $!a.P$, that are generally taken to be bisimilar, and would re-integrate auto-concurrent systems—that are, after all, unjustifiably excluded—in the realm of comparable systems.

Representing man-in-the-middle One could add to the testing operators an operator $\nabla a.P$, which would forbid P to act silently on channel a . This novel operator would add the possibility for the environment to “spy” on a determined channel, as if the environment was controlling (part of) the router of the tested system. One could then reduce “normally” in a context $\nabla a[\square]$ if the channel is still secure:

$$\nabla a(b.Q \mid \bar{b}.P) \rightarrow^\tau \nabla a(Q \mid P) \quad (\text{If } a \neq b)$$


But in the case where $a = b$, the environment could intercept the communication and then decide to forward, prevent, or alter it. Adding this operator to the set of testing operators would for instance open up the possibility of interpreting $\nabla a(P)$ as an operation securing the channel a in P , enabling the study of relations \sim that could include e.g.

$$\begin{aligned} \nabla a(\nabla a(P \mid Q) \sim \nabla a(\nabla b(P[a/b] \mid Q[a/b])) & \quad (\text{For } b \text{ not in the free names of } P \text{ nor } Q) \\ \nabla a(\nabla a(P \mid Q)) \sim \nabla a(P \mid Q) & \quad (\text{Uselessness of securing a hacked channel}) \end{aligned}$$

While the first rule enforces that, once secured, channel names are α -equivalent (the process can decide to migrate to a different channel without being spied on), the second illustrates that, once a channel is tapped, a process cannot obtain any confidentiality on it anymore.

In π -calculus, all tests are instantiating contexts (in the sense that the term tested needs to be either already closed, or to be closed by the context), and all instantiating contexts use only construction operators, and hence are “construction contexts”. This situation corresponds to Situation A in Figure 1. We believe the picture could be much more general, with tests having access to *more constructors*, and not needing to be instantiating—in the sense that completion can be different from closedness—, so that we would obtain Situation B in Figure 1. While we believe this remark applies to most of the process algebras we have discussed so far, it is particularly salient in π -calculus, where the match and mismatch operators have been used “to internalize a lot of meta theory” [32, p. 57], standing “inside” the “Construction operators” circle while most authors seem to agree that they would prefer not to add it to the internals of the language⁵. It should also be noted that the mismatch operator—in its “intuitionistic” version—furthermore “tried to escape the realm of instantiating contexts” by being tightly connected [42] to *quasi-open bisimilarities* [78, p. 300, Definition 6], which is a subtle variation on how substitutions can be applied by context to the terms being tested.

⁵To be more precise: while “most occurrences of matching can be encoded by parallel composition [...] mismatching cannot be encoded in the original π -calculus” [69, p. 526], which makes it somehow suspicious.

Figure 1: Opening up the testing capacities of π -calculus

Having a notion of “being complete” not requiring closedness could be useful when representing distributed programming, where “one often wants to send a piece of code to a remote site and execute it there. [...] [T]his feature will greatly enhance the expressive power of distributed programming[by] send[ing] an open term and to make the necessary binding at the remote site.” [35, p. 250] We believe that maintaining the possibility of testing “partially closed”—but still complete—terms would enable a more theoretical understanding of distributed programming and remote compilation.

Distributed π -calculus, could explore the possible differences between two parallelisms: between threads in the same process—in the Unix sense—and between units of computation. Such a distinction could be rephrased thanks to two parallel operators, one on processes and the other on systems. Such a distinction would allow to observationally distinguish e.g. the execution of a program with two threads on a dual-core computer and the execution of two single thread programs on two single-core computers.

For cryptographic protocols, we could imagine representing encryption of data as a special context $\mathcal{E}[\Box]$ that would transform a process P into an encrypted system $\mathcal{E}[P]$, and make it un-executable unless “plugged” in an environment $\mathcal{D}[\Box]$ that could decrypt it. This could allow the applied π -calculus [1] to become more expressive and to be treated as a decoration of the pure π -calculus more effectively. This could also, as the authors wish, make “the formalization of attackers as contexts [...] continue to play a role in the analysis of security protocols” [1, p. 35].

Recent progresses in the field of verification of cryptographic protocols [12] hinted in this direction as well. By taking “[t]he notion of test [to] be relative to an environment” [12, p. 12], a formal development involving “frames” [12, Definition 2.3] can emerge and give flesh to some ideas expressed in our proposal. It should be noted that this work also “enrich[...] processes with a success construct” [12, p. 12], that cannot be used to construct processes, to construct “experiments”.

9 Concluding Remarks

We conclude by discussing related approaches, by casting a new light on a technical issue related to barbed congruences, by offering the context lemmas a new interpretation, and by coming back to our motivations.

9.1 An Approved and Promising Perspective

We would like to stress that our proposal resonates with previous comments, and should not be treated as an isolated historical perspective that will have no impact on the future.

In the study of process algebras, in addition to the numerous hints toward our formalism that we already discussed, there are at least two instances when the power of the “testing suite” was explicitly

discussed [75, Remark 5.2.21]. In a 1981 article, it is assumed that “by varying the ambient (‘weather’) conditions, an experimenter” [55, p. 32] can observe and discriminate better than a simple user could. Originally, this idea seemed to encapsulate two orthogonal dimensions: the first was that the tester could run the tested system any number of times, something that would now be represented by the addition of the replication operator $!$ to the set of testing operators. The second was that the tester could enumerate all possible non-deterministic transitions of the tested system. This second dimension gave birth to “a language for testing concurrent processes” [49, p. 1] that is more powerful than the language used to write the programs being tested. In this particular example, the tester has access to a termination operator and probabilistic features that are not available to the programmer: as a result, the authors “may distinguish non-bisimilar processes through testing” [49, p. 19].

Looking forward, the vibrant field of secure compilation made a clear-cut distinction between “target language contexts” representing adversarial code and programmers’ “source context” to explore property preservation of programs [3]. This perspective was already partially at play in the spi calculus for cryptographic protocols [2, p. 1], where the attacker is represented as the “environment of a protocol”. We believe that both approaches—coming from the secure compilation, from the concurrency community, but also from other fields—concur to the same observation that the environment—formally captured by a particular notion of context—deserves an explicit and technical study to model different interactions with systems, and need to be detached from “construction” contexts.

9.2 When Should Contexts Come into Play?

The interesting question of *when* to use contexts when testing terms [79, pp. 116–117, Section 2.4.4] raises a technical question that is put under a different perspective by our analysis. Essentially, the question is whether the congruences under study should being *defined* as congruences (e.g. reduction-closed barbed congruence [79, p. 116]), or being defined in two steps, i.e. as the contextual closure of a pre-existing relation (e.g. strong barbed congruence [79, p. 61, Definition 2.1.17], which is the contextual closure of strong barbed bisimilarity [79, p. 57, Definition 2.1.7])?

Indeed, bisimulations can be presented as an “interaction game” [82] generally played as 1. Pick an environment for both terms (i.e., complete them, then embed them in the same testing environment), 2. Have them “play” (i.e. have them try to match each other’s step). But a more dynamic version of the game let picking an environment *be part of the game*, so that each process can not only pick the next step, *but also in which environment it needs to be performed*. This version of the game, called “dynamic observational congruence” [62], provides a better software modularity and reusability, as it allows to study the similarity of terms that can be re-configured “on the fly”. Embedding the contexts in the definitions of the relations is a strategy that was also used to obtain behavioral characterization of theories [40, p. 455, Proposition 3.24], and that corresponds to open bisimilarities [73, p. 77, Proposition 3.12]

Those two approaches have been extensively compared and studied—still are [1, p. 24]—but to our knowledge they rarely co-exist, as if one had to take a side at the early stage of the language design, instead of letting the tester decide later on which approach is best suited for what they wish to observe. We argue that both approaches are equally valid, *provided we acknowledge they play different roles*.

This question of *when are the terms completed?* can be rephrased as *what is it that you are trying to observe?*, or even *who is completing them?*: is the completion provided by the programmer, once and for all, or is the tester allowed to explore different completions and to change them as the tests unfold? Looking back at our Java example from Sect. 2, this corresponds to letting the tester repeatedly tweak the parameter or return type of the wrapping from `int` to `long`, allowing them to have finer comparisons between snippets. In this frame, moving from the *static* definition of congruence to *dynamic* one would

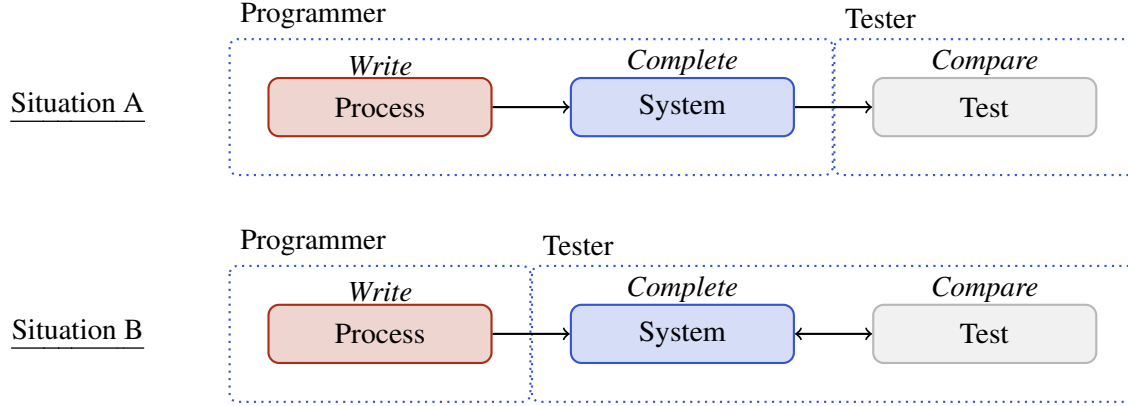


Figure 2: Distinguishing between completing strategies

corresponds to going from Situation A to Situation B in Figure 2. This illustrates two aspects worth highlighting:

1. Playing on the variation “*should I complete the terms before or during their comparison?*” is not simply a technical question, but reflects a choice between two different situations equally interesting.
2. This choice can appeal to different notions of systems, completions and tests: for instance, while completing a process before testing it (Situation A) may indeed be needed when the environment represents an external deployment platform, it makes less sense if we think of the environment as part of the development workflow, in charge of providing feedback to the programmer or as a powerful attacker than can manipulate the conditions in which the process is executed (Situation B).

If completion is seen as compilation, this opens up the possibility of studying how the bindings performed *by the user*, on *their* particular set-up, during a *remote* compilation, can alter a program. One can then compare different relations—some comparing source code’s releases, some comparing binaries’ releases—to get a better, fuller, picture of the program.

9.3 Penetrating Context Lemmas’ Meanings

What is generally refereed to as *the* context lemma⁶ is actually a series of results stating that considering all the operators when constructing the context for a congruence may not be needed. For instance, it is equivalent to define the barbed congruence [79, p. 95, Definition 2.4.5] as the closure of barbed bisimilarity under all context, or only under contexts of the form $[\Box]\sigma \mid P$ for all substitution σ and term P . In its first version [72, p. 432, Lemma 5.2.2], this lemma had additional requirements e.g. on sorting contexts, but the core idea is always the same: “*there is no need to consider all contexts to determine if a relation is a congruence, you can consider only contexts of a particular form*”.

The “flip side” of the context lemma is what we would like to call the “anti-context pragmatism”: whenever a particular type of operator or context prevents a relation from being a congruence, it is tempting to simply exclude it. For instance, contexts like $[\Box] + 0$ are routinely removed—as we discussed in Sect. 4—to define the barbed congruence of π -calculus, or contexts were restricted to what is called harnesses in the mobile ambients calculus [34] before proving such results. As strong bisimulation [68, p. 514,

⁶At least, in process algebra, as the same name is used for a different type of meaning in e.g. λ -calculus [53, p. 6].

Definition 1] is not preserved by input prefix [68, p. 515, Proposition 4] but is by all the other operators, it is sometimes tempting to simply remove input prefix from the set of constructors allowed at top-level in contexts, which is what non-input contexts [79, p. 62, Definition 2.1.22] do, and then to establish a context lemma for this limited notion of context.

Taken together, those two remarks produce a strange impression: while it is mathematically elegant and interesting to prove that weaker conditions are enough to satisfy an interesting property, it seems to us that this result is sometimes “forced” into the process algebra by having ahead of time excluded all the operators that would not fit, hence producing a result that is not only weaker, but also somehow artificial, or even tautological. Furthermore the criteria of “not adding any discriminating power” should not be a positive criterion when deciding if a context belongs to the algebra: on the opposite, one would want contexts to *increase* the discriminating power—as for the mismatch operator—and not to “conform” to what substitution and parallel composition have already decided.

Context lemmas seem to embrace an uncanny perspective: instead of being used to prove properties about tests more easily, they should be considered from the perspective of the ease of use of testing systems. Stated differently, we believe that the set of testing operators should come first, and then *then*, if the language designer wishes to add operators to ease the testers’ life, they can do so providing they obtain a context lemma proving that those operators do not alter the original testing capacities. Once again, varying the testing suite is perfectly acceptable, but once fixed, *the context lemma is simply present to show that adding some testing operators is innocent, that it will simply make testing certain properties easier.*

9.4 Embracing the Diversity

Before daring to submit a non-technical paper, we tried to conceive a technical construction that could convey our ideas. In particular we tried to build a syntactic (even categorical) meta-theory of processes, systems and tests. We wanted to define congruences in this meta-theory, and to answer the following question: what could be the minimal requirements on contexts and operators to prove a generic form of context lemma for concurrent languages?

However, as the technical work unfolded, we realized that the definitions of contexts, observations, and operators, were so deeply interwoven that it was nearly impossible to extract any general or useful principle. Context lemmas use specific features of languages, in a narrow sense, as for instance no context lemma can exist in the “Situation B” of Figure 2 [79, p. 117], and we were not able to find a unifying framework. This also suggests that context lemmas are often *fit* for particular process algebras *by chance*, and dependent intrinsically of the language considered, for no deep reasons.

This was also liberating, as all the nuances of languages we had been fighting against started to form a regular pattern: every single language we considered exhibited (at least parts of) the structure we sketched in the present proposal. Furthermore, our framework was a good lens to read and answer some of the unspoken questions suggested in the margin or the footnotes—but rarely upfront—of the multiple research papers, lecture notes and books we consulted. So, even without mathematical proofs, we consider this contribution a good way of stirring the community, and to question the traditional wisdom.

It seems indeed to us that there is nothing but benefits in altering the notion of context, as it is actually routine to do so, even recently [39], and that stating the variations used will only improve the expressiveness of the testing capacities and the clarity of the exposition.

It is a common trope to observe the immense variety of process calculi, and to sometimes wish there could be a common formalism to capture them all—to this end, *the* π -calculus is often considered the best candidate. Acknowledging this diversity is already being one step ahead of the λ -calculus—that keeps

forgetting that there is more than one λ -calculus, depending on the evaluation strategy and on features such as sharing [5]—and this proposal encourages to push the decomposition into smaller languages even further, as well as it encourages to see whole theories as simple “completion” of standard languages. As we defended, breaking the monolithic status of context⁷ will actually make the theory and presentation follow more closely the technical developments, and liberate from the goal of having to find *the* process algebra with *its unique* observation technique that would capture all possible needs.

10 Postscript

The tone and scope of this essay may seem “timeless”, but discussions during the *14th Interaction and Concurrency Experience* (particularly with Alceste Scalas and Emilio Tuosto) and with colleagues (among which David Baelde and Ross Horne) finished anchoring it to very concrete and timely questions. We would like to briefly sketch a handful of future themes, hoping to faithfully record the fascinating discussions that led to them.

1. Semantical approaches to security revealed multiple small gaps that are hard to explain or solve with usual notions of contexts and equivalences. For instance, the study of equivalence properties led to distinguish between semantics where all the communications have to be made via the environment (i.e., the attacker) and semantics where communications inside the process and communications between the process and its environment are treated differently. The (surprising) result [11] is that while both treatments coincide for reachability properties, they are incomparable for indistinguishability properties. The resulting *classical*, *private* and *eavesdropping* semantics each yield their own may-testing and observational equivalences, that could possibly be elegantly captured by our frame. The interested reader should refer to a recent discussion [12, Section 4.1.6] for more insight on this issue.
2. Our discussion could also be used to bridge possible gaps between process algebras and game semantics, or at least introduce a gradation between them. Indeed, while “usual” algebraic contexts are extremely passive and with limited (static) knowledge of the term placed inside of them, strategies generally have all laterality to evolve, and to store information about the play unfolding. We would like to argue, following a classical paper [44, Section 5], that maybe processes (or, in our case, contexts) could represent strategies, and that there could be a gradation in the “dynamicity” of contexts: from very statical ones to more dynamical and reactive ones. This could be used e.g. to study the open barbed bisimilarity between processes with free variables, by allowing contexts to use and retain information about permitted submissions.
3. Reversible CCS (RCCS) [24] and CCS with keys (CCSK) [70] are two extensions to CCS aiming at formalizing reversible concurrent computation. They have been compared and studied, but, as it turns out, they actually are the two faces of the same coin [46]. However, when defining contexts, and despite the similarities in the definitions, two recent studies [10, 47] came to exactly opposite conclusions: while the RCCS-inspired system seems to provide definitive evidence that *no relation can be a congruence if the context can change the history of the process* [10, Theorem 2], the

⁷This may be a good place to mention that this monolithic status probably comes in part from the original will of making e.g. CCS a programming *and* specification language—an original perspective that was reminded to us by Ilaria Castellani, who we wish to thank. The specification was supposed to be the program itself, that would be easy to check for correctness: the goal was to make it “possible to describe existing systems, to specify and program new systems, and to argue mathematically about them, all without leaving the notational framework of the calculus” [56, p. 1]. This original research project slightly shifted—from specifying programs to specifying behaviors—but that original perspective remained.

CCSK approach came to the conclusion that a particular barbed bisimilarity was a congruence [47, Corollary 5.12]. Does the difference rest only on the definition of context, or is there a more subtle distinction at work? It is our hope that the approach sketched here could help to solve this mystery.

Last but not least, an “easy” exercise to illustrate the relevance of our perspective would be to study a Morris-style equivalence of the programming language Java, using the convenient, “simplified”, Featherweight Java [43]. It is our hope that this lead, on which we will be working in the near future, will fruitfully illustrate the benefits of the distinction we detailed in this paper.

Acknowledgements The author wish to thank the organizer of ICE 2021 for organizing this welcoming, open workshop, as well as the reviewers who kindly shared their comments, suggestions and insights with us. This paper benefited a lot from them, as well as from the discussions mentioned in Sect. 10.

References

- [1] Martín Abadi, Bruno Blanchet & Cédric Fournet (2018): *The Applied Pi Calculus: Mobile Values, New Names, and Secure Communication*. *J. ACM* 65(1), pp. 1:1–1:41, doi:10.1145/3127586.
- [2] Martín Abadi & Andrew D. Gordon (1999): *A Calculus for Cryptographic Protocols: The spi Calculus*. *Inf. Comput.* 148(1), pp. 1–70, doi:10.1006/inco.1998.2740.
- [3] Carmine Abate, Roberto Blanco, Deepak Garg, Catalin Hritcu, Marco Patrignani & Jérémy Thibault (2019): *Journey Beyond Full Abstraction: Exploring Robust Property Preservation for Secure Compilation*. In: *32nd IEEE Computer Security Foundations Symposium, CSF 2019, Hoboken, NJ, USA, June 25-28, 2019*, IEEE, pp. 256–271, doi:10.1109/CSF.2019.00025.
- [4] Beniamino Accattoli (2013): *Evaluating functions as processes*. In Rachid Echahed & Detlef Plump, editors: *TERMGRAPH 2013, EPTCS* 110, pp. 41–55, doi:10.4204/EPTCS.110.6.
- [5] Beniamino Accattoli (2019): *A Fresh Look at the lambda-Calculus (Invited Talk)*. In Herman Geuvers, editor: *CSL, LIPIcs* 131, Schloss Dagstuhl, pp. 1:1–1:20, doi:10.4230/LIPIcs.FSCD.2019.1.
- [6] Beniamino Accattoli & Ugo Dal Lago (2014): *Beta reduction is invariant, indeed*. In Thomas A. Henzinger & Dale Miller, editors: *CSL, ACM*, p. 8, doi:10.1145/2603088.2603105.
- [7] Beniamino Accattoli & Ugo Dal Lago (2012): *On the Invariance of the Unitary Cost Model for Head Reduction*. In Ashish Tiwari, editor: *23rd International Conference on Rewriting Techniques and Applications (RTA'12), RTA 2012, May 28 - June 2, 2012, Nagoya, Japan, LIPIcs* 15, Schloss Dagstuhl, pp. 22–37, doi:10.4230/LIPIcs.RTA.2012.22.
- [8] Roberto M. Amadio (2016): *Operational methods in semantics*. Lecture notes, Université Denis Diderot Paris 7. Available at <https://hal.archives-ouvertes.fr/ce1-01422101>.
- [9] Clément Aubert & Ioana Cristescu (2020): *How Reversibility Can Solve Traditional Questions: The Example of Hereditary History-Preserving Bisimulation*. In Igor Konnov & Laura Kovács, editors: *CONCUR, LIPIcs* 2017, Schloss Dagstuhl, pp. 13:1–13:24, doi:10.4230/LIPIcs.CONCUR.2020.13.
- [10] Clément Aubert & Doriana Medić (2021): *Explicit Identifiers and Contexts in Reversible Concurrent Calculus*. In Shigeru Yamashita & Tetsuo Yokoyama, editors: *Reversible Computation - 13th International Conference, RC 2021, Virtual Event, July 7-8, 2021, Proceedings, LNCS* 12805, Springer, pp. 144–162, doi:10.1007/978-3-030-79837-6_9.
- [11] Kushal Babel, Vincent Cheval & Steve Kremer (2020): *On the semantics of communications when verifying equivalence properties*. *Journal of Computer Security* 28, pp. 71–127, doi:10.3233/JCS-191366. 1.
- [12] David Baelde (2021): *Contributions à la Vérification des Protocoles Cryptographiques*. Habilitation à diriger des recherches, Université Paris-Saclay. Available at http://www.lsv.fr/~baelde/hdr/habilitation_baelde.pdf.
- [13] Hendrik Pieter Barendregt (1984): *The Lambda Calculus – Its Syntax and Semantics*. *Studies in Logic and the Foundations of Mathematics* 103, North-Holland, doi:10.1016/B978-0-444-87508-2.50006-X.
- [14] Emmanuel Beffara & Virgile Mogbil (2012): *Proofs as executions*. In Jos C. M. Baeten, Thomas Ball & Frank S. de Boer, editors: *IFIP TCS, LNCS* 7604, Springer, pp. 280–294, doi:10.1007/978-3-642-33475-7_20.
- [15] Jan A. Bergstra, Alban Ponse & Scott A. Smolka, editors (2001): *Handbook of Process Algebra*. Elsevier Science, Amsterdam, doi:10.1016/B978-0-444-82830-9/50017-5.
- [16] Bruno Blanchet (2016): *Modeling and Verifying Security Protocols with the Applied Pi Calculus and ProVerif*. *Foundations and Trends in Privacy and Security* 1(1-2), pp. 1–135, doi:10.1561/33000000004.
- [17] Olivier Bodini (2021): Personal communication.
- [18] Mirna Bognar (2002): *Contexts in Lambda Calculus*. Ph.D. thesis, Vrije Universiteit Amsterdam. Available at https://www.cs.vu.nl/en/Images/bognar_thesis_tcm210-92584.pdf.
- [19] Mirna Bognar & Roel C. de Vrijer (2001): *A Calculus of Lambda Calculus Contexts*. *J. Autom. Reasoning* 27(1), pp. 29–59, doi:10.1023/A:1010654904735.

- [20] Michele Boreale & Rocco De Nicola (1995): *Testing Equivalence for Mobile Processes*. *Inf. Comput.* 120(2), pp. 279–303, doi:10.1006/inco.1995.1114.
- [21] Flavien Breuvert (2015): *Dissecting denotational semantics*. Ph.D. thesis, Université Paris Diderot — Paris VII. Available at https://www.lipn.univ-paris13.fr/~breuvert/These_breuvert.pdf.
- [22] Antonio Bucciarelli, Alberto Carraro, Thomas Ehrhard & Giulio Manzonetto (2011): *Full Abstraction for Resource Calculus with Tests*. In Marc Bezem, editor: *CSL, LIPIcs* 12, Schloss Dagstuhl, Dagstuhl, Germany, pp. 97–111, doi:10.4230/LIPIcs.CSL.2011.97.
- [23] Luís Caires, Frank Pfenning & Bernardo Toninho (2016): *Linear logic propositions as session types*. *MSCS* 26(3), pp. 367–423, doi:10.1017/S0960129514000218.
- [24] Vincent Danos & Jean Krivine (2004): *Reversible Communicating Systems*. In Philippa Gardner & Nobuko Yoshida, editors: *CONCUR, LNCS* 3170, Springer, pp. 292–307, doi:10.1007/978-3-540-28644-8_19.
- [25] Rocco De Nicola & Matthew Hennessy (1984): *Testing Equivalences for Processes*. *Theor. Comput. Sci.* 34, pp. 83–133, doi:10.1016/0304-3975(84)90113-0.
- [26] Rocco De Nicola, Ugo Montanari & Frits W. Vaandrager (1990): *Back and Forth Bisimulations*. In Jos C. M. Baeten & Jan Willem Klop, editors: *CONCUR '90, LNCS* 458, Springer, pp. 152–165, doi:10.1007/BFb0039058.
- [27] Edsger W. Dijkstra (1968): *Letters to the editor: go to statement considered harmful*. *Commun. ACM* 11(3), pp. 147–148, doi:10.1145/362929.362947.
- [28] Uffe Engberg & Mogens Nielsen (2000): *A calculus of communicating systems with label passing - ten years after*. In Gordon D. Plotkin, Colin Stirling & Mads Tofte, editors: *Proof, Language, and Interaction, Essays in Honour of Robin Milner*, The MIT Press, pp. 599–622.
- [29] Claudia Faggian & Simona Ronchi Della Rocca (2019): *Lambda Calculus and Probabilistic Computation*. In: *LICS*, IEEE, pp. 1–13, doi:10.1109/LICS.2019.8785699.
- [30] Cédric Fournet & Georges Gonthier (2005): *A hierarchy of equivalences for asynchronous calculi*. *J. Log. Algebr. Methods Program.* 63(1), pp. 131–173, doi:10.1016/j.jlap.2004.01.006.
- [31] Simon Fowler, Sam Lindley & Philip Wadler (2017): *Mixing Metaphors: Actors as Channels and Channels as Actors*. In Peter Müller, editor: *ECOOP 2017, LIPIcs* 74, Schloss Dagstuhl, pp. 11:1–11:28, doi:10.4230/LIPIcs.ECOOP.2017.11.
- [32] Yuxi Fu & Zhenrong Yang (2003): *Tau laws for pi calculus*. *Theor. Comput. Sci.* 308(1-3), pp. 55–130, doi:10.1016/S0304-3975(03)00202-0.
- [33] Robert J. van Glabbeek (1993): *The Linear Time - Branching Time Spectrum II*. In Eike Best, editor: *CONCUR '93, LNCS* 715, Springer, pp. 66–81, doi:10.1007/3-540-57208-2_6.
- [34] Andrew D. Gordon & Luca Cardelli (2003): *Equational Properties Of Mobile Ambients*. *MSCS* 13(3), pp. 371–408, doi:10.1017/S0960129502003742.
- [35] Masatomo Hashimoto & Atsushi Ohori (2001): *A typed context calculus*. *Theor. Comput. Sci.* 266(1-2), pp. 249–272, doi:10.1016/S0304-3975(00)00174-2.
- [36] Matthew Hennessy (2007): *A distributed Pi-calculus*. CUP, doi:10.1017/CBO9780511611063.
- [37] Bas van den Heuvel & Jorge A. Pérez (2020): *Session Type Systems based on Linear Logic: Classical versus Intuitionistic*. In Stephanie Balzer & Luca Padovani, editors: *PLACES@ETAPS 2020, EPTCS* 314, pp. 1–11, doi:10.4204/EPTCS.314.1.
- [38] Carl Hewitt, Peter Boehler Bishop, Irene Greif, Brian Cantwell Smith, Todd Matson & Richard Steiger (1973): *Actor Induction and Meta-Evaluation*. In Patrick C. Fischer & Jeffrey D. Ullman, editors: *POPL*, ACM Press, pp. 153–168, doi:10.1145/512927.512942.
- [39] Daniel Hirschhoff, Enguerrand Prebet & Davide Sangiorgi (2020): *On the Representation of References in the Pi-Calculus*. In Igor Konnov & Laura Kovács, editors: *CONCUR, LIPIcs* 2017, Schloss Dagstuhl, pp. 34:1–34:20, doi:10.4230/LIPIcs.CONCUR.2020.34.

- [40] Kohei Honda & Nobuko Yoshida (1995): *On Reduction-Based Process Semantics*. *Theor. Comput. Sci.* 151(2), pp. 437–486, doi:10.1016/0304-3975(95)00074-7.
- [41] Eiichi Horita & Ken Mano (1997): *A Metric Semantics for the π -Calculus Extended with External Events*. *Kôkyûroku* 996, pp. 67–81. Available at <http://hdl.handle.net/2433/61239>.
- [42] Ross Horne, Ki Yung Ahn, Shang-Wei Lin & Alwen Tiu (2018): *Quasi-Open Bisimilarity with Mismatch is Intuitionistic*. In Anuj Dawar & Erich Grädel, editors: *LICS*, ACM, pp. 26–35, doi:10.1145/3209108.3209125.
- [43] Atsushi Igarashi, Benjamin C. Pierce & Philip Wadler (2001): *Featherweight Java: A Minimal Core Calculus for Java and GJ*. *ACM Trans. Program. Lang. Syst.* 23(3), p. 396–450, doi:10.1145/503502.503505.
- [44] Jim Laird (2005): *A Game Semantics of the Asynchronous π -Calculus*. In Martín Abadi & Luca de Alfaro, editors: *CONCUR*, LNCS 3653, Springer, pp. 51–65, doi:10.1007/11539452.8.
- [45] Ivan Lanese, Michael Lienhardt, Claudio Antares Mezzina, Alan Schmitt & Jean-Bernard Stefani (2013): *Concurrent Flexible Reversibility*. In Matthias Felleisen & Philippa Gardner, editors: *ESOP*, LNCS 7792, Springer, pp. 370–390, doi:10.1007/978-3-642-37036-6_21.
- [46] Ivan Lanese, Dorian Medić & Claudio Antares Mezzina (2019): *Static versus dynamic reversibility in CCS*. *Acta Inform.*, doi:10.1007/s00236-019-00346-6.
- [47] Ivan Lanese & Iain Phillips (2021): *Forward-Reverse Observational Equivalences in CCSK*. In Shigeru Yamashita & Tetsuo Yokoyama, editors: *Reversible Computation - 13th International Conference, RC 2021, Virtual Event, July 7-8, 2021, Proceedings*, LNCS 12805, Springer, pp. 126–143, doi:10.1007/978-3-030-79837-6_8.
- [48] Ivan Lanese, Cátia Vaz & Carla Ferreira (2010): *On the Expressive Power of Primitives for Compensation Handling*. In Andrew D. Gordon, editor: *ESOP*, LNCS 6012, Springer, pp. 366–386, doi:10.1007/978-3-642-11957-6_20.
- [49] Kim Guldstrand Larsen & Arne Skou (1991): *Bisimulation through Probabilistic Testing*. *Inf. Comput.* 94(1), pp. 1–28, doi:10.1016/0890-5401(91)90030-6.
- [50] Jean-Marie Madiot (2015): *Higher-order languages: dualities and bisimulation enhancements*. Ph.D. thesis, École Normale Supérieure de Lyon, Università di Bologna. Available at <https://hal.archives-ouvertes.fr/tel-01141067>.
- [51] Massimo Merro (1998): *On the Expressiveness of Chi, Update, and Fusion calculi*. In Ilaria Castellani & Catuscia Palamidessi, editors: *EXPRESS*, *Electron. Notes Theor. Comput. Sci.* 16, Elsevier, pp. 133–144, doi:10.1016/S1571-0661(04)00122-7.
- [52] Massimo Merro & Francesco Zappa Nardelli (2005): *Behavioral theory for mobile ambients*. *J. ACM* 52(6), pp. 961–1023, doi:10.1145/1101821.1101825.
- [53] Robin Milner (1977): *Fully abstract models of typed λ -calculi*. *Theor. Comput. Sci.* 4(1), pp. 1–22, doi:10.1016/0304-3975(77)90053-6.
- [54] Robin Milner (1980): *A Calculus of Communicating Systems*. LNCS, Springer-Verlag, doi:10.1007/3-540-10235-3.
- [55] Robin Milner (1981): *A Modal Characterisation of Observable Machine-Behaviour*. In Egidio Astesiano & Corrado Böhm, editors: *CAAP '81, Trees in Algebra and Programming, 6th Colloquium, Genoa, Italy, March 5-7, 1981, Proceedings*, LNCS 112, Springer, pp. 25–34, doi:10.1007/3-540-10828-9_52.
- [56] Robin Milner (1986): *A Calculus of Communicating Systems*. LFCS Report Series ECS-LFCS-86-7, The University of Edinburgh. Available at <http://www.lfcs.inf.ed.ac.uk/reports/86/ECS-LFCS-86-7/>.
- [57] Robin Milner (1989): *Communication and Concurrency*. PHI Series in computer science, Prentice-Hall.
- [58] Robin Milner (1993): *Elements of Interaction: Turing Award Lecture*. *Commun. ACM* 36(1), p. 78–89, doi:10.1145/151233.151240.
- [59] Robin Milner, Joachim Parrow & David Walker (1992): *A Calculus of Mobile Processes, I*. *Inf. Comput.* 100(1), pp. 1–40, doi:10.1016/0890-5401(92)90008-4.

- [60] Robin Milner, Joachim Parrow & David Walker (1992): *A Calculus of Mobile Processes, II*. *Inf. Comput.* 100(1), pp. 41–77, doi:10.1016/0890-5401(92)90009-5.
- [61] Robin Milner & Davide Sangiorgi (1992): *Barbed Bisimulation*. In Werner Kuich, editor: *ICALP, LNCS 623*, Springer, pp. 685–695, doi:10.1007/3-540-55719-9.114.
- [62] Ugo Montanari & Vladimiro Sassone (1992): *Dynamic congruence vs. progressing bisimulation for CCS*. *Fund. Inform.* 16(1), pp. 171–199, doi:10.3233/FI-1992-16206.
- [63] Aleksandar Nanevski, Frank Pfenning & Brigitte Pientka (2008): *Contextual Modal Type Theory*. *ACM Trans. Comput. Log.* 9(3), doi:10.1145/1352582.1352591.
- [64] Mogens Nielsen & Christian Clausen (1994): *Bisimulation for Models in Concurrency*. In Bengt Jonsson & Joachim Parrow, editors: *CONCUR '94, LNCS 836*, Springer, pp. 385–400, doi:10.1007/BFb0015021.
- [65] Mogens Nielsen, Uffe Engberg & Kim S. Larsen (1989): *Fully abstract models for a process language with refinement*. In J. W. de Bakker, Willem P. de Roever & Grzegorz Rozenberg, editors: *Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency, School/Workshop, Noordwijkerhout, The Netherlands, May 30 - June 3, 1988, Proceedings, LNCS 354*, Springer, pp. 523–548, doi:10.1007/BFb0013034.
- [66] Patrick Niemeyer & Daniel Leuck (2013): *Learning Java*, 4th edition. O'Reilly Media, Incorporated.
- [67] Catuscia Palamidessi & Frank D. Valencia (2005): *Recursion vs Replication in Process Calculi: Expressiveness*. *Bull. EATCS* 87, pp. 105–125. Available at <http://eatcs.org/images/bulletin/beatcs87.pdf>.
- [68] Joachim Parrow (2001): *An Introduction to the π -Calculus*. In Jan A. Bergstra, Alban Ponse & Scott A. Smolka, editors: *Handbook of Process Algebra*, North-Holland / Elsevier, pp. 479–543, doi:10.1016/b978-044482830-9/50026-6.
- [69] Joachim Parrow & Davide Sangiorgi (1993): *Algebraic Theories for Name-Passing Calculi*. In J. W. de Bakker, Willem P. de Roever & Grzegorz Rozenberg, editors: *A Decade of Concurrency, Reflections and Perspectives, REX School/Symposium, Noordwijkerhout, The Netherlands, June 1-4, 1993, Proceedings, LNCS 803*, Springer, pp. 509–529, doi:10.1007/3-540-58043-3.27.
- [70] Iain Phillips & Irek Ulidowski (2006): *Reversing Algebraic Process Calculi*. In Luca Aceto & Anna Ingólfssdóttir, editors: *FoSSaCS, LNCS 3921*, Springer, pp. 246–260, doi:10.1007/11690634_17.
- [71] Iain Phillips & Irek Ulidowski (2007): *Reversibility and Models for Concurrency*. *Electron. Notes Theor. Comput. Sci.* 192(1), pp. 93–108, doi:10.1016/j.entcs.2007.08.018.
- [72] Benjamin C. Pierce & Davide Sangiorgi (1996): *Typing and Subtyping for Mobile Processes*. *MSCS* 6(5), pp. 409–453, doi:10.1017/S096012950007002X.
- [73] Davide Sangiorgi (1996): *A Theory of Bisimulation for the π -Calculus*. *Acta Inform.* 33(1), pp. 69–97, doi:10.1007/s002360050036.
- [74] Davide Sangiorgi (1999): *The Name Discipline of Uniform Receptiveness*. *Theor. Comput. Sci.* 221(1-2), pp. 457–493, doi:10.1016/S0304-3975(99)00040-7.
- [75] Davide Sangiorgi (2011): *Introduction to Bisimulation and Coinduction*. CUP, doi:10.1017/CBO9780511777110.
- [76] Davide Sangiorgi (2011): *Pi-Calculus*. In David A. Padua, editor: *Encyclopedia of Parallel Computing*, Springer, pp. 1554–1562, doi:10.1007/978-0-387-09766-4.202.
- [77] Davide Sangiorgi & Jan Rutten, editors (2011): *Advanced Topics in Bisimulation and Coinduction*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, doi:10.1017/CBO9780511792588.
- [78] Davide Sangiorgi & David Walker (2001): *On Barbed Equivalences in π -Calculus*. In Kim Guldstrand Larsen & Mogens Nielsen, editors: *CONCUR, LNCS 2154*, Springer, pp. 292–304, doi:10.1007/3-540-44685-0_20.
- [79] Davide Sangiorgi & David Walker (2001): *The Pi-calculus*. CUP.
- [80] Vladimiro Sassone, Mogens Nielsen & Glynn Winskel (1996): *Models for Concurrency: Towards a Classification*. *Theor. Comput. Sci.* 170(1-2), pp. 297–348, doi:10.1016/S0304-3975(96)80710-9.

- [81] Peter Selinger & Benoît Valiron (2009): *Quantum Lambda Calculus*. In Simon Gay & Ian Mackie, editors: *Semantic Techniques in Quantum Computation*, Cambridge University Press, p. 135–172, doi:10.1017/CBO9781139193313.005.
- [82] Colin Stirling (1995): *Modal and Temporal Logics for Processes*. In Faron Moller & Graham M. Birtwistle, editors: *Logics for Concurrency - Structure versus Automata (8th Banff Higher Order Workshop, Banff, Canada, August 27 - September 3, 1995, Proceedings)*, LNCS 1043, Springer, pp. 149–237, doi:10.1007/3-540-60915-6_5.
- [83] Paul Taylor: *Comment to "Substitution is pullback"*. Available at <http://math.andrej.com/2012/09/28/substitution-is-pullback/>.
- [84] André van Tondervan (2004): *A Lambda Calculus for Quantum Computation*. *SIAM J. Comput.* 33(5), pp. 1109–1135, doi:10.1137/S0097539703432165.
- [85] Carlos A. Varela (2013): *Programming Distributed Computing Systems: A Foundational Approach*. The MIT Press.
- [86] Glynn Winskel (2017): *Event Structures, Stable Families and Concurrent Games*. Lecture notes, University of Cambridge. Available at <https://www.cl.cam.ac.uk/~gw104/ecsym-notes.pdf>.
- [87] Wang Yi (1991): *CCS + Time = An Interleaving Model for Real Time Systems*. In Javier Leach Albert, Burkhard Monien & Mario Rodríguez-Artalejo, editors: *ICALP*, LNCS 510, Springer, pp. 217–228, doi:10.1007/3-540-54233-7_136.