MODAL LOGICS FOR NOMINAL TRANSITION SYSTEMS

JOACHIM PARROW, JOHANNES BORGSTRÖM, LARS-HENRIK ERIKSSON, RAMŪNAS GUTKOVAS, TJARK WEBER

Uppsala University, Sweden

ABSTRACT. We define a general notion of transition system where states and action labels can be from arbitrary nominal sets, actions may bind names, and state predicates from an arbitrary logic define properties of states. A Hennessy-Milner logic for these systems is introduced, and proved adequate and expressively complete for bisimulation equivalence. A main technical novelty is the use of finitely supported infinite conjunctions. We show how to treat different bisimulation variants such as early, late, open and weak in a systematic way, explore the folklore theorem that state predicates can be replaced by actions, and make substantial comparisons with related work. The main definitions and theorems have been formalised in Nominal Isabelle.

1. INTRODUCTION

Transition systems are ubiquitous in models of computing. Specifications about what may and must happen during executions are often formulated in a modal logic. There is a plethora of different versions of both transition systems and logics, including a variety of higher-level constructs such as updatable data structures, new name generation, alias generation, dynamic topologies for parallel components etc. In this paper we formulate a general kind of transition system where such aspects can be treated uniformly, and define accompanying modal logics. Our results are on adequacy, i.e., that logical equivalence coincides with bisimilarity, and expressive completeness, i.e., that any bisimulation-preserving property can be expressed.

States. In any transition system there is a set of *states* $P, Q, \ldots$ representing the configurations a system can reach, and a relation telling how a computation can move between them. Many formalisms, for example all process algebras, define languages for expressing states, but in the present paper we shall make no assumptions about any such syntax and just assume that the states form a set.

Actions. In systems describing communicating parallel processes, the transitions are often labelled with *actions* α, β, representing the externally observable effect of the transition. A transition $P \xrightarrow{\alpha} P'$ thus says that in state P the execution can progress to P' while conducting the action α, which is visible to the rest of the world. For example, in CCS

these actions are atomic and partitioned into output and input communications. In value-passing calculi the actions can be more complicated, consisting of a channel designation and a value from some data structure to be sent along that channel.

Scope openings. With the advent of the pi-calculus [MPW92] an important aspect of transitions was introduced: that of name generation and scope opening. The main idea is that names (i.e., atomic identifiers) can be scoped to represent local resources. They can also be transmitted in actions, to give the environment access to this resource. In the monadic pi-calculus such an action is written $\overline{a}(\nu b)$, to mean that the local name b is exported along the channel a. These names can be subjected to alpha-conversion: if $P \xrightarrow{\overline{a}(\nu b)} P'$ and c is a fresh name then also $P \xrightarrow{\overline{a}(\nu c)} P'\{c/b\}$, where $P'\{c/b\}$ is P' with all bs replaced by cs. Making this idea fully formal is not entirely trivial and many papers gloss over it. In the polyadic pi-calculus several names can be exported in one action, and in psi-calculi arbitrary data structures may contain local names. In this paper we make no assumptions about how actions are expressed, and just assume that for any action α there is a finite set of names $\mathrm{bn}(\alpha)$, the *binding names*, representing exported names. In our formalisation we use nominal sets [Pit13], an attractive theory to reason about objects depending on names on a high level and in a fully rigorous way.

State predicates. The final general components of our transition systems are the *state predicates* ranged over by φ, representing what can be concluded in a given state. For example state predicates can be equality tests of expressions, or connectivity between communication channels. We write $P \vdash \varphi$ to mean that in state P the state predicate φ holds. We make no assumptions about what the state predicates are, beyond the fact that they form a nominal set, and that $\vdash$ does not depend on particular names.

A structure with states, transitions, and state predicates as discussed above we call a *nominal transition system*.

Hennessy-Milner Logic. Modal logic has been used since the 1970s to describe how facts evolve through computation. We use the popular and general branching time logic known as Hennessy-Milner Logic [HM85] (HML). Here the idea is that an action modality $\langle \alpha \rangle$ expresses a possibility to perform an action α. If A is a formula then $\langle \alpha \rangle A$ says that it is possible to perform α and reach a state where A holds. With conjunction and negation this gives a powerful logic shown to be *adequate* for bisimulation equivalence: two processes satisfy the same formulas exactly if they are bisimilar. In the general case, conjunction must take an infinite number of operands when the transition systems have states with an infinite number of outgoing transitions.

Contributions. Our definition of nominal transition systems is very general since we leave open what the states, transitions and predicates are. The only requirement is that transitions satisfy alpha-conversion. A technically important point is that we do not assume the usual *name preservation principle*, that if $P \xrightarrow{\alpha} P'$ then the names occurring in P' must be a subset of those occurring in P and α. This means that the results are applicable to a wide range of calculi. For example, the pi-calculus represents a trivial instance where there are no state predicates. CCS represent an even more trivial instance where bn always returns the empty set. In the fusion calculus and the applied pi-calculus the state contains an environmental part which tells which expressions are equal to which. In the general

framework of psi-calculi the states are processes together with assertions describing their environments. All of these are special cases of nominal transition systems.

We define a modal logic with the $\langle\alpha\rangle$ operator that binds the names in $\mathrm{bn}(\alpha)$. The fully formal treatment of this requires care in ensuring that infinite conjunctions do not exhaust all names, leaving none available for alpha-conversion. All previous works that have considered this issue properly have used uniformly bounded conjunction, i.e., the set of all names in all conjuncts is finite. Instead of this we use the notion of finite support from nominal sets: the conjunction of an infinite set of formulas is admissible if the set has finite support. This results in a much greater generality and expressiveness. For example, we can now define quantifiers and the next-step modalities as derived operators. Also the traditional fixpoint operators from the modal mu-calculus are definable through an infinite set of approximants. The main technical difficulty is to ensure that formulas and their alpha-equivalence classes are finitely supported throughout, i.e., only depend on a finite set of names, even in the presence of infinite conjunction.

We establish *adequacy*: that logical equivalence coincides with bisimilarity. Compared to previous such adequacy results our proof takes a new twist. We further establish *expressive completeness*: that all properties (i.e., subsets of the set of states) that are bisimulation-closed can be expressed as formulas. To our knowledge this result is the first of its kind.

We provide versions of the logic adequate for a whole family of bisimulation equivalences, including late, early, their corresponding congruences, open, and hyper. Traditionally these differ in how they take name substitutions into account. We define a general kind of effect function, including many different kinds of substitution, and show how all variants can be obtained by varying it. We also show how such effects can be represented directly as transitions. Thus these different kinds of bisimulation can all be considered the same, only on different transition systems.

Weak bisimulation, where the so called silent actions do not count, identifies many more states than strong. We provide adequate and expressively complete logics for weak bisimulation. In the presence of arbitrary state predicates this is a particularly challenging area and there seems to be more than one alternative formulation. We formally prove the folklore theorem that for strong and weak bisimulation, state predicates can be encoded as actions on self-loops. The counterpart for weak logic is less clear and there appear a few different possibilities.

Finally we compare our logic to several other proposed logics for CCS and developments of the pi-calculus. A conclusion is that we can easily represent most of them. The correspondence is not exact because of our slightly different treatment of conjunction, but we certainly gain simplicity and robustness in otherwise complicated logics. We also show how our framework can be applied to obtain logics where none have been suggested previously.

Our main definitions and theorems have been formalised in Nominal Isabelle [UK12]. This has required significant new ideas to represent data types with infinitary constructors like infinite conjunction, and their alpha-equivalence classes. As a result we corrected several details in our formulations and proofs, and now have very high confidence in their correctness. The formalisation effort has been substantial, and we consider it a very worthwhile investment. It is hard to measure it precisely since five persons worked on this on and off since 2015, but at a very rough estimate it is less than half of the total effort. The main hurdle was in the very beginning, to get a representation of the infinitely wide formulas with bound names, and accompanying induction schemes.

Exposition. In the following section we provide the necessary background on nominal sets. In Section 3 we present our main definitions and results on nominal transition systems and modal logics. In Sections 4 and 5 we derive useful operators such as quantifiers and fixpoints, and indicate some practical uses. Section 6 shows how to treat variants of bisimilarity such as late and open in a uniform way, and in Section 7 we treat a logic for weak bisimilarity. Section 8 presents an encoding of state predicates as actions. In Sections 9 and 10 we compare with related work and demonstrate how our framework can be applied to recover earlier results uniformly. Finally Section 11 concludes with some remarks on the formalisation in Nominal Isabelle.

This paper is an extended version of [PBE+15] and [PWBE17].

2. Background on nominal sets

Nominal sets [Pit13] is a general theory of objects which depend on names, and in particular formulates the notion of alpha-equivalence when names can be bound. The reader need not know nominal set theory to follow this paper, but some key definitions will make it easier to appreciate our work and we recapitulate them here.

We assume a countably infinite multi-sorted set of atomic identifiers or *names* $\mathcal{N}$ ranged over by $a, b, \ldots$. [1] A *permutation* is a bijection on names that leaves all but finitely many names invariant. The singleton permutation which swaps names a and b and has no other effect is written $(a\,b)$, and the identity permutation that swaps nothing is written id. Permutations are ranged over by π, π'. The effect of applying a permutation π to an object X is written $\pi \cdot X$. Formally, the permutation action $\cdot$ can be any operation that satisfies $\mathrm{id} \cdot X = X$ and $\pi \cdot (\pi' \cdot X) = (\pi \circ \pi') \cdot X$, but a reader may comfortably think of $\pi \cdot X$ as the object obtained by permuting all names in X according to π.

A set of names N *supports* an object X if for all π that leave all members of N invariant it holds $\pi \cdot X = X$. In other words, if N supports X then names outside N do not matter to X. If a finite set supports X then there is a unique minimal set supporting X, called the *support* of X, written $\mathrm{supp}(X)$, intuitively consisting of exactly the names that matter to X. As an example the set of names textually occurring in a datatype element is the support of that element, and the set of free names is the support of the alpha-equivalence class of the element. Note that in general, the support of a set is not the same as the union of the support of its members. An example is the set of all names; each element has itself as support, but the whole set has empty support since $\pi \cdot \mathcal{N} = \mathcal{N}$ for any π.

We write $a\#X$, pronounced "a is fresh for X", for $a \notin \mathrm{supp}(X)$. The intuition is that if $a\#X$ then X does not depend on a in the sense that a can be replaced with any fresh name without affecting X. If N is a set of names we write $N\#X$ for $\forall a \in N \,.\, a\#X$.

A *nominal set* S is a set with a permutation action such that $X \in S \Rightarrow \pi \cdot X \in S$, and where each member $X \in S$ has finite support. A main point is that then each member has infinitely many fresh names available for alpha-conversion. Similarly, a set of names N supports a function f on a nominal set if for all π that leave N invariant it holds $\pi \cdot f(X) = f(\pi \cdot X)$, and similarly for relations and functions of higher arity. Thus we extend the notion of support to finitely supported functions and relations as the minimal finite support, and can derive general theorems such as $\mathrm{supp}(f(X)) \subseteq \mathrm{supp}(f) \cup \mathrm{supp}(X)$.

---

[1] In [Pit13] they are called *atoms* and the set of atoms is written $\mathbb{A}$.

An object that has empty support we call *equivariant.* For example, a unary function f is equivariant if $\pi \cdot f(X) = f(\pi \cdot X)$ for all π, X. The intuition is that an equivariant object does not treat any name special.

In order to reason about bound names we adopt the notion of generalised name abstraction for finite sets of names (see [Pit13, Chapter 4.6]). Let S be a nominal set. The *nominal alpha-equivalence* $=_\alpha$ is the binary relation on $\mathcal{P}_{\mathrm{fin}}(\mathcal{N}) \times S$ defined by $(N, X) =_\alpha (N', X')$ if there is a permutation π such that $\pi{\cdot}(N, X) = (N', X')$ and $\pi(a) = a$ for all $a \in \mathrm{supp}(X)\backslash N$. It is easily proven that $=_\alpha$ is an equivariant equivalence. The intuition is that the pair (N, X) means X with bound names N, and that alpha-converting the bound names is allowed as long as there are no captures, i.e., collisions with names free in X.

The set of equivalence classes $(\mathcal{P}_{\mathrm{fin}}(\mathcal{N}) \times S)/{=_\alpha}$ is traditionally written $[\mathcal{P}_{\mathrm{fin}}(\mathcal{N})]S$, and the equivalence class containing (N, X) is written $<N{>}X$. This is an unfortunate clash of notation with Hennessy-Milner logics, where these brackets signify a modal action operator.

3. Nominal transition systems and Hennessy-Milner logic

3.1. **Basic definitions.** We define nominal transition systems, bisimilarity, and a corresponding Hennessy-Milner logic in this subsection.

Definition 1. A *nominal transition system* is characterised by the following
- STATES: A nominal set of *states* ranged over by P, Q.
- PRED: A nominal set of *state predicates* ranged over by φ.
- An equivariant binary relation $\vdash$ on STATES and PRED. We write $P \vdash \varphi$ to mean that in state P the state predicate φ holds.
- ACT: A nominal set of *actions* ranged over by α.
- An equivariant function bn from ACT to finite sets of names, which for each α returns a subset of $\mathrm{supp}(\alpha)$, called the *binding names*.
- An equivariant subset of STATES $\times [\mathcal{P}_{\mathrm{fin}}(\mathcal{N})](\mathrm{ACT} \times \mathrm{STATES})$, called the *transition relation*, written $\rightarrow$. If $(P, <N{>}(\alpha, Q)) \in\ \rightarrow$ it must hold that $\mathrm{bn}(\alpha) = N$.

We call $<\mathrm{bn}(\alpha){>}(\alpha, P')$ a *residual*. A residual is thus an alpha-equivalence class of a pair of an action and a state, where the scope of the binding names in the action also contains the state. For $(P, <\mathrm{bn}(\alpha){>}(\alpha, P')) \in\ \rightarrow$ we write $P \xrightarrow{\alpha} P'$. This follows the traditional notation in process algebras like the pi-calculus, although it hides the fact that the scope of the names bound by the action extends into the target state. Note that the transition relation satisfies alpha-conversion of residuals: If $a \in \mathrm{bn}(\alpha)$, $b\#\alpha, P'$ and $P \xrightarrow{\alpha} P'$ then also $P \xrightarrow{(a\,b){\cdot}\alpha} (a\,b) \cdot P'$.

As an example, basic CCS from [Mil89] is a trivial nominal transition system. Here the STATES are the CCS agents, ACT the CCS actions, $\mathrm{bn}(\alpha) = \emptyset$ for all actions, and PRED $= \emptyset$. For the pi-calculus, STATES are the pi-calculus agents, and ACT the four kinds of pi-calculus actions (silent, output, input, bound output). In the early semantics bn returns the empty set except for bound outputs where $\mathrm{bn}(\bar{a}(\nu x)) = \{x\}$. In the late semantics there are actions like $a(x)$ where x is a placeholder so also $\mathrm{bn}(a(x)) = \{x\}$. In the polyadic pi-calculus each action may bind a finite set of names.

In the original terminology of these and similar calculi these occurrences of x are referred to as "bound." We believe a better terminology is "binding," since they bind into the target

state. For higher-order calculi this distinction is important. Consider an example where objects transmitted in a communication are processes, and a communicated object contains a bound name:

$$P \xrightarrow{\overline{a}(\nu b)Q} R$$

The action here transmits the process $(\nu b)Q$ along the channel a. The name b is local to Q, so alpha-converting b to some new name affects only Q. Normally agents are considered up to alpha-equivalence; this means that b is not in the support of the action, and we have $\mathrm{bn}(\overline{a}(\nu b)Q) = \emptyset$.

In the same calculus we may also have a different action

$$P \xrightarrow{(\nu b)\overline{a}Q} R$$

Again this transmits a process along the channel a, but the process here is just Q. The name b is shared between Q and R, and is extruded in the action. An alpha-conversion of b thus affects both Q and R simultaneously. In the action b is a free name, in the sense that b is in its support, and it cannot be replaced by another name in the action alone. Here we have $\mathrm{bn}((\nu b)\overline{a}Q) = \{b\}$.

In all of the above we have PRED $= \emptyset$ since communication is the only way a process may influence a parallel process, and thus communications are the only things that matter for process equivalence. More general examples come from psi-calculi [BJPV11] where there are so called "conditions" representing what holds in different states; those would then correspond to PRED. Other calculi, e.g. [WG05, BM07], also have mechanisms where processes can influence each other without explicit communication, such as fusions and updates of a constraint store. All of these are straightforward to accommodate as nominal transition systems. Section 10 contains further descriptions of these and other examples.

Definition 2. A *bisimulation* R is a symmetric binary relation on states in a nominal transition system satisfying the following two criteria: $R(P, Q)$ implies

(1) *Static implication*: $P \vdash \varphi$ implies $Q \vdash \varphi$.
(2) *Simulation*: For all α, P' such that $\mathrm{bn}(\alpha) \# Q$ there exist Q' such that if $P \xrightarrow{\alpha} P'$ then $Q \xrightarrow{\alpha} Q'$ and $R(P', Q')$

We write $P \mathrel{\dot\sim} Q$ to mean that there exists a bisimulation R such that $R(P, Q)$.

Static implication means that bisimilar states must satisfy the same state predicates; this is reasonable since these can be tested by an observer or parallel process. The simulation requirement is familiar from the pi-calculus. Note that this definition corresponds to "early" bisimulation in the pi-calculus. In Section 6 we will consider other variants of bisimilarity.

Proposition 1. $\dot\sim$ is an equivariant equivalence relation.

Proof. The proof has been formalised in Isabelle. Equivariance is a simple calculation, based on the observation that if R is a bisimulation, then $\pi \cdot R$ is a bisimulation. To prove reflexivity of $\dot\sim$, we note that equality is a bisimulation. Symmetry is immediate from Def. 2. To prove transitivity, we show that the composition of $\dot\sim$ with itself is a bisimulation; the simulation requirement is proved by a considering an alpha-variant of $P \xrightarrow{\alpha} P'$ where $\mathrm{bn}(\alpha)$ is fresh for Q. $\qquad\square$

We shall now define a Hennessy-Milner logic for nominal transition systems. As mentioned above, we shall use conjunctions with an infinite set of conjuncts, and need to take care to avoid set-theoretic paradoxes. For example, if we allow conjunctions over arbitrary subsets of formulas then the formulas will not form a set (because its cardinality would then be the same as the cardinality of the set of its subsets). We thus fix an infinite cardinal κ to bound the conjunctions. Our logic for nominal transition systems is the following.

Definition 3. The nominal set of formulas $\mathcal{A}$ ranged over by A is defined by induction as follows:

$$A \quad ::= \quad \bigwedge_{i \in I} A_i \quad | \quad \neg A \quad | \quad \varphi \quad | \quad \langle \alpha \rangle A$$

Support and name permutation are defined as usual (permutation distributes over all formula constructors). In $\bigwedge_{i \in I} A_i$ it is required that the indexing set I has bounded cardinality, by which we mean that $|I| < \kappa$. We assume that κ is sufficiently large; specifically, we require $\kappa > \aleph_0$ (so that we may form countable conjunctions) and $\kappa > |\text{STATES}|$. It is also required that the set of conjuncts $\{A_i \mid i \in I\}$ has finite support; this is then the support of the conjunction.

Alpha-equivalent formulas are identified; the only binding construct is in $\langle \alpha \rangle A$ where $\text{bn}(\alpha)$ binds into α and A. To be completely formally correct we should write the final clause of Definition 3 as $<\text{bn}(\alpha)>(\langle \alpha \rangle A)$. As with the transitions, we abbreviate this to just $\langle \alpha \rangle A$, letting the scope of the names bound in α tacitly extend into A. To avoid notation clashes we shall in the following let name abstractions be implicit in all transitions and logical formulas.

For a simple example related to the pi-calculus with the late semantics, consider the formula $\langle a(x) \rangle \langle \bar{b}x \rangle \top$ where $\top$ is the empty conjunction and thus always true. We have $\text{bn}(a(x))) = \{x\}$ and therefore an alpha-variant of the formula is $\langle a(y) \rangle \langle \bar{b}y \rangle \top$. It says that it must be possible to input something along a and then output it along b. In the early semantics, where the input action contains the object received rather than a placeholder, the corresponding formula is

$$\bigwedge_{x \in \mathcal{N}} \langle ax \rangle \langle \bar{b}x \rangle \top$$

in other words, the conjunction is over the set of formulas $S = \{\langle ax \rangle \langle \bar{b}x \rangle \top \mid x \in \mathcal{N}\}$. This set has finite support, in fact the support is just $\{a, b\}$. The reason is that for $c, d \# \{a, b\}$ we have $(c\,d) \cdot S = S$. Note that the formulas in this set have no finite common support, i.e., there is no finite set of names that supports all elements, and thus the conjunction is not expressible in the usual logics for the pi-calculus.

This example highlights one of the main novelties in Definition 3, that we use conjunction of a possibly infinite and finitely supported set of conjuncts. In comparison, the earliest HML for CCS, Hennessy and Milner (1985) [HM85], uses finite conjunction, meaning that the logic is adequate only for finitely branching transition systems. In his subsequent book (1989) [Mil89] Milner admits arbitrary infinite conjunction, disregarding the danger of running into paradoxes. Abramsky (1991) [Abr91] employs a kind of uniformly bounded conjunction, with a finite set of names that supports all conjuncts, an idea that is also used in the first HML for the pi-calculus (1993) [MPW93]. All subsequent developments follow one of these three approaches.

Our main point is that both finite and uniformly bounded conjunction are expressively weak, in that the logic is not adequate for the full range of nominal transition systems, and

in that quantifiers over infinite structures are not definable. In contrast, our use of finitely supported sets of conjuncts is adequate for all nominal transition systems (cf. Theorems 1 and 2 below) and admits quantifiers as derived operators (cf. Section 4 below). Universal quantification over names $\forall x \in \mathcal{N}.A(x)$ is usually defined to mean that $A(n)$ must hold for all $n \in \mathcal{N}$. We can define this as the (infinite) conjunction of all these $A(n)$. This set of conjuncts is not uniformly bounded if $n \in \mathrm{supp}(A(n))$. But it is supported by $\mathrm{supp}(A)$ since, for any permutation π not affecting $\mathrm{supp}(A)$ we have $\pi \cdot A(n) = A(\pi(n))$ which is also a conjunct; thus the set of conjuncts is unaffected by π.

Another novelty is the use of a nominal set of actions α with binders, and the formal definition of alpha-equivalence. We define it by structural recursion over formulas. Two conjunctions $\bigwedge_{i \in I} A_i$ and $\bigwedge_{i \in I} B_i$ are alpha-equivalent if for every conjunct A_i there is an alpha-equivalent conjunct B_j, and vice versa. The other cases are standard; two formulas $\langle \alpha \rangle A$ and $\langle \beta \rangle B$ are alpha-equivalent if there exists a permutation π, renaming the binding names of α to those of β, such that $\pi \cdot A$ and B are alpha-equivalent, and $\pi \cdot \alpha = \beta$. Moreover, π must leave names that are free in A invariant. The free names in a formula are *not* defined by (mutual) structural recursion, but are simply given by the support of the formula's alpha-equivalence class. Notably, the free names in a conjunction may be different from the union of free names in all conjuncts. This is analogous to the situation for nominal sets in general, where the support of a set is not necessarily the same as the union of the supports of its members. It follows from this definition that a name is free in $\langle \alpha \rangle A$ if it is in $\mathrm{supp}(\alpha)$ or free in A, and not contained in $\mathrm{bn}(\alpha)$. Fortunately, our formalisation in Isabelle proves that we need not keep the details of this construction in mind, but can simply identify alpha-equivalent formulas. The notions of free names and support then coincide.

The validity of a formula A for a state P is written $P \models A$ and is defined by recursion over A as follows.

Definition 4.

$$
\begin{aligned}
P &\models \bigwedge_{i \in I} A_i &&\text{if for all } i \in I \text{ it holds that } P \models A_i \\
P &\models \neg A &&\text{if not } P \models A \\
P &\models \varphi &&\text{if } P \vdash \varphi \\
P &\models \langle \alpha \rangle A &&\text{if there exists } P' \text{ such that } P \xrightarrow{\alpha} P' \text{ and } P' \models A
\end{aligned}
$$

In the last clause we assume that $\langle \alpha \rangle A$ is a representative of its alpha-equivalence class such that $\mathrm{bn}(\alpha) \# P$.

Lemma 1. $\models$ is equivariant.

Proof. By the Equivariance Principle in Pitts (2013) [Pit13, page 21]. A more detailed proof that verifies $P \models A \iff \pi \cdot P \models \pi \cdot A$ for any permutation π has been formalised in Isabelle. The proof proceeds by structural induction on A, using equivariance of all involved relations. For the case $\langle \alpha \rangle A$ in particular, we use the fact that if $\langle \alpha' \rangle A' = \langle \alpha \rangle A$, then $\langle \pi \cdot \alpha' \rangle (\pi \cdot A') = \langle \pi \cdot \alpha \rangle (\pi \cdot A)$. $\square$

Definition 5. Two states P and Q are *logically equivalent*, written $P \doteq Q$, if for all A it holds that $P \models A$ iff $Q \models A$.

3.2. **Logical adequacy.** We show that the logic defined in Section 3.1 is adequate for bisimilarity; that is, bisimilarity and logical equivalence coincide.

Theorem 1. $P \mathbin{\dot{\sim}} Q \implies P \mathbin{\dot{=}} Q$

Proof. The proof has been formalised in Isabelle. Assume $P \mathbin{\dot{\sim}} Q$. We show $P \models A \iff Q \models A$ by structural induction on A.

(1) Base case: $A = \varphi$. Then $P \models A \iff P \vdash \varphi \iff Q \vdash \varphi \iff Q \models A$ by static implication and symmetry of $\dot{\sim}$.

(2) Inductive steps $\bigwedge_{i \in I} A_i$ and $\neg A$: immediate by induction.

(3) Inductive step $\langle \alpha \rangle A$: Assume $P \models \langle \alpha \rangle A$. Then for some alpha-variant $\langle \alpha' \rangle A' = \langle \alpha \rangle A$, $\exists P' . P \xrightarrow{\alpha'} P'$ and $P' \models A'$. Without loss of generality we assume also $\mathrm{bn}(\alpha') \# Q$, otherwise just find an alpha-variant of $\langle \alpha' \rangle A'$ where this holds. Then by simulation $\exists Q' . Q \xrightarrow{\alpha'} Q'$ and $P' \mathbin{\dot{\sim}} Q'$. By induction and $P' \models A'$ we get $Q' \models A'$, whence by definition $Q \models \langle \alpha \rangle A$. The proof of $Q \models \langle \alpha \rangle A \implies P \models \langle \alpha \rangle A$ is symmetric, using the fact that $P \mathbin{\dot{\sim}} Q$ entails $Q \mathbin{\dot{\sim}} P$. $\square$

The converse result uses the idea of *distinguishing formulas*.

Definition 6. A *distinguishing formula* for P and Q is a formula A such that $P \models A$ and not $Q \models A$.

The following lemma says that we can find such a formula where, a bit surprisingly, the support does not depend on Q.

Lemma 2. If $P \mathbin{\dot{\neq}} Q$ then there exists a distinguishing formula B for P and Q such that $\mathrm{supp}(B) \subseteq \mathrm{supp}(P)$.

Proof. The proof has been formalised in Isabelle. If $P \mathbin{\dot{\neq}} Q$ then there exists a distinguishing formula A for P and Q, i.e., $P \models A$ and not $Q \models A$. Let $\Pi_P = \{\pi \mid n \in \mathrm{supp}(P) \Rightarrow \pi(n) = n\}$ be the group of name permutations that leave $\mathrm{supp}(P)$ invariant and let $\mathcal{B}$ be the Π_P-orbit of A, i.e.,

$$\mathcal{B} = \{\pi \cdot A \mid \pi \in \Pi_P\}$$

In the terminology of Pitts [Pit13] ch. 5, $\mathcal{B}$ is $\mathrm{hull}_{\mathrm{supp}(P)} A$. Clearly, if $a, b \# P$ and $\pi \in \Pi_P$ then $(a\,b) \circ \pi \in \Pi_P$. Thus $(a\,b) \cdot \mathcal{B} = \mathcal{B}$ and therefore $\mathrm{supp}(\mathcal{B}) \subseteq \mathrm{supp}(P)$. This means that the formula $B = \bigwedge \mathcal{B}$ is well-formed and $\mathrm{supp}(B) \subseteq \mathrm{supp}(P)$. For all $\pi \in \Pi_P$ we have by definition $P = \pi \cdot P$ and by equivariance $\pi \cdot P \models \pi \cdot A$, i.e., $P \models \pi \cdot A$. Therefore $P \models B$. Furthermore, since the identity permutation is in Π_P and not $Q \models A$ we get not $Q \models B$. $\square$

Note that in this proof, B uses a conjunction that is not uniformly bounded.

Theorem 2. $P \mathbin{\dot{=}} Q \implies P \mathbin{\dot{\sim}} Q$

Proof. The proof has been formalised in Isabelle. We establish that $\dot{=}$ is a bisimulation. Obviously it is symmetric. So assume $P \mathbin{\dot{=}} Q$, we need to prove the two requirements on a bisimulation.

(1) Static implication. $P \vdash \varphi$ iff $P \models \varphi$ iff $Q \models \varphi$ iff $Q \vdash \varphi$.

(2) Simulation. The proof is by contradiction. Assume that $\dot{=}$ does not satisfy the simulation requirement. Then there exist P, Q, P', α with $\mathrm{bn}(\alpha) \# Q$ such that $P \mathbin{\dot{=}} Q$ and $P \xrightarrow{\alpha} P'$ and, letting $\mathcal{Q} = \{Q' \mid Q \xrightarrow{\alpha} Q'\}$, for all $Q' \in \mathcal{Q}$ it holds that $P' \mathbin{\dot{\neq}} Q'$. Assume $\mathrm{bn}(\alpha) \# P$, otherwise just find an alpha-variant of the transition satisfying this. By $P' \mathbin{\dot{\neq}} Q'$, for all $Q' \in \mathcal{Q}$ there exists a distinguishing formula for P' and Q'.

The formula may depend on Q', and by Lemma 2 we can find such a distinguishing formula $B_{Q'}$ for P' and Q' with $\mathrm{supp}(B_{Q'}) \subseteq \mathrm{supp}(P')$. Therefore the formula

$$B = \bigwedge_{Q' \in \mathcal{Q}} B_{Q'}$$

is well-formed with support included in $\mathrm{supp}(P')$. We thus get that $P \models \langle \alpha \rangle B$ but not $Q \models \langle \alpha \rangle B$, contradicting $P \doteq Q$. □

This proof of the simulation property is different from other such proofs in the literature. For finitely branching transition systems, $\mathcal{Q}$ is finite so finite conjunction is enough to define B. For transition systems with the name preservation property, i.e., that if $P \xrightarrow{\alpha} P'$ then $\mathrm{supp}(P') \subseteq \mathrm{supp}(P) \cup \mathrm{supp}(\alpha)$, uniformly bounded conjunction suffices with common support $\mathrm{supp}(P) \cup \mathrm{supp}(Q) \cup \mathrm{supp}(\alpha)$. Without the name preservation property, we here use a not uniformly bounded conjunction in Lemma 2.

3.3. **Expressive completeness.** Theorem 2 shows that for every pair of non-bisimilar states, there is a formula that distinguishes them. We can prove an even stronger result: the logic contains formulas that precisely characterise the equivalence classes of bisimilarity. Moreover, every finitely supported set of states that is closed under bisimilarity can be described by a formula.

We first strengthen Lemma 2 by showing that there is an equivariant function that yields distinguishing formulas for non-equivalent states.

Lemma 3. If $P \not\doteq Q$, write $B_{P,Q}$ for a distinguishing formula for P and Q such that $\mathrm{supp}(B_{P,Q}) \subseteq \mathrm{supp}(P)$. Then

$$D(P, Q) := \bigwedge_{\pi} \pi^{-1} \cdot B_{\pi \cdot P, \pi \cdot Q}$$

defines a distinguishing formula for P and Q with support included in $\mathrm{supp}(P)$. Moreover, the function D is equivariant.

Proof. The proof has been formalised in Isabelle. Assume $P \not\doteq Q$. For any permutation π, $\pi \cdot P \not\doteq \pi \cdot Q$ by equivariance of $\doteq$. Thus the formulas $B_{\pi \cdot P, \pi \cdot Q}$ exist by Lemma 2, and $\mathrm{supp}(\pi^{-1} \cdot B_{\pi \cdot P, \pi \cdot Q}) \subseteq \mathrm{supp}(P)$. Hence the conjunction $D(P, Q)$ is well-formed with support included in $\mathrm{supp}(P)$.

From $\pi \cdot P \models B_{\pi \cdot P, \pi \cdot Q}$ we have $P \models \pi^{-1} \cdot B_{\pi \cdot P, \pi \cdot Q}$ by equivariance of $\models$. Hence $P \models D(P, Q)$. Also not $Q \models B_{P,Q}$, hence (by considering the identity permutation) not $Q \models D(P, Q)$. Therefore $D(P, Q)$ is a distinguishing formula for P and Q.

For equivariance, we have

$$
\begin{aligned}
\rho \cdot D(P, Q) \;&=\; \bigwedge_{\pi} (\rho \circ \pi^{-1}) \cdot B_{\pi \cdot P, \pi \cdot Q} \\
\{\sigma = \pi \circ \rho^{-1}\} \;&=\; \bigwedge_{\sigma} \sigma^{-1} \cdot B_{(\sigma \circ \rho) \cdot P, (\sigma \circ \rho) \cdot Q} \\
&=\; D(\rho \cdot P, \rho \cdot Q)
\end{aligned}
$$

as required. □

Definition 7. A *characteristic formula* for P is a formula A such that for all Q,

$$P \mathbin{\dot\sim} Q \quad \text{iff} \quad Q \models A$$

Characteristic formulas can be obtained as a (possibly infinite) conjunction of distinguishing formulas.

Lemma 4. Let D be defined as in Lemma 3. The formula

$$\mathrm{Char}(P) := \bigwedge_{P \not\mathrel{\dot=} Q} D(P, Q)$$

is a characteristic formula for P.

Proof. The proof has been formalised in Isabelle. By Lemma 3, $\mathrm{supp}(D(P, Q)) \subseteq \mathrm{supp}(P)$ for all Q with $P \not\mathrel{\dot=} Q$. Thus the conjunction is well-formed with support included in $\mathrm{supp}(P)$.

We show that $P \mathrel{\dot=} Q$ iff $Q \models \mathrm{Char}(P)$. The lemma then follows because bisimilarity and logical equivalence coincide (Theorems 1 and 2).

$\Longrightarrow$: Assume $P \mathrel{\dot=} Q$. By the definition of distinguishing formulas, $P \models D(P, Q')$ for all Q' with $P \not\mathrel{\dot=} Q'$. Hence $P \models \mathrm{Char}(P)$. Therefore $Q \models \mathrm{Char}(P)$ by logical equivalence.

$\Longleftarrow$: Assume $Q \models \mathrm{Char}(P)$ and $P \not\mathrel{\dot=} Q$. From the definition of $\mathrm{Char}(P)$ we have $Q \models D(P, Q)$. This is a contradiction since $D(P, Q)$ is a distinguishing formula for P and Q. $\qquad\square$

Lemma 5. Let $\mathrm{Char}(P)$ be defined as in Lemma 4. The function

$$\mathrm{Char} \colon P \mapsto \mathrm{Char}(P)$$

is equivariant.

Proof. The proof has been formalised in Isabelle. We verify that $\pi \cdot \mathrm{Char}(P) = \mathrm{Char}(\pi \cdot P)$ for any permutation π using the equivariance of $\bigwedge$, $\not\mathrel{\dot=}$, and D. $\qquad\square$

Definition 8. A *property* of states is a subset of the states. A property S is *well-formed* if it is finitely supported and bisimulation closed, i.e., if $P \in S$ and $P \sim Q$ the also $Q \in S$.

Well-formed properties can be described as a (possibly infinite) disjunction of characteristic formulas. (Disjunction $\bigvee_{i \in I} A_i$ is defined in the usual way as $\neg \bigwedge_{i \in I} \neg A_i$; see Section 4 for further details.)

Theorem 3 (Expressive Completeness)**.** Let S be a well-formed property. The there exists a formula A such that

$$P \models A \quad \text{iff} \quad P \in S$$

Proof. Let

$$A = \bigvee_{P' \in S} \mathrm{Char}(P')$$

The proof has been formalised in Isabelle. By the equivariance of Char (Lemma 5) $\mathrm{supp}(\{\mathrm{Char}(P') \mid P' \in S\}) \subseteq \mathrm{supp}(S)$ which is finite by assumption, so the disjunction is well-formed.

Assume $P \in S$. Since $P \models \mathrm{Char}(P)$ we get $P \models \bigvee_{P' \in S} \mathrm{Char}(P')$. Conversely, assume that $P \models \bigvee_{P' \in S} \mathrm{Char}(P')$. Then for some $P' \in S$, $P \models \mathrm{Char}(P')$. By Lemma 4, $P' \mathbin{\dot\sim} P$. Hence $P \in S$ because S is closed under bisimilarity. $\qquad\square$

4. Derived formulas

Dual connectives. We define logical disjunction $\bigvee_{i \in I} A_i$ in the usual way as $\neg \bigwedge_{i \in I} \neg A_i$, when the indexing set I has bounded cardinality and $\{A_i \mid i \in I\}$ has finite support. A special case is $I = \{1, 2\}$: we then write $A_1 \wedge A_2$ instead of $\bigwedge_{i \in I} A_i$, and dually for $A_1 \vee A_2$. We write $\top$ for the empty conjunction $\bigwedge_{i \in \emptyset}$, and $\bot$ for $\neg \top$. We also write $A \implies B$ for $B \vee \neg A$. The must modality $[\alpha]A$ is defined as $\neg \langle \alpha \rangle \neg A$, and requires A to hold after every possible α-labelled transition from the current state. Note that $\mathrm{bn}(\alpha)$ bind into A. For example, $[\alpha](A \wedge B)$ is equivalent to $[\alpha]A \wedge [\alpha]B$, and dually $\langle \alpha \rangle(A \vee B)$ is equivalent to $\langle \alpha \rangle A \vee \langle \alpha \rangle B$.

Quantifiers.   Let S be any finitely supported set of bounded cardinality and use v to range over members of S. Write $A\{v/x\}$ for the substitution of v for x in A, and assume this substitution function is equivariant. Then we define $\forall x \in S.A$ as $\bigwedge_{v \in S} A\{v/x\}$. There is not necessarily a common finite support for the formulas $A\{v/x\}$, for example if S is some term algebra over names, but the set $\{A\{v/x\} \mid v \in S\}$ has finite support bounded by $\{x\} \cup \mathrm{supp}(S) \cup \mathrm{supp}(A)$. In our examples in Section 10, substitution is defined inductively on the structure of formulas, based on primitive substitution functions for actions and state predicates, avoiding capture and preserving the binding names of actions.

Existential quantification $\exists x \in S.A$ is defined as the dual $\neg \forall x \in S. \neg A$. When X is a metavariable used to range over a nominal set $\mathcal{X}$, we simply write X for "$X \in \mathcal{X}$". As an example, $\forall a.A$ means that the formula $A\{n/a\}$ holds for all names $n \in \mathcal{N}$.

New name quantifier. The new name quantifier $\mathsf{N}x.A$ intuitively states that $P \models A\{n/x\}$ holds where n is a fresh name for P. For example, suppose we have actions of the form $a\,b$ for input, and $\overline{a}\,b$ for output where a and b are free names, then the formula $\mathsf{N}x.[a\,x]\langle \overline{b}\,x \rangle \top$ expresses that whenever a process inputs a fresh name x on channel a, it has to be able to output that name on channel b. If the name received is not fresh (i.e., already present in P) then P is not required to do anything. Therefore this formula is weaker than $\forall x.[a\,x]\langle \overline{b}\,x \rangle \top$.

Since A and P have finite support, if $P \models A\{n/x\}$ holds for some n fresh for P, by equivariance it also holds for almost all n, i.e., all but finitely many n. Conversely, if it holds for almost all n, it must hold for some $n \# \mathrm{supp}(P)$. Therefore $\mathsf{N}x$ is often pronounced "for almost all x". In other words, $P \models \mathsf{N}x.A$ holds if $\{x \mid P \models A(x)\}$ is a cofinite set of names [Pit13, Definition 3.8].

To avoid the need for a substitution function, we here define the new name quantifier using name swapping $(a\,n)$. Letting $\mathrm{COF} = \{S \subseteq \mathcal{N} \mid \mathcal{N} \setminus S \text{ is finite}\}$ we thus encode $\mathsf{N}x.A$ as $\bigvee_{S \in \mathrm{COF}} \bigwedge_{n \in S} (x\,n) \cdot A$. This formula states there is a cofinite set of names n such that for all of them $(x\,n) \cdot A$ holds. The support of $\bigwedge_{n \in S} (x\,n) \cdot A$ is bounded by $(\mathcal{N} \setminus S) \cup \mathrm{supp}(A)$ where $S \in \mathrm{COF}$, and the support of the encoding $\bigvee_{S \in \mathrm{COF}} \bigwedge_{n \in S} (x\,n) \cdot A$ is bounded by $\mathrm{supp}(A)$.

Next step. We can generalise the action modality to sets of actions: if T is a finitely supported set of actions that has bounded cardinality, we write $\langle T \rangle A$ for $\bigvee_{\alpha \in T} \langle \alpha \rangle A$. The support of $\{\langle \alpha \rangle A \mid \alpha \in T\}$ is bounded by $\mathrm{supp}(T) \cup \mathrm{supp}(A)$ and thus finite. Dually, we write $[T]A$ for $\neg \langle T \rangle \neg A$, denoting that A holds after all transitions with actions in T. Note that binding names of actions in T bind into A, and so $\langle \alpha \rangle A$ is equivalent to $\langle \{\alpha\} \rangle A$ for any α.

To encode the next-step modality, let $\mathrm{ACT}_A = \{\alpha \mid \mathrm{bn}(\alpha) \# A\}$. Note that $\mathrm{supp}(\mathrm{ACT}_A) \subseteq \mathrm{supp}(A)$ is finite. If κ (Definition 3) is larger than $|\mathrm{ACT}_A|$, we write $\langle \rangle A$ for $\langle \mathrm{ACT}_A \rangle A$,

meaning that we can make some (non-capturing) transition to a state where A holds. As an example, $\langle\rangle\top$ means that the current state is not deadlocked. The dual modality $[\,]A = \neg\langle\rangle\neg A$ means that A holds after every transition from the current state. Larsen [Lar88] uses the same approach to define next-step operators in HML, though his version is less expressive since he uses a finite action set to define the next-step modality.

5. Fixpoint operators

Fixpoint operators are a way to introduce recursion into a logic. For example, they can be used to concisely express safety and liveness properties of a transition system, where by safety we mean that some invariant holds for all reachable states, and by liveness that some property will eventually hold. Kozen [Koz83] introduced the least ($\mu X.A$) and the greatest ($\nu X.A$) fixpoint operators in modal logic. Intuitively, the least fixpoint operator states a property that holds for states of a finite path, while the greatest holds for states of an infinite path.

By combining the fixpoint and next-step operators, we can encode the temporal logic CTL [CE82], following Emerson [Eme96]. The CTL formula $\mathsf{AG}\,A$, which states that A holds along all paths, is defined as $\nu X.A \wedge [\,]X$. Dually the formula $\mathsf{EF}\,A$, stating the there is some path where A eventually holds, is defined $\mu X.A \vee \langle\rangle X$. These are special cases of more general formulas: the formula $\mathsf{A}[A\,\mathsf{U}\,B]$ states that for all paths A holds until B holds, and dually $\mathsf{E}[A\,\mathsf{U}\,B]$ states that there is a path along which A holds until B. They are encoded as $\nu X.B \vee ([\,]X \wedge A)$ and $\mu X.B \vee (\langle\rangle X \wedge A)$, respectively. For example, deadlock-freedom is given by the CTL formula $\mathsf{AG}\langle\rangle\top$ expressing that every reachable state has a transition. The encoding of this formula is $\nu X.\langle\rangle\top \wedge [\,]X$.

We extend the logic of Definition 3 with the least fixpoint operator and give meaning to formulas as sets of satisfying states. We show that the meaning of the fixpoint operator is indeed a fixpoint. We proceed to show that the least fixpoint operator can be directly encoded in the logic. The greatest fixpoint operator can then be expressed as the dual of the least fixpoint.

5.1. **Logic with the least fixpoint operator.**

Definition 9. We extend the nominal set of formulas with the least fixpoint operator:

$$A \quad ::= \quad \bigwedge_{i\in I} A_i \quad | \quad \neg A \quad | \quad \varphi \quad | \quad \langle\alpha\rangle A \quad | \quad X \quad | \quad \mu X.A$$

where X ranges over a countable set of equivariant variables. We require that all occurrences of a variable X in a formula $\mu X.A$ are in the scope of an even number of negations.

An occurrence of a variable X in A is said to be free if it is not a subterm of some $\mu X.B$. We say that a formula A is closed if for every variable X, none of its occurrences in A are free. We use a capture-avoiding substitution function $[A/X]$ on formulas that substitutes A for the free occurrences of the variable X. In particular, $(\langle\alpha\rangle B)[A/X] = \langle\alpha\rangle(B[A/X])$ when $\mathrm{bn}(\alpha)$ is fresh for A.

We give a semantics to formulas containing variables and fixpoint modalities as sets of satisfying states.

Definition 10. A *valuation function* ε is a finitely supported map from variables to (finitely supported) sets of states. We write $\varepsilon[X \mapsto S]$ for the valuation function that maps X to S, and any variable $X' \neq X$ to $\varepsilon(X')$.

We define the *interpretation* of formula A under valuation ε by structural induction as the set of states $[\![A]\!]_\varepsilon$:

$$
\begin{array}{rcl}
[\![\bigwedge_{i \in I} A_i]\!]_\varepsilon &=& \bigcap_{i \in I} [\![A_i]\!]_\varepsilon \\
[\![\neg A]\!]_\varepsilon &=& \mathrm{STATES} - [\![A]\!]_\varepsilon \\
[\![\varphi]\!]_\varepsilon &=& \{P \mid P \vdash \varphi\} \\
[\![\langle\alpha\rangle A]\!]_\varepsilon &=& \left\{P \mid \exists \alpha' \, A' \, P' . \, \langle\alpha\rangle A = \langle\alpha'\rangle A' \wedge \mathrm{bn}(\alpha')\#P, \varepsilon \wedge P \xrightarrow{\alpha'} P' \wedge P' \in [\![A']\!]_\varepsilon\right\} \\
[\![X]\!]_\varepsilon &=& \varepsilon(X) \\
[\![\mu X.A]\!]_\varepsilon &=& \bigcap \left\{S \in \mathcal{P}_{\mathrm{fs}}(\mathrm{STATES}) \mid [\![A]\!]_{\varepsilon[X \mapsto S]} \subseteq S\right\}
\end{array}
$$

We write $[\![\;]\!]$ for the function $(A, \varepsilon) \mapsto [\![A]\!]_\varepsilon$.

Lemma 6. $[\![\;]\!]$ is equivariant.

Proof. By the Equivariance Principle [Pit13, page 21]. $\qquad\square$

Lemma 7. For any formula A and valuation function ε, $[\![A]\!]_\varepsilon \in \mathcal{P}_{\mathrm{fs}}(\mathrm{STATES})$.

Proof. By equivariance of $[\![\;]\!]$, $\mathrm{supp}([\![A]\!]_\varepsilon) \subseteq \mathrm{supp}(A) \cup \mathrm{supp}(\varepsilon)$. $\qquad\square$

Temporal operators such as "eventually" can be encoded using the least fixpoint operator: the formula $\mu X.\langle\alpha\rangle X \vee A$ states that eventually A holds along some path labelled with α. We define the greatest fixpoint operator $\nu X.A$ in terms of the least as $\neg\mu X.\neg A[\neg X/X]$. Using the greatest fixpoint operator we can state global invariants: $\nu X.[\alpha]X \wedge A$ expresses that A holds along all paths labelled with α. We can freely mix the fixpoint operators to obtain formulas like $\nu X.[\alpha]X \wedge (\mu Y.\langle\beta\rangle Y \vee A)$, which means that for each state along any path labelled with α, a state where A holds is reachable along a path labelled with β.

As sanity checks for our definition, we prove that the interpretation of formulas without fixpoint modalities is unchanged, and that the interpretation of the formula $\mu X.A$ is indeed the least fixpoint of the function $F_A^\varepsilon \colon S \mapsto [\![A]\!]_{\varepsilon[X \mapsto S]}$.

Proposition 2. Let A be a formula as in Definition 3. Then for any valuation function ε and state P, $P \models A$ if and only if $P \in [\![A]\!]_\varepsilon$.

Proof. By structural induction on A. The clauses for X and $\mu X.A'$ in Definition 10 are not used. The interesting case is

Case $\langle\alpha\rangle A'$**:** Assume $P \models \langle\alpha\rangle A'$. Without loss of generality assume also $\mathrm{bn}(\alpha)\#P, \varepsilon$, otherwise just find an alpha-variant of $\langle\alpha\rangle A'$ where this holds. From Definition 4 we obtain P' such that $P \xrightarrow{\alpha} P'$ and $P' \models A'$. Then $P' \in [\![A']\!]_\varepsilon$ by the induction hypothesis, hence $P \in [\![\langle\alpha\rangle A']\!]_\varepsilon$ by Definition 10.

Next, assume $P \in [\![\langle\alpha\rangle A']\!]_\varepsilon$. From Definition 10 we obtain an alpha-variant $\langle\alpha'\rangle A'' = \langle\alpha\rangle A'$ and P' such that $\mathrm{bn}(\alpha')\#P$, $P \xrightarrow{\alpha'} P'$ and $P' \in [\![A'']\!]_\varepsilon$. Then $P' \models A''$ by the induction hypothesis. Hence $P \models \langle\alpha'\rangle A'' = \langle\alpha\rangle A'$ by Definition 4.

$\qquad\square$

Lemma 8. For any formula $\mu X.A$ and valuation function ε, F_A^ε has finite support.

Proof. By equivariance of $[\![\;]\!]$, $\mathrm{supp}(F_A^\varepsilon) \subseteq \mathrm{supp}(A) \cup \mathrm{supp}(\varepsilon)$. $\qquad\square$

Lemma 9. For any formula $\mu X.A$ and valuation function ε, the function $F_A^\varepsilon \colon \mathcal{P}_{\mathrm{fs}}(\textsc{states}) \to \mathcal{P}_{\mathrm{fs}}(\textsc{states})$ is monotonic with respect to subset inclusion.

Proof. By structural induction on A, for arbitrary ε. Let $S, T \in \mathcal{P}_{\mathrm{fs}}(\textsc{states})$ such that $S \subseteq T$. We prove a more general statement: if all occurrences of X in A are positive (i.e., within the scope of an even number of negations), $F_A^\varepsilon(S) \subseteq F_A^\varepsilon(T)$, and if all occurrences of X in A are negative (i.e., within the scope of an odd number of negations), $F_A^\varepsilon(T) \subseteq F_A^\varepsilon(S)$. The interesting case is

Case $\mu X'.A'$: If $X = X'$, note that for any $U, V \in \mathcal{P}_{\mathrm{fs}}(\textsc{states})$, $\varepsilon[X \mapsto U][X \mapsto V] = \varepsilon[X \mapsto V]$. Therefore, $[\![\mu X'.A']\!]_{\varepsilon[X \mapsto S]} = [\![\mu X'.A']\!]_{\varepsilon[X \mapsto T]}$ is immediate from Definition 10.

Otherwise, $X \neq X'$. Suppose that all occurrences of X in A are positive. Then all occurrences of X in A' are positive, and for any $V \in \mathcal{P}_{\mathrm{fs}}(\textsc{states})$, $[\![A']\!]_{\varepsilon[X' \mapsto V][X \mapsto S]} \subseteq [\![A']\!]_{\varepsilon[X' \mapsto V][X \mapsto T]}$ by the induction hypothesis applied to A' and $\varepsilon[X' \mapsto V]$. Since $X \neq X'$, for any $U, V \in \mathcal{P}_{\mathrm{fs}}(\textsc{states})$, $\varepsilon[X \mapsto U][X' \mapsto V] = \varepsilon[X' \mapsto V][X \mapsto U]$. Thus,

$$[\![\mu X'.A']\!]_{\varepsilon[X \mapsto S]} = \bigcap \left\{ S' \in \mathcal{P}_{\mathrm{fs}}(\textsc{states}) \mid [\![A']\!]_{\varepsilon[X \mapsto S][X' \mapsto S']} \subseteq S' \right\} \subseteq$$

$$\bigcap \left\{ S' \in \mathcal{P}_{\mathrm{fs}}(\textsc{states}) \mid [\![A']\!]_{\varepsilon[X \mapsto T][X' \mapsto S']} \subseteq S' \right\} = [\![\mu X'.A']\!]_{\varepsilon[X \mapsto T]}.$$

The case where all occurrences of X in A are negative is similar. $\qquad\square$

We use a nominal version of Tarski's fixpoint theorem [Tar55] to show existence, uniqueness, and the construction of the least fixpoint of F_A^ε. Note that the usual Tarski fixpoint theorem does not apply, since the lattice $\mathcal{P}_{\mathrm{fs}}(\textsc{states})$ is not necessarily complete.

Theorem 4. Suppose X is a nominal set, and $f \colon \mathcal{P}_{\mathrm{fs}}(X) \to \mathcal{P}_{\mathrm{fs}}(X)$ is finitely supported and monotonic with respect to subset inclusion. Then f has a least fixpoint $\mathrm{lfp}\, f$, and

$$\mathrm{lfp}\, f = \bigcap \{ S \in \mathcal{P}_{\mathrm{fs}}(\textsc{states}) \mid f(S) \subseteq S \}.$$

Proof. (Due to Pitts [Pit15])
Since f is finitely supported and $\bigcap$ is equivariant, also $\bigcap \{ S \in \mathcal{P}_{\mathrm{fs}}(\textsc{states}) \mid f(S) \subseteq S \}$ is finitely supported (with support bounded by $\mathrm{supp}(f)$). It then follows by a replay of the usual Tarski argument that $\bigcap \{ S \in \mathcal{P}_{\mathrm{fs}}(\textsc{states}) \mid f(S) \subseteq S \}$ is the least fixpoint of f. $\quad\square$

Finally, we can show that the interpretation of a fixpoint formula $\mu X.A$ is the least fixpoint of the function F_A^ε.

Proposition 3. For any formula $\mu X.A$ and valuation function ε, $[\![\mu X.A]\!]_\varepsilon = \mathrm{lfp}\, F_A^\varepsilon$.

Proof. Using Lemmas 8 and 9, the proposition is immediate from Theorem 4. $\qquad\square$

5.2. **Encoding the least fixpoint operator.** The least fixpoint operator can be encoded in our logic of Section 3. The basic idea here is simple: we translate the fixpoint modality into a transfinite disjunction that at each step α unrolls the formula α times. This then semantically corresponds to a limit of an increasing chain generated by a monotonic function, i.e., a least fixpoint.

Recall that the cardinality of a set of conjuncts—and thus also of a set of disjuncts—must be less than some fixed infinite cardinal κ (see Definition 3).

Definition 11. We define the formula $\mathsf{unroll}_\alpha(\mu X.A)$ for all ordinals $\alpha < \kappa$ by transfinite induction.

$$
\begin{aligned}
\mathsf{unroll}_0(\mu X.A) &= \bot \\
\mathsf{unroll}_{\alpha+1}(\mu X.A) &= A[\mathsf{unroll}_\alpha(\mu X.A)/X] \\
\mathsf{unroll}_\lambda(\mu X.A) &= \textstyle\bigvee_{\alpha<\lambda} \mathsf{unroll}_\alpha(\mu X.A) \quad \text{when } \lambda \text{ is a limit ordinal}
\end{aligned}
$$

Since unroll_α is equivariant, the disjunction in the limit case has finite support bounded by $\mathrm{supp}(A)$. Note that if A does not contain any fixpoint modalities then $\mathsf{unroll}_\alpha(\mu X.A)$ also does not contain any fixpoint modalities.

To show that for some ordinal $\gamma < \kappa$, the interpretation of $\mathsf{unroll}_\gamma(\mu X.A)$ indeed is the least fixpoint of F_A^ε, we use a nominal version of a chain fixpoint theorem for sets by Kuratowski (1922) [Kur22], augmented with a bound on the depth of the unrolling.

Theorem 5. Suppose X is a nominal set, and $f \colon \mathcal{P}_{\mathrm{fs}}(X) \to \mathcal{P}_{\mathrm{fs}}(X)$ is finitely supported and monotonic with respect to subset inclusion. Set $f^0 = \emptyset$, $f^{\alpha+1} = f(f^\alpha)$, and $f^\lambda = \bigcup_{\alpha<\lambda} f^\alpha$ for limit ordinals λ. Then f has a least fixpoint $\mathrm{lfp}\, f$, and if ν is a cardinal with $\nu > |X|$, there exists an ordinal $\gamma < \nu$ such that

$$
\mathrm{lfp}\, f = f^\gamma.
$$

Proof. First, each f^α is finitely supported, since $\mathrm{supp}(f^\alpha) \subseteq \mathrm{supp}(f)$ for every ordinal α by transfinite induction. Also, using monotonicity of f, we have $f^\alpha \subseteq f^\beta$ for all $\alpha \leq \beta$.

We then show that f has a fixpoint f^γ for some ordinal $\gamma < \nu$, by contradiction. Otherwise, for each $\gamma < \nu$ there is $x_\gamma \in f^{\gamma+1} \setminus f^\gamma$. This yields an injective function $g \colon \nu \to X$ with $g(\gamma) = x_\gamma$, which is a contradiction since $\nu > |X|$.

Let y be any fixpoint of f. For every ordinal α, $f^\alpha \subseteq y$ by transfinite induction, so in particular $f^\gamma \subseteq y$. Thus f^γ is the least fixpoint of f. $\square$

Let $\nu = \max\{|\mathrm{STATES}|^+, \aleph_0\}$ denote the least infinite cardinal larger than $|\mathrm{STATES}|$. Note that $\nu \leq \kappa$ by assumption.

Lemma 10. For any formula $\mu X.A$ and valuation function ε, there exists an ordinal $\gamma < \nu$ such that $\mathrm{lfp}\, F_A^\varepsilon = (F_A^\varepsilon)^\alpha$ for all ordinals $\alpha \geq \gamma$.

Proof. Since $\nu > |\mathrm{STATES}|$, the lemma is an immediate consequence of Theorem 5, whose other assumptions follow from Lemmas 8 and 9. $\square$

From this lemma, we obtain an equivariant function conv that maps each formula $\mu X.A$ and each valuation function ε to the least ordinal $\mathsf{conv}(\mu X.A, \varepsilon) < \nu$ such that

$$
\mathrm{lfp}\, F_A^\varepsilon = (F_A^\varepsilon)^{\mathsf{conv}(\mu X.A,\varepsilon)}.
$$

When $\mathcal{E}$ is a non-empty set of valuation functions, we write $\mathsf{Conv}(\mu X.A, \mathcal{E})$ for $\sup_{\varepsilon \in \mathcal{E}} \mathsf{conv}(\mu X.A, \varepsilon)$.

Lemma 11. Let $\mathcal{E}$ be a non-empty set of valuation functions such that $|\mathcal{E}| < \nu$. Then $\mathsf{Conv}(\mu X.A, \mathcal{E}) < \nu$.

Proof. Note that $\aleph_0$ is regular, and every successor cardinal is regular. Hence ν is regular. Therefore, the set $\{\mathsf{conv}(\mu X.A, \varepsilon) \mid \varepsilon \in \mathcal{E}\}$ (whose cardinality is less than ν) is not cofinal in ν. $\square$

Lemma 12. For any formulas A, B and valuation function ε, if A does not contain any fixpoint operators, then $[\![A[B/X]]\!]_\varepsilon = [\![A]\!]_{\varepsilon[X \mapsto [\![B]\!]_\varepsilon]}$.

Proof. By structural induction on A. The clause for $\mu X.A'$ in Definition 10 is not used. $\square$

Lemma 13. For any formula A and valuation function ε, if A does not contain any fixpoint operators, then $[\![\mathsf{unroll}_\alpha(\mu X.A)]\!]_\varepsilon = (F_A^\varepsilon)^\alpha$ for all ordinals $\alpha < \kappa$.

Proof. By transfinite induction on α.
  (1) Base case: $[\![\mathsf{unroll}_0(\mu X.A)]\!]_\varepsilon = [\![\bot]\!]_\varepsilon = \emptyset = (F_A^\varepsilon)^0$ by definition.
  (2) Inductive step:

$$[\![\mathsf{unroll}_{\alpha+1}(\mu X.A)]\!]_\varepsilon = [\![A[\mathsf{unroll}_\alpha(\mu X.A)/X]]\!]_\varepsilon$$

$$\overset{(1)}{=} [\![A]\!]_{\varepsilon[X \mapsto [\![\mathsf{unroll}_\alpha(\mu X.A)]\!]_\varepsilon]}$$

$$\overset{(2)}{=} [\![A]\!]_{\varepsilon[X \mapsto (F_A^\varepsilon)^\alpha]}$$

$$= F_A^\varepsilon((F_A^\varepsilon)^\alpha)$$

$$= (F_A^\varepsilon)^{\alpha+1}$$

as required. Above, equality (1) follows from Lemma 12 and equality (2) follows from the induction hypothesis for α.
  (3) Limit case: The limit case is straightforward. We have

$$[\![\mathsf{unroll}_\lambda(\mu X.A)]\!]_\varepsilon = [\![\bigvee_{\alpha<\lambda} \mathsf{unroll}_\alpha(\mu X.A)]\!]_\varepsilon$$

$$= \bigcup_{\alpha<\lambda} [\![\mathsf{unroll}_\alpha(\mu X.A)]\!]_\varepsilon$$

$$\overset{(1)}{=} \bigcup_{\alpha<\lambda} (F_A^\varepsilon)^\alpha$$

$$= (F_A^\varepsilon)^\lambda$$

where equality (1) follows from the induction hypothesis for all $\alpha < \lambda$. $\qquad\square$

Definition 12. Given any non-empty set of valuation functions $\mathcal{E}$ with $|\mathcal{E}| < \nu$, we define the formula $\underline{A}_\mathcal{E}$ homomorphically on the structure of A. The encoding $\underline{\mu X.A}_\mathcal{E}$ of a fixpoint formula is its unrolling up to a sufficiently large ordinal.

$$
\begin{aligned}
\underline{\bigwedge_{i \in I} A_i}_\mathcal{E} &= \bigwedge_{i \in I} \underline{A_i}_\mathcal{E} \\
\underline{\neg A}_\mathcal{E} &= \neg \underline{A}_\mathcal{E} \\
\underline{\varphi}_\mathcal{E} &= \varphi \\
\underline{\langle \alpha \rangle A}_\mathcal{E} &= \langle \alpha \rangle \underline{A}_\mathcal{E} \\
\underline{X}_\mathcal{E} &= X \\
\underline{\mu X.A}_\mathcal{E} &= \mathsf{unroll}_\gamma(\mu X.\underline{A}_{\mathcal{E}'}) \quad \text{where } \mathcal{E}' = \{\varepsilon[X \mapsto (F_A^\varepsilon)^\alpha] \mid \varepsilon \in \mathcal{E}, \alpha \le \mathsf{conv}(\mu X.A, \varepsilon)\} \\
&\qquad\qquad\qquad\qquad\qquad \text{and } \gamma = \mathsf{Conv}(\mu X.\underline{A}_{\mathcal{E}'}, \mathcal{E})
\end{aligned}
$$

In the fixpoint case, since $|\mathcal{E}| < \nu$ and $\mathsf{Conv}(\mu X.A, \mathcal{E}) < \nu$ (Lemma 11), we also have $|\mathcal{E}'| < \nu \cdot \nu = \nu$. Since the encoding function is equivariant, it preserves the finite support property for conjunctions. Clearly, $\underline{A}_\mathcal{E}$ does not contain any fixpoint operators. Moreover, if A is closed, then $\underline{A}_\mathcal{E}$ does not contain any variables, and is therefore a formula in the sense of Definition 3.

Theorem 6. Let $\mathcal{E}$ be a non-empty set of valuation functions such that $|\mathcal{E}| < \nu$. For any formula A and valuation function $\varepsilon \in \mathcal{E}$, $[\![\underline{A}_\mathcal{E}]\!]_\varepsilon = [\![A]\!]_\varepsilon$.

Proof. By structural induction on A, for arbitrary $\mathcal{E}$ and ε. The interesting case is

   Case $\mu X.A'$: We need to show that $[\![\underline{\mu X.A'}_{\mathcal{E}}]\!]_\varepsilon = [\![\mu X.A']\!]_\varepsilon$. Let $\mathcal{E}'$ and γ be as in Definition 12. First, we compute the left-hand side to $[\![\underline{\mu X.A'}_{\mathcal{E}}]\!]_\varepsilon = [\![\mathsf{unroll}_\gamma(\mu X.\underline{A'}_{\mathcal{E}'})]\!]_\varepsilon = (F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^\gamma = \mathrm{lfp}\, F^\varepsilon_{\underline{A'}_{\mathcal{E}'}}$ by Lemma 13 and Lemma 10, where we use the fact that $\gamma \geq \mathsf{conv}(\mu X.\underline{A'}_{\mathcal{E}'}, \varepsilon)$ since $\varepsilon \in \mathcal{E}$. For the right-hand side, we have $[\![\mu X.A']\!]_\varepsilon = \mathrm{lfp}\, F^\varepsilon_{A'}$ by Proposition 3.

   We now show $(F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^\alpha = (F^\varepsilon_{A'})^\alpha$ for all ordinals $\alpha \leq \mathsf{conv}(\mu X.A', \varepsilon) + 1$ by transfinite induction on α. The base case ($\alpha = 0$) and limit case ($\alpha = \lambda$) are straightforward. For the inductive step, we have

$$\begin{aligned}
(F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^{\alpha+1} &= F^\varepsilon_{\underline{A'}_{\mathcal{E}'}}((F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^\alpha)\\
&\overset{(1)}{=} F^\varepsilon_{\underline{A'}_{\mathcal{E}'}}((F^\varepsilon_{A'})^\alpha)\\
&= [\![\underline{A'}_{\mathcal{E}'}]\!]_{\varepsilon[X \mapsto (F^\varepsilon_{A'})^\alpha]}\\
&\overset{(2)}{=} [\![A']\!]_{\varepsilon[X \mapsto (F^\varepsilon_{A'})^\alpha]}\\
&= F^\varepsilon_{A'}((F^\varepsilon_{A'})^\alpha)\\
&= (F^\varepsilon_{A'})^{\alpha+1}
\end{aligned}$$

as required. Above, equality (1) follows from the induction hypothesis for α, and equality (2) follows from the outer induction hypothesis applied to $\mathcal{E}'$ and $\varepsilon[X \mapsto (F^\varepsilon_{A'})^\alpha]$. Note that $\varepsilon[X \mapsto (F^\varepsilon_{A'})^\alpha] \in \mathcal{E}'$ since $\alpha \leq \mathsf{conv}(\mu X.A', \varepsilon)$.

   It follows with Lemma 10 that $(F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^{\mathsf{conv}(\mu X.A', \varepsilon)} = \mathrm{lfp}\, F^\varepsilon_{A'} = (F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^{\mathsf{conv}(\mu X.A', \varepsilon)+1}$. Hence $\mathrm{lfp}\, F^\varepsilon_{A'}$ is a fixpoint of $F^\varepsilon_{\underline{A'}_{\mathcal{E}'}}$, and thus $\mathrm{lfp}\, F^\varepsilon_{\underline{A'}_{\mathcal{E}'}} \subseteq \mathrm{lfp}\, F^\varepsilon_{A'}$. Moreover, $(F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^\alpha \subseteq \mathrm{lfp}\, F^\varepsilon_{\underline{A'}_{\mathcal{E}'}}$ for any ordinal α (as in the proof of Theorem 5), so in particular $\mathrm{lfp}\, F^\varepsilon_{A'} = (F^\varepsilon_{\underline{A'}_{\mathcal{E}'}})^{\mathsf{conv}(\mu X.A', \varepsilon)} \subseteq \mathrm{lfp}\, F^\varepsilon_{\underline{A'}_{\mathcal{E}'}}$. By combining both inclusions, $\mathrm{lfp}\, F^\varepsilon_{\underline{A'}_{\mathcal{E}'}} = \mathrm{lfp}\, F^\varepsilon_{A'}$. □

   If A is closed, its semantics does not depend on ε. In this case, we can pick an arbitrary valuation function to perform the encoding: e.g., let $\varepsilon_\emptyset$ be the valuation that maps every variable to $\emptyset$. We simply write $\underline{A}$ for $\underline{A}_{\{\varepsilon_\emptyset\}}$.

   Every closed formula containing fixpoint operators can be translated into an equivalent formula without fixpoint operators.

Corollary 1. For any ε, P and closed formula A, we have $P \models \underline{A}$ iff $P \in [\![A]\!]_\varepsilon$.

Proof. By Theorem 6 and Proposition 2.   □

6. Logics for variants of bisimilarity

6.1. **Variants of bisimilarity.** There are variants of bisimilarity, differing in the effect the binding in a transition can have on the target state. A typical example is in the pi-calculus where the input action binds a name, and the target state must be considered for all possible instantiations of it. There is then a difference between the so-called *late* bisimilarity, where the target states must bisimulate before instantiating the input, and *early* bisimilarity, where it is enough to bisimulate after each instantiation. There are also corresponding congruences, obtained by closing bisimilarity under all substitutions of names for names.

The original value-passing variant of CCS from 1989 [Mil89] uses early bisimilarity. The original bisimilarity for the pi-calculus (1992) is of the late kind [MPW92], where it also was noted that late equivalence is the corresponding congruence. Early bisimilarity and congruence in the pi-calculus were introduced in 1993 [MPW93], where HMLs adequate for a few different bisimilarity definitions are explored. Other ways to treat the name instantiations include Sangiorgi's *open bisimilarity* (1993) [San93] and Parrow and Victor's *hyperbisimilarity* (1998) [PV98]. Hyperbisimilarity is the requirement that the bisimulation relation is closed under all name instantiations, and corresponds to a situation where the environment of a process may at any time instantiate any name. Open bisimilarity is between late and hyper: the environment may at any time instantiate any name except those that have previously been extruded or declared constant.

In our definition of nominal transition systems there are no particular input variables in the states or in the actions, and thus no a priori concept of replacing a name by something else. In order to cover all of the above variants of the pi-calculus and also of high-level extensions of it, we generalise name instantiation using a notion of effect functions. An *effect* is simply a finitely supported function from states to states. For instance, in the monadic pi-calculus the effects would be the functions replacing one name by another. In a value-passing calculus the effects would be substitutions of values for variables. In the psi-calculi framework the effects would be sequences of parallel substitutions. Variants of bisimilarity then correspond to the use of various effects. For instance, if the action contains an input variable x, then the effects appropriate for late bisimilarity would be substitutions for x. Our only requirement is that the effects form a nominal set, which we designate by $\mathcal{F}$. In the following let $\mathcal{P}_{\mathrm{fs}}(\mathcal{F})$, ranged over by F, be the finitely supported subsets of $\mathcal{F}$.

The definition of bisimilarity will now be through a $\mathcal{P}_{\mathrm{fs}}(\mathcal{F})$-indexed family $\{R_F\}$ of bisimulations. The index F ("first") simply says which effects must be taken into account before performing a transition. There is a function L ("later"), which determines what effects should be considered after taking a transition. In its simplest form these effects depend only on the action of the transition, that is, L has type $\mathrm{ACT} \to \mathcal{P}_{\mathrm{fs}}(\mathcal{F})$. We require that L is equivariant.

Definition 13 (Simple L-bisimulation and $\overset{F/L}{\sim}$). A *simple L-bisimulation*, for $L\colon \mathrm{ACT} \to \mathcal{P}_{\mathrm{fs}}(\mathcal{F})$ equivariant, is a $\mathcal{P}_{\mathrm{fs}}(\mathcal{F})$-indexed family $\{R_F\}$ of symmetric binary relations on states satisfying the following:

For all $F \in \mathcal{P}_{\mathrm{fs}}(\mathcal{F})$, $R_F(P,Q)$ implies

(1) *Static implication*: For all $f \in F$, $f(P) \vdash \varphi$ implies $f(Q) \vdash \varphi$.
(2) *Simulation*: For all $f \in F$ and α, P' such that $\mathrm{bn}(\alpha)\#f(Q), F$ there exists Q' such that
$$\text{if } f(P) \xrightarrow{\alpha} P' \text{ then } f(Q) \xrightarrow{\alpha} Q' \text{ and } R_{L(\alpha)}(P', Q')$$

We write $P \overset{F/L}{\sim} Q$, called F/L-bisimilarity, to mean that there exists a simple L-bisimulation $\{R_F\}_{F \in \mathcal{P}_{\mathrm{fs}}(\mathcal{F})}$ such that $R_F(P,Q)$.

To exemplify F/L-bisimilarity we shall consider some of the popular bisimulation equivalences in the monadic pi-calculus; the ideas obviously generalize to more advanced settings. Thus the states are pi-calculus agents, and there are no state predicates. In the monadic pi-calculus there are input actions written $a(x)$, where a and x are names and the intention is that the binding input object x shall be instantiated with another name received in a

communication with a parallel process. There are also binding output actions $\bar{a}(x)$ signifying the output of a local name x. We let $\mathrm{bn}(a(x)) = \mathrm{bn}(\bar{a}(x)) = \{x\}$, and for all other actions α we let $\mathrm{bn}(\alpha) = \emptyset$.

The set of effects $\mathcal{F}$ is the set of name substitutions, i.e., functions σ from names to names that are identity almost everywhere. Clearly if a permutation π only permutes names a such that $\sigma(a) = a$ then $\pi \cdot \sigma = \sigma$, thus each substitution has finite support and $\mathcal{F}$ is a nominal set. The effect of applying σ to a state P is notated $P\sigma$, in conformance with most of the literature on the pi-calculus.

Write id for the identity function on names, and let $\mathrm{subst}_x = \{\sigma \mid \forall y \neq x. \ \sigma(y) = y\}$ be the set of substitutions for x. Note that $\mathrm{supp}(\mathrm{subst}_x) = \{x\}$. We now get:

- **Early bisimilarity** $\dot{\sim}_\mathrm{E}$ does not use binding input actions, instead it uses non-binding actions where the received object is already present. Early bisimilarity thus is precisely as defined in Definition 2, which is the same as $\{\mathrm{id}\} / L_\mathrm{E}$-bisimilarity where $L_\mathrm{E}(\alpha) = \{\mathrm{id}\}$ for all α. No substitutive effect is needed; the substitution of output object for input object is included already in the semantics and thus already present in the corresponding nominal transition system.
- **Early equivalence** $\sim_\mathrm{E}$ is early bisimilarity closed under all possible substitutions, i.e., $P \sim_\mathrm{E} Q$ if for all σ it holds $P\sigma \dot{\sim}_\mathrm{E} Q\sigma$. Therefore $\sim_\mathrm{E}$ is $\mathcal{F} / L_\mathrm{E}$-bisimilarity where L_E is as above. Any substitution can be applied initially, and thereafter no substitution is needed. Early equivalence has been shown to be the smallest congruence including early bisimilarity.
- **Late bisimilarity** $\dot{\sim}_\mathrm{L}$ has a binding input action and should consider all possible instantiations of the bound input object before the next transition. In other words, $\dot{\sim}_\mathrm{L}$ is $\{\mathrm{id}\} / L_\mathrm{L}$-bisimilarity where $L_\mathrm{L}(a(x)) = \mathrm{subst}_x$ and $L_\mathrm{L}(\alpha) = \{\mathrm{id}\}$ for all other actions α.
- **Late equivalence** $\sim_\mathrm{L}$ is late bisimilarity closed under all possible substitutions, i.e., $P \sim_\mathrm{L} Q$ if for all σ it holds $P\sigma \dot{\sim}_\mathrm{L} Q\sigma$. Therefore $\sim_\mathrm{L}$ is $\mathcal{F} / L_\mathrm{L}$-bisimilarity where L_L is as above. Late equivalence has been shown to be the smallest congruence including late bisimilarity.
- **Hyperbisimilarity** $\sim_\mathrm{H}$ means that any name can be substituted at any time, thus it is $\mathcal{F} / L_\mathrm{H}$-bisimilarity where $L_\mathrm{H}(\alpha) = \mathcal{F}$ for all α.

Open bisimilarity is more involved and requires a generalisation of the L-function to take additional parameters. We begin by quoting the definition from Sangiorgi 1993 [San93].

Definition 14 (Open bisimilarity, Sangiorgi)**.** A *distinction* is a finite symmetric and irreflexive relation on names. A substitution σ *respects* a distinction D if $(a, b) \in D$ implies $\sigma(a) \neq \sigma(b)$. A distinction-indexed family of symmetric relations $\{S_D\}_D$ is an *open bisimulation* if for all S_D and for each σ which respects D, $(P, Q) \in S_D$ implies

  (1) If $P\sigma \xrightarrow{\bar{a}(b)} P'$ with b fresh then Q' exists s.t. $Q\sigma \xrightarrow{\bar{a}(b)} Q'$ and $(P', Q') \in S_{D'}$ where $D' = D\sigma \cup (\{b\} \times \mathrm{fn}(P\sigma, Q\sigma))$ (with symmetric closure)
  (2) If $P\sigma \xrightarrow{\alpha} P'$ with α not a binding output and $\mathrm{bn}(\alpha)$ fresh then Q' exists s.t. $Q\sigma \xrightarrow{\alpha} Q'$ and $(P', Q') \in S_{D\sigma}$

Write $P \sim_\mathrm{O} Q$ to mean that $(P, Q) \in S_\emptyset$ for some open bisimulation $\{S_D\}_D$.

In this definition, "b is fresh" means that it "is supposed to be different from any other name appearing in the objects of the statement, like processes or distinctions." The function fn extracts the free names of a process; in nominal terms this is the support.

Clearly, the distinctions here correspond to our effect sets since they determine which substitutions should be taken into account. A complication is then that the distinction D' after the transition in clause (1) depends not only on the action but also on D, σ, P and Q. We therefore present an alternative definition of open bisimulation, where distinctions may be infinite but still finitely supported. Write $D + b$ for the distinction $D \cup (\{b\} \times (\mathcal{N} - \{b\}))$ (and symmetric closure) and $D - b$ for the distinction $\{(x, y) \in D \mid x, y \neq b\}$. Note that $\operatorname{supp}(D + b) \subseteq \operatorname{supp}(D) \cup \{b\}$ and $\operatorname{supp}(D - b) \subseteq \operatorname{supp}(D) \cup \{b\}$. It is easy to see that the function that maps a distinction to the set of substitutions respecting it is an equivariant injection. Therefore we let sets of substitutions be represented as distinctions, and say "$\sigma \in D$" instead of "σ respects D."

Definition 15 (Alternative definition of open bisimulation). A distinction-indexed family of symmetric relations $\{S_D\}_D$ is an open bisimulation if for all S_D and for each $\sigma \in D$, $(P, Q) \in S_D$ implies

(1) If $P\sigma \xrightarrow{\overline{a}(b)} P'$ with $b \# Q\sigma, D, \sigma$ then Q' exists s.t. $Q\sigma \xrightarrow{\overline{a}(b)} Q'$ and $(P', Q') \in S_{D\sigma + b}$

(2) If $P\sigma \xrightarrow{\alpha} P'$ with $\operatorname{bn}(\alpha) = \emptyset$ then Q' exists s.t. $Q\sigma \xrightarrow{\alpha} Q'$ and $(P', Q') \in S_{D\sigma}$

(3) If $P\sigma \xrightarrow{a(b)} P'$ with $b \# Q\sigma, D, \sigma$ then Q' exists s.t. $Q\sigma \xrightarrow{a(b)} Q'$ and $(P', Q') \in S_{D\sigma - b}$

Compared to Definition 14, clause (1) now requires b to be distinct from *all* names, not just the names in $P\sigma$ and $Q\sigma$. In the pi-calculus, it is known that if $P\sigma \xrightarrow{\overline{a}(b)} P'$ then $\operatorname{supp}(P') \subseteq \operatorname{supp}(P\sigma) \cup \{b\}$. Therefore, the difference between the clauses only concerns names outside $\operatorname{supp}(P', Q')$, meaning that the additional substitutions allowed by clause (1) of Definition 14 will be injective when restricted to $\operatorname{supp}(P', Q')$, and open bisimilarity (with any distinction) is closed under injective substitutions.

All names "occur" in $D\sigma + b$ even though it has finite support. Thus, in a subsequent input $a(b)$ it will be impossible to choose b fresh according to Sangiorgi. Instead we get the same effect with clause (3): removing an input binder b from the distinction is the same as choosing a new one that does not occur there. We also tighten and make explicit the necessary freshness condition on b.

To capture open bisimulation in our framework, we extend the simple L-functions to take more parameters as follows:

Definition 16 (L-bisimulation). A *(general) L-bisimulation*, for $L \colon \text{ACT} \times \mathcal{P}_{\text{fs}}(\mathcal{F}) \times \mathcal{F} \to \mathcal{P}_{\text{fs}}(\mathcal{F})$ equivariant, is a $\mathcal{P}_{\text{fs}}(\mathcal{F})$-indexed family $\{R_F\}$ of symmetric binary relations on states satisfying the following:
For all $F \in \mathcal{P}_{\text{fs}}(\mathcal{F})$, $R_F(P, Q)$ implies

(1) *Static implication*: For all $f \in F$, $f(P) \vdash \varphi$ implies $f(Q) \vdash \varphi$.

(2) *Simulation*: For all $f \in F$ and α, P' such that $\operatorname{bn}(\alpha) \# f(Q), F, f$ there exist Q' such that

$$\text{if } f(P) \xrightarrow{\alpha} P' \text{ then } f(Q) \xrightarrow{\alpha} Q' \text{ and } R_{L(\alpha, F, f)}(P', Q')$$

The simple L-bisimulation of Definition 13 is thus the special case when L does not depend on F or f. In the following when we write L-bisimulation we always refer to the general case of Definition 16.

To represent open bisimulation as an L-bisimulation we let $\mathcal{F}$ be the set of all substitutions and use distinctions to represent sets of substitutions (thus $\mathcal{F}$ is the empty distinction),

and define the function L_O by:

$$L_O(\alpha, D, \sigma) = \begin{cases} D\sigma + b & \text{if } \alpha = \overline{a}(b) \\ D\sigma & \text{if } \text{bn}(\alpha) = \emptyset \\ D\sigma - b & \text{if } \alpha = a(b) \end{cases}$$

Strictly speaking this L_O is partial since there are sets of substitutions in $\mathcal{P}_{\text{fs}}(\mathcal{F})$ that cannot be represented by a distinction. These substitution sets will not matter and we can make L_O total by assigning an arbitrary value in those cases. We now immediately get, using Definition 15:

Proposition 4. $P \sim_O Q \quad$ iff $\quad P \overset{\mathcal{F}/L_O}{\sim} Q$

In conclusion we have a uniform framework to co-inductively define most forms of bisimulation. The common claim that open bisimulation is the only co-inductively defined congruence in the pi-calculus is somewhat contradicted by our framework, which represents both late and early equivalence conveniently.

6.2. **Variants of the logic.** In view of the previous subsection we only need to provide a modal logic adequate for (general) F/L-bisimilarity; it can then immediately be specialised to all of the above variants. For this we introduce a new kind of logical *effect consequence* operator @, which appears in front of state predicates and actions in formulas. We define the formulas that can directly use effects from F and after actions use effects according to L, ranged over by $A^{F/L}$, in the following way:

Definition 17. Given L as in Definition 16, for all $F \in \mathcal{P}_{\text{fs}}(\mathcal{F})$ define $\mathcal{A}^{F/L}$ as the set of formulas given by the mutually recursive definitions:

$$A^{F/L} \quad ::= \quad \bigwedge_{i \in I} A_i^{F/L} \quad | \quad \neg A^{F/L} \quad | \quad f@\varphi \quad | \quad f@\langle\alpha\rangle A^{L(\alpha,F,f)/L}$$

where we require $f \in F$ and $\text{bn}(\alpha)\#f, F$ and that the conjunction has bounded cardinality and finite support.

Validity of a formula for a state P is defined as in Definition 4, where in the last clause we assume that the formula $\langle\alpha\rangle A$ is a representative of its alpha-equivalence class such that $\text{bn}(\alpha)\#P, F, f$. Validity of formulas involving the effect consequence operator is defined as follows.

Definition 18. For each $f \in \mathcal{F}$,

$$P \models f@A \quad \text{if} \quad f(P) \models A$$

Thus the formula $f@A$ means that A holds when the effect f is applied to the state. The effect consequence operator is similar in spirit to the action modalities: both $f@A$ and $\langle\alpha\rangle A$ assert that something (an effect or action) must be possible and that A holds afterwards. Indeed, effects can be viewed as a special case of transitions (as formalised in Definition 19 below).

Lemma 14. If $A \in \mathcal{A}^{F/L}$ is a distinguishing formula for P and Q, then there exists a distinguishing formula $B \in \mathcal{A}^{F/L}$ for P and Q such that $\text{supp}(B) \subseteq \text{supp}(P, F)$.

Proof. The proof has been formalised in Isabelle. It is an easy generalisation of the proof of Lemma 2, just replace $\mathcal{A}$ by $\mathcal{A}^{F/L}$ and $\text{supp}(P)$ by $\text{supp}(P, F)$ everywhere. We additionally need to prove that $B \in \mathcal{A}^{F/L}$. Since L is equivariant, $A \in \mathcal{A}^{F/L}$ implies $\pi \cdot A \in \mathcal{A}^{(\pi \cdot F)/L} = \mathcal{A}^{F/L}$ for all $\pi \in \Pi_{(P,F)}$, this establishes $B \in \mathcal{A}^{F/L}$. □

Let $P \overset{F/L}{=} Q$ mean that P and Q satisfy the same formulas in $\mathcal{A}^{F/L}$.

Theorem 7. $P \overset{F/L}{\sim} Q$ iff $P \overset{F/L}{=} Q$

Proof. The proof has been formalised in Isabelle. Direction $\Rightarrow$ is a generalisation of Theorem 1.

(1) Base case: $A = f@\varphi$ and $f \in F$. Then $f(P) \vdash \varphi$. By static implication $f(Q) \vdash \varphi$, which means $Q \models A$.

(2) Inductive step $f@\langle\alpha\rangle A$ where $A \in \mathcal{A}^{L(\alpha,F,f)/L}$ with $f \in F$ and $\text{bn}(\alpha)\#f, F$: Assume $P \models f@\langle\alpha\rangle A$. Then $\exists P'. f(P) \overset{\alpha}{\to} P'$ and $P' \models A$. Without loss of generality we assume also $\text{bn}(\alpha)\#f(Q)$, otherwise just find an alpha-variant of the transition where this holds. Then by simulation $\exists Q'. f(Q) \overset{\alpha}{\to} Q'$ and $P' \overset{L(\alpha,F,f)/L}{\sim} Q'$. By induction and $P' \models A$ we get $Q' \models A$, whence by definition $Q \models f@\langle\alpha\rangle A$.

The direction $\Leftarrow$ is a generalisation of Theorem 2: we prove that $\overset{F/L}{=}$ is an F/L-bisimulation. The modified clauses are:

(1) Static implication. Assume $f \in F$, then $f(P) \vdash \varphi$ iff $P \models f@\varphi$ iff $Q \models f@\varphi$ iff $f(Q) \vdash \varphi$.

(2) Simulation. The proof is by contradiction. Assume that $\overset{F/L}{=}$ does not satisfy the simulation requirement. Then there exist $f \in F, P, Q, P', \alpha$ with $\text{bn}(\alpha)\#f(Q), F, f$ such that $P \overset{F/L}{=} Q$ and $f(P) \overset{\alpha}{\to} P'$ and, letting $\mathcal{Q} = \{Q' \mid f(Q) \overset{\alpha}{\to} Q'\}$, for all $Q' \in \mathcal{Q}$ it holds that not $P' \overset{L(\alpha,F,f)/L}{=} Q'$. Thus, for all $Q' \in \mathcal{Q}$ there exists a distinguishing formula in $\mathcal{A}^{L(\alpha,F,f)/L}$ for P' and Q'. The formula may depend on Q', and by Lemma 14 we can find such a distinguishing formula $B_{Q'} \in \mathcal{A}^{L(\alpha,F,f)/L}$ for P' and Q' with $\text{supp}(B_{Q'}) \subseteq \text{supp}(P', L(\alpha, F, f))$. Therefore the formula

$$B = \bigwedge_{Q' \in \mathcal{Q}} B_{Q'}$$

is well formed in $\mathcal{A}^{L(\alpha,F,f)/L}$ with support included in $\text{supp}(P', L(\alpha, F, f))$. We thus get that $P \models f@\langle\alpha\rangle B$ but not $Q \models f@\langle\alpha\rangle B$, contradicting $P \overset{F/L}{=} Q$. □

6.3. **Effects as transitions.** It is possible to view effects as transitions, and thus not use the effect consequence operator explicitly. The idea is to let the effect functions be part of the transition relation, thus $f(P) = P'$ becomes $P \overset{f}{\to} P'$. We here make this idea fully formal. Given a transition system $\mathbf{T}$ we construct a new transition system $L(\mathbf{T})$ without effects. The effects in $\mathbf{T}$ are included among the actions in $L(\mathbf{T})$. The states in $L(\mathbf{T})$ are of two kinds. One is $\text{EF}(F, P)$ where P is a state in $\mathbf{T}$ and $F \in \mathcal{P}_{\text{fs}}(\mathcal{F})$, corresponding to the state P where F contains the effects that should be considered before taking an action or checking a state predicate. The transitions from $\text{EF}(F, P)$ are effects in F. The other kind

is $\text{AC}(f, F, P)$, corresponding to a state P where $f \in F$ has been applied; the transitions from $\text{AC}(f, F, P)$ are the actions from P.

To define this formally, and to talk about different nominal transition systems in the same definition, we write $\text{STATES}_\mathbf{T}$ for the states of the nominal transition system $\mathbf{T}$ and similarly for actions, predicates, etc.

Definition 19. Let $\mathbf{T}$ be a nominal transition system with a set of effects $\mathcal{F}$. Let L be as in Definition 16. The *L-transform* $L(\mathbf{T})$ of a nominal transition system $\mathbf{T}$ is a nominal transition system where:

- $\text{STATES}_{L(\mathbf{T})} = \{\text{AC}(f, F, P) \mid f \in \mathcal{F}, F \in \mathcal{P}_{\text{fs}}(\mathcal{F}), P \in \text{STATES}_\mathbf{T}\}$
  $\cup \{\text{EF}(F, P) \mid F \in \mathcal{P}_{\text{fs}}(\mathcal{F}), P \in \text{STATES}_\mathbf{T}\}$
- $\text{PRED}_{L(\mathbf{T})} = \text{PRED}_\mathbf{T}$
- $\text{AC}(f, F, P) \vdash_{L(\mathbf{T})} \varphi$ if $P \vdash_\mathbf{T} \varphi$, and $\text{EF}(F, P) \vdash_{L(\mathbf{T})} \varphi$ never holds.
- $\text{ACT}_{L(\mathbf{T})} = \text{ACT}_\mathbf{T} \uplus \mathcal{F}$.
- $\text{bn}_{L(\mathbf{T})}(\alpha) = \text{bn}_\mathbf{T}(\alpha)$ for $\alpha \in \text{ACT}_\mathbf{T}$; $\text{bn}_{L(\mathbf{T})}(f) = \emptyset$ for $f \in \mathcal{F}$.
- The transitions in $L(\mathbf{T})$ are of two kinds. If in $\mathbf{T}$ we have $P \xrightarrow{\alpha} P'$ with $\text{bn}(\alpha)\#f, F$, then in $L(\mathbf{T})$ there is a transition $\text{AC}(f, F, P) \xrightarrow{\alpha} \text{EF}(L(\alpha, F, f), P')$. Additionally, for each $f \in F$ it holds $\text{EF}(F, P) \xrightarrow{f} \text{AC}(f, F, f(P))$.

The intuition is that states of kind AC can perform ordinary actions, and states of kind EF can commit effects. The analogous transform of modal formulas in $\mathbf{T}$ to formulas in $L(\mathbf{T})$ simply replaces effects by actions: $L(f@A) = \langle f \rangle L(A)$, and L is homomorphic on all other formula constructors.

The L-transform preserves satisfaction of formulas in the following sense:

Theorem 8. Assume $A \in \mathcal{A}^{F/L}$. Then $P \models A$ iff $\text{EF}(F, P) \models L(A)$.

Proof. The proof has been formalised in Isabelle. It is by induction over formulas in $\mathcal{A}^{F/L}$, for arbitrary P. The cases conjunction and negation are immediate by induction.

Case $A = f@\varphi$: We then know $f \in F$, and have $P \models f@\varphi$ iff $f(P) \vdash_\mathbf{T} \varphi$ iff, by construction of $\vdash_{L(\mathbf{T})}$, $\text{AC}(f, F, f(P)) \models \varphi$. Since $\text{EF}(F, P) \xrightarrow{f} \text{AC}(f, F, f(P))$, it follows that $\text{EF}(F, P) \models \langle f \rangle \varphi = L(f@\varphi)$. Conversely, from $\text{EF}(F, P) \models \langle f \rangle \varphi$ we obtain P' with $\text{EF}(F, P) \xrightarrow{f} P'$ and $P' \models \varphi$. Since $\text{bn}_{L(\mathbf{T})}(f) = \emptyset$, there are no other alpha-variants of this transition. It follows that $P' = \text{AC}(f, F, f(P))$.

Case $A = f@\langle \alpha \rangle A'$: We then know $A' \in \mathcal{A}^{L(\alpha, F, f)/L}$ and $f \in F$ and $\text{bn}(\alpha)\#f, F$. Without loss of generality we assume also $\text{bn}(\alpha)\#P$, otherwise just find an alpha-variant of the formula where this holds.

Assume $P \models f@\langle \alpha \rangle A'$. Then $\exists P'. f(P) \xrightarrow{\alpha} P'$ and $P' \models A'$. We thus have $\text{EF}(L(\alpha, F, f), P') \models L(A')$ by induction. With $\text{AC}(f, F, f(P)) \xrightarrow{\alpha} \text{EF}(L(\alpha, F, f), P')$ we have $\text{AC}(f, F, f(P)) \models \langle \alpha \rangle L(A')$. Since $\text{EF}(F, P) \xrightarrow{f} \text{AC}(f, F, f(P))$, it follows that $\text{EF}(F, P) \models \langle f \rangle \langle \alpha \rangle L(A') = L(f@\langle \alpha \rangle A')$.

Conversely, from $\text{EF}(F, P) \models \langle f \rangle \langle \alpha \rangle L(A')$ we obtain P' with $\text{EF}(F, P) \xrightarrow{f} P'$ and $P' \models \langle \alpha \rangle L(A')$. Since $\text{bn}_{L(\mathbf{T})}(f) = \emptyset$, there are no other alpha-variants of this transition. It follows that $P' = \text{AC}(f, F, f(P))$. Note that $\text{bn}(\alpha)\#f, F, f(P)$. By construction of $L(\mathbf{T})$ we obtain P'' with $\text{AC}(f, F, f(P)) \xrightarrow{\alpha} \text{EF}(L(\alpha, F, f), P'')$ and $f(P) \xrightarrow{\alpha} P''$ and $\text{EF}(L(\alpha, F, f), P'') \models L(A')$. Thus $P'' \models A'$ by induction. Hence $f(P) \models \langle \alpha \rangle A'$. Hence $P \models f@\langle \alpha \rangle A'$. $\square$

As an immediate corollary we get:

Theorem 9. Let the L-transform be as above. Then $\mathrm{EF}(F,P) \mathrel{\dot\sim} \mathrm{EF}(F,Q) \Longrightarrow P \overset{F/L}{\sim} Q$.

Proof. The proof has been formalised in Isabelle. We first apply Theorem 1 to get $\mathrm{EF}(F,P) \mathrel{\dot=} \mathrm{EF}(F,Q)$, then Theorem 8 to get $\mathrm{EF}(F,P) \overset{F/L}{=} \mathrm{EF}(F,Q)$, and finally Theorem 7 to get $P \overset{F/L}{\sim} Q$. $\qquad\square$

The converse cannot be proved in the same way since L is not a surjection on formulas. Nevertheless it holds:

Theorem 10. Let the L-transform be as above. Then $P \overset{F/L}{\sim} Q \Longrightarrow \mathrm{EF}(F,P) \mathrel{\dot\sim} \mathrm{EF}(F,Q)$.

Proof. The proof has been formalised in Isabelle. Define a binary relation R on $\mathrm{STATES}_{L(\mathbf{T})}$ by including $(\mathrm{EF}(F,P), \mathrm{EF}(F,Q)) \in R$ and $(\mathrm{AC}(f,F,f(P)), \mathrm{AC}(f,F,f(Q))) \in R$ for all f, F, P, Q with $f \in F$ and $P \overset{F/L}{\sim} Q$. We prove that R is a bisimulation.

Symmetry is immediate since $\overset{F/L}{\sim}$ is symmetric. Now assume $R(S,T)$.

(1) Static implication. Assume $S \vdash_{L(\mathbf{T})} \varphi$. Then $S = \mathrm{AC}(f,F,f(P))$ for some F, $f \in F$ and P with $f(P) \vdash_{\mathbf{T}} \varphi$, and $T = \mathrm{AC}(f,F,f(Q))$ with $P \overset{F/L}{\sim} Q$. Thus $f(Q) \vdash_{\mathbf{T}} \varphi$, hence $T \vdash_{L(\mathbf{T})} \varphi$.

(2) Simulation. Assume $S \overset{\alpha}{\to} S'$. By construction of $L(\mathbf{T})$ there are two cases:

- $S = \mathrm{EF}(F,P) \overset{f}{\to} \mathrm{AC}(f,F,f(P)) = S'$ and $f \in F$. Then $T = \mathrm{EF}(F,Q)$ with $P \overset{F/L}{\sim} Q$. We get $T \overset{f}{\to} \mathrm{AC}(f,F,f(Q)) =: T'$ and $R(S',T')$. Note here that $\mathrm{bn}_{L(\mathbf{T})}(f) = \emptyset$.

- $S = \mathrm{AC}(f,F,f(P)) \overset{\alpha}{\to} \mathrm{EF}(L(\alpha,F,f),P') = S'$ and $f(P) \overset{\alpha}{\to} P'$ and $f \in F$. Then $T = \mathrm{AC}(f,F,f(Q))$ with $P \overset{F/L}{\sim} Q$. We may assume $\mathrm{bn}(\alpha)\#T$, hence also $\mathrm{bn}(\alpha)\#f(Q),F,f$. Then by simulation $\exists Q' . f(Q) \overset{\alpha}{\to} Q'$ and $P' \overset{L(\alpha,F,f)/L}{\sim} Q'$. Thus $T \overset{\alpha}{\to} \mathrm{EF}(L(\alpha,F,f),Q') =: T'$, and $R(S',T')$ as required. $\qquad\square$

A consequence is that for the variants of bisimilarity considered in Section 6.1 no new machinery is really needed; they can all be obtained by extending the transition relation. For the examples on the monadic pi-calculus, there would be substitution transitions which takes a state to another state where the substitution has been performed. Here the L function determines exactly which substitutions are applicable at which states, and by varying it we can obtain e.g. late, open and hyperbisimulation.

7. Unobservable actions

The logics and bisimulations considered so far are all of the strong variety, in the sense that all transitions are regarded as equally significant. In many models of concurrent computation there is a special action which is *unobservable* in the sense that in a bisimulation, and also in the definition of the action modalities, the presence of extra such transitions does not matter. This leads to notions of *weak* bisimulation and accompanying weak modal logics. For example, a process that has no transitions is weakly bisimilar with any process that has only unobservable transitions, and these satisfy the same weak modal logic

formulas. We shall here introduce these ideas into the nominal transition systems. One main source of complication over similar treatments in process algebras turns out to be the presence of state predicates.

To cater for unobservable transitions assume a special action τ with empty support. The following definitions are standard:

Definition 20.
  (1) $P \Rightarrow P'$ is defined by induction to mean $P = P'$ or $P \xrightarrow{\tau} \circ \Rightarrow P'$.
  (2) $P \xrightarrow{\alpha} P'$ means $P \Rightarrow \circ \xrightarrow{\alpha} \circ \Rightarrow P'$.
  (3) $P \xrightarrow{\hat{\alpha}} P'$ means $P \Rightarrow P'$ if $\alpha = \tau$ and $P \xrightarrow{\alpha} P'$ otherwise.

Intuitively $P \xrightarrow{\hat{\alpha}} P'$ means that P can evolve to P' through transitions with the only observable content α. We call this a weak action α and it will be the basis for the semantics in this section.

7.1. **Weak bisimilarity.** The normal way to define weak bisimilarity is to weaken $Q \xrightarrow{\alpha} Q'$ to $Q \xrightarrow{\hat{\alpha}} Q'$ in the simulation requirement. This results in the weak simulation criterion:

Definition 21. A binary relation R on STATES is a *weak simulation* if $R(P, Q)$ implies that for all α, P' with $\mathrm{bn}(\alpha)\#Q$ there exists Q' such that

$$\text{if } P \xrightarrow{\alpha} P' \text{ then } Q \xrightarrow{\hat{\alpha}} Q' \text{ and } R(P', Q')$$

However, just replacing the simulation requirement with weak simulation in Definition 2 will not suffice. The reason is that through the static implication criterion in Definition 2, an observer can still observe the state predicates directly, and thus distinguish between a state that satisfies φ and a state that does not but can silently evolve to another state that satisfies φ:

Example 1.

$$P \xrightarrow{\quad \tau \quad} \underset{\varphi}{Q}$$

Certainly $\{(P, Q), (Q, Q)\}$ is a weak simulation according to Definition 21. But $P \nvdash \varphi$ and $Q \vdash \varphi$, thus they are in no static implication. We argue that if φ is the *only* state predicate (in particular, there is no predicate $\neg\varphi$), then the only test that an observer can apply is "if φ then ...," and here P and Q will behave the same; P can pass the test after an unobservable delay. Thus P and Q should be deemed weakly bisimilar, and static implication as in Definition 2 is not appropriate.

Therefore we need a weak counterpart of static implication where τ transitions are admitted before checking predicates, that is, if $P \vdash \varphi$ then $Q \Rightarrow Q' \vdash \varphi$. In other words, Q can unobservably evolve to a state that satisfies φ. However, this is not quite enough by itself. Consider the following example where $P \vdash \varphi_0$, $P \vdash \varphi_1$, $R \vdash \varphi_1$ and $Q \vdash \varphi_0$, with transitions $P \xrightarrow{\tau} R$ and $Q \xrightarrow{\tau} R$:

Example 2.

$$P_{\substack{\varphi_0 \\ \varphi_1}} \xrightarrow{\;\tau\;} R_{\varphi_1} \xleftarrow{\;\tau\;} Q_{\varphi_0}$$
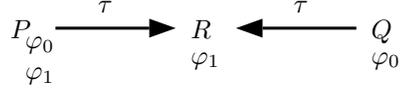
Here we do not want to regard P and Q as weakly bisimilar. They do have the same transitions and can satisfy the same predicates, possibly after a τ transition. But an observer of P can first determine that φ_1 holds, and then determine that φ_0 holds. This is not possible for Q: an observer who concludes φ_1 must already have evolved to R.

Similarly, consider the following example where the only difference between P and Q is that $P \vdash \varphi$ but not $Q \vdash \varphi$:

Example 3.

$$P_{\varphi} \xrightarrow{\alpha} P_0 \xleftarrow{\alpha} Q \qquad P_{\varphi} \xrightarrow{\tau} P_{1\,\varphi} \xleftarrow{\tau} Q$$
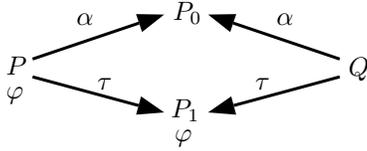
Again we do not want to regard P and Q as weakly bisimilar. Intuitively, an observer of Q who determines that φ holds must already be at P_1 and thus have preempted the possibility to do α, whereas for P, the predicate φ holds while retaining the possibility to do α. For instance, P in parallel with a process of kind "if φ then γ" can perform γ followed by α, but Q in parallel with the same cannot do that sequence.

In conclusion, the weak counterpart of static implication should allow the simulating state to proceed through unobservable actions to a state that *both* satisfies the same predicate *and* continues to bisimulate. This leads to the following:

Definition 22. A binary relation R on states is a *weak static implication* if $R(P, Q)$ implies that for all φ there exists Q' such that

$$\text{if } P \vdash \varphi \text{ then } Q \Rightarrow Q' \text{ and } Q' \vdash \varphi \text{ and } R(P, Q')$$

Definition 23. A *weak bisimulation* is a symmetric binary relation on states satisfying both weak simulation and weak static implication. We write $P \mathbin{\dot{\approx}} Q$ to mean that there exists a weak bisimulation R such that $R(P, Q)$.

In Example 1, $\{(P, Q), (Q, P), (Q, Q)\}$ is a weak bisimulation. In Examples 2 and 3, P and Q are not weakly bisimilar.

It is interesting to compare this with weak bisimilarities defined for psi-calculi [JBPV10]. A psi-calculus contains a construct of kind "if φ then $\dots$" to test if a state predicate is true. These constructs may be nested; for instance, "if φ_0 then if φ_1 then $\dots$" effectively tests if both φ_o and φ_1 are true simultaneously. If state predicates are closed under conjunction, Definition 23 coincides with the definition of simple weak bisimulation in [JBPV10]. In general, however, Definition 23 is less discriminating. Consider $P_0 \xrightarrow{\tau} P_1 \xrightarrow{\tau} P_0$ where for $i = 0, 1$: $P_i \vdash \varphi_i$. Compare it to Q with no transitions where both $Q \vdash \varphi_0$ and $Q \vdash \varphi_1$:

Example 4.

$$P_0 \xrightarrow{\tau} P_1 \qquad\qquad Q$$

Here all of P_0, P_1 and Q are weakly bisimilar, unless the predicates are closed under conjunction, in which case the predicate $\varphi_0 \wedge \varphi_1$ distinguishes between them. In psi-calculi Q would not be simply weakly bisimilar to P_0 or P_1 for the same reason.

We proceed to establish some expected properties of weak bisimilarity.

Lemma 15. If $P \stackrel{.}{\approx} Q$ and $P \stackrel{\hat{\alpha}}{\Rightarrow} P'$ with $\mathrm{bn}(\alpha)\#Q$ then for some Q' it holds $P' \stackrel{.}{\approx} Q'$ and $Q \stackrel{\hat{\alpha}}{\Rightarrow} Q'$.

Proof. The proof has been formalised in Isabelle. We first prove the lemma for the case $\alpha = \tau$ by induction on the derivation of $P \Rightarrow P'$. The base case is immediate. In the inductive step, we have $P \stackrel{\tau}{\rightarrow} P'' \Rightarrow P'$ for some P''. By weak simulation we obtain Q'' with $Q \Rightarrow Q''$ and $P'' \stackrel{.}{\approx} Q''$. It then follows from the induction hypothesis that there exists Q' with $Q'' \Rightarrow Q'$ and $P' \stackrel{.}{\approx} Q'$.

In case $\alpha \neq \tau$, we obtain P_1, P_2 with $P \Rightarrow P_1$ and $P_1 \stackrel{\alpha}{\rightarrow} P_2$ and $P_2 \Rightarrow P'$. By the previous case there exists Q_1 with $Q \Rightarrow Q_1$ and $P_1 \stackrel{.}{\approx} Q_1$. Without loss of generality we assume $\mathrm{bn}(\alpha)\#Q_1$, otherwise just find an alpha-variant of the transition where this holds. By weak simulation we obtain Q_2 with $Q_1 \stackrel{\hat{\alpha}}{\Rightarrow} Q_2$ and $P_2 \stackrel{.}{\approx} Q_2$. It then follows from the previous case that there exists Q' with $Q_2 \Rightarrow Q'$ and $P' \stackrel{.}{\approx} Q'$.   □

Lemma 16. $\stackrel{.}{\approx}$ is an equivariant equivalence relation.

Proof. The proof has been formalised in Isabelle. Equivariance is a simple calculation, based on the observation that if R is a weak bisimulation, then $\pi \cdot R$ is a weak bisimulation. To prove reflexivity of $\stackrel{.}{\approx}$, we note that equality is a weak bisimulation. Symmetry is immediate from Definition 23. To prove transitivity, we show that the composition of $\stackrel{.}{\approx}$ with itself is a bisimulation. Assume $P \stackrel{.}{\approx} S \stackrel{.}{\approx} Q$. The weak simulation requirement is proved by considering a transition $P \stackrel{\alpha}{\rightarrow} P'$ where $\mathrm{bn}(\alpha)$ is fresh for Q. Without loss of generality we also assume $\mathrm{bn}(\alpha)\#S$, otherwise just find an alpha-variant of the transition where this holds. Since $\stackrel{.}{\approx}$ satisfies weak simulation, we obtain S' with $S \stackrel{\hat{\alpha}}{\Rightarrow} S'$ and $P' \stackrel{.}{\approx} S'$. Since $S \stackrel{.}{\approx} Q$, Lemma 15 then implies that there exists Q' with $Q \stackrel{\hat{\alpha}}{\Rightarrow} Q'$ and $S' \stackrel{.}{\approx} Q'$. To prove weak static implication, assume $P \vdash \varphi$. We use $P \stackrel{.}{\approx} S$ and the fact that $\stackrel{.}{\approx}$ satisfies weak static implication to obtain some S' with $S \Rightarrow S'$ and $P \stackrel{.}{\approx} S'$ and $S' \vdash \varphi$. Since $S \stackrel{.}{\approx} Q$, Lemma 15 then implies that there exists Q' with $Q \Rightarrow Q'$ and $S' \stackrel{.}{\approx} Q'$. By weak static implication of $\stackrel{.}{\approx}$ we obtain Q'' with $Q' \Rightarrow Q''$ and $S' \stackrel{.}{\approx} Q''$ and $Q'' \vdash \varphi$. So we have $Q \Rightarrow Q''$ and $P \stackrel{.}{\approx}\circ\stackrel{.}{\approx} Q''$, as required.   □

7.2. **Weak logic.** We here define a Hennessy-Milner logic adequate for weak bisimilarity. Since weak bisimilarity identifies more states than strong bisimilarity, the logic needs to be correspondingly less expressive: it must not contain formulas that distinguish between

weakly bisimilar states. Our approach is to keep the definition of formulas (Definition 3) and identify an adequate sublogic.

One main idea is to restrict the action modalities $\langle \alpha \rangle$ to occur only in accordance with the requirement of a weak bisimulation, thus checking for $\overset{\hat{\alpha}}{\Rightarrow}$ rather than for $\overset{\alpha}{\rightarrow}$. We therefore define the derived *weak action* modal operator $\langle\!\langle \alpha \rangle\!\rangle$ in the following way, where $\langle \tau \rangle^i A$ is defined to mean A if $i = 0$ and $\langle \tau \rangle \langle \tau \rangle^{i-1} A$ otherwise.

Definition 24 (Weak action modality).

$$\langle\!\langle \tau \rangle\!\rangle A = \bigvee_{i \in \omega} \langle \tau \rangle^i A \qquad\qquad \langle\!\langle \alpha \rangle\!\rangle A = \langle\!\langle \tau \rangle\!\rangle \langle \alpha \rangle \langle\!\langle \tau \rangle\!\rangle A \quad \text{for } \alpha \neq \tau$$

Note that in $\langle\!\langle \alpha \rangle\!\rangle A$ the names in $\mathrm{bn}(\alpha)$ bind into A. As usual we consider formulas up to alpha-conversion in the standard sense, i.e., to prove a property of a formula it is enough to prove a property of an alpha-variant. It is then straightforward to show (and formalize in Isabelle) that $\langle\!\langle \alpha \rangle\!\rangle A$ corresponds to the weak transitions used in the definition of weak bisimilarity:

Proposition 5. Assume $\mathrm{bn}(\alpha)\#P$. Then

$$P \models \langle\!\langle \alpha \rangle\!\rangle A \quad \text{iff} \quad \exists P'.\, P \overset{\hat{\alpha}}{\Rightarrow} P' \text{ and } P' \models A$$

In particular, for $\alpha = \tau$, we have that $\langle\!\langle \tau \rangle\!\rangle A$ holds iff A holds after zero or more τ transitions.

Thus a first step towards a weak sublogic is to replace $\langle \alpha \rangle$ by $\langle\!\langle \alpha \rangle\!\rangle$ in Definition 3. By itself this is not enough; that sublogic is still too expressive. For instance, the formula φ asserts that φ holds in a state; this holds for Q but not for P in Example 1, even though they are weakly bisimilar.

To disallow φ as a weak formula we require that state predicates only occur guarded by a weak action $\langle\!\langle \tau \rangle\!\rangle$. This solves part of the problem. In Example 1 we can no longer use φ as a formula, and the formula $\langle\!\langle \tau \rangle\!\rangle \varphi$ holds of both P and Q. Still, in Example 1 there would be the formula $\langle\!\langle \tau \rangle\!\rangle \neg \varphi$ which holds for P but not for Q, and in Example 4 the formula $\langle\!\langle \tau \rangle\!\rangle (\varphi_0 \wedge \varphi_1)$ holds for Q but not for P_0. Clearly a logic adequate for weak bisimulation cannot have such formulas. The more draconian restriction that state predicates occur *immediately* under $\langle\!\langle \tau \rangle\!\rangle$ would indeed disallow both $\langle\!\langle \tau \rangle\!\rangle \neg \varphi$ and $\langle\!\langle \tau \rangle\!\rangle (\varphi_0 \wedge \varphi_1)$ but would also disallow any formula distinguishing between P and Q in Examples 2 and 3.

A solution is to allow state predicates under $\langle\!\langle \tau \rangle\!\rangle$, and never directly under negation or in conjunction with another state predicate. The logic is:

Definition 25 (Weak formulas). The set of *weak formulas* is the sublogic of Definition 3 given by

$$A \quad ::= \quad \bigwedge_{i \in I} A_i \quad | \quad \neg A \quad | \quad \langle\!\langle \alpha \rangle\!\rangle A \quad | \quad \langle\!\langle \tau \rangle\!\rangle (A \wedge \varphi)$$

Note that since $P \overset{\hat{\alpha}}{\Rightarrow} \circ \Rightarrow P'$ holds iff $P \overset{\hat{\alpha}}{\Rightarrow} P'$ we have that $\langle\!\langle \alpha \rangle\!\rangle \langle\!\langle \tau \rangle\!\rangle A$ is logically equivalent to $\langle\!\langle \alpha \rangle\!\rangle A$. We thus abbreviate $\langle\!\langle \alpha \rangle\!\rangle \langle\!\langle \tau \rangle\!\rangle (A \wedge \varphi)$ to $\langle\!\langle \alpha \rangle\!\rangle (A \wedge \varphi)$. We also abbreviate $\langle\!\langle \alpha \rangle\!\rangle (\top \wedge \varphi)$ to $\langle\!\langle \alpha \rangle\!\rangle \varphi$.

Compared to Definition 3, the state predicates can now only occur in formulas of the form $\langle\!\langle \tau \rangle\!\rangle (A \wedge \varphi)$, i.e., under a weak action, and not under negation or conjunction with another predicate. For instance, in Example 1 above, neither φ nor $\langle\!\langle \tau \rangle\!\rangle \neg \varphi$ are weak formulas, and in fact there is no weak formula to distinguish between P and Q. Similarly,

in Example 4 $\langle\!\langle\tau\rangle\!\rangle(\varphi_0 \wedge \varphi_1)$ is not a weak formula, and no weak formula distinguishes between Q and P_i.

To argue that the logic still is expressive enough to provide distinguishing formulas for states that are not weakly bisimilar, consider Example 2 and the formula $\langle\!\langle\tau\rangle\!\rangle((\langle\!\langle\tau\rangle\!\rangle\varphi_0)\wedge\varphi_1)$ which holds for P but not for Q. Similarly, in Example 3 $\langle\!\langle\tau\rangle\!\rangle((\langle\!\langle\alpha\rangle\!\rangle\top)\wedge\varphi)$ holds for P but not for Q.

Definition 26. Two states P and Q are *weakly logically equivalent*, written $P \stackrel{.}{\equiv} Q$, if for all weak formulas A it holds that $P \models A$ iff $Q \models A$.

7.3. **Logical adequacy and expressive completeness.** We show that the logic defined in Section 7.2 is adequate for weak bisimilarity. Moreover, every finitely supported set of states that is closed under weak bisimilarity can be described by a weak formula.

Theorem 11. $P \stackrel{.}{\approx} Q \implies P \stackrel{.}{\equiv} Q$

Proof. The proof has been formalised in Isabelle. We prove by induction over weak formulas that $P \stackrel{.}{\approx} Q$ implies that $P \models A$ iff $Q \models A$. The cases for conjunction and negation are immediate by induction.

Case $\langle\!\langle\alpha\rangle\!\rangle A$: Assume $P \models \langle\!\langle\alpha\rangle\!\rangle A$. Then for some alpha-variant $\langle\!\langle\alpha'\rangle\!\rangle A' = \langle\!\langle\alpha\rangle\!\rangle A$, $\exists P'.\, P \stackrel{\hat{\alpha'}}{\Rightarrow} P'$ and $P' \models A'$. Without loss of generality we assume also $\mathrm{bn}(\alpha')\#Q$, otherwise just find an alpha-variant of $\langle\alpha'\rangle A'$ where this holds. Then, by Lemma 15, $\exists Q'.\, Q \stackrel{\hat{\alpha'}}{\Rightarrow} Q'$ and $P' \stackrel{.}{\approx} Q'$. By induction and $P' \models A'$ we get $Q' \models A'$, hence by definition $Q \models \langle\!\langle\alpha\rangle\!\rangle A$. The proof of $Q \models \langle\!\langle\alpha\rangle\!\rangle A \implies P \models \langle\!\langle\alpha\rangle\!\rangle A$ is symmetric, using the fact that $P \stackrel{.}{\approx} Q$ entails $Q \stackrel{.}{\approx} P$.

Case $A = \langle\!\langle\tau\rangle\!\rangle(B \wedge \varphi)$. Assume $P \stackrel{.}{\approx} Q$ and $P \models A$. Then $P \Rightarrow P'$ such that $P' \models B$ and $P' \vdash \varphi$. By $P \stackrel{.}{\approx} Q$ and Lemma 15 we obtain Q' such that $Q \Rightarrow Q'$ and $P' \stackrel{.}{\approx} Q'$. By weak static implication we obtain Q'' with $Q' \Rightarrow Q''$ and $P' \stackrel{.}{\approx} Q''$ and $Q'' \vdash \varphi$. By induction $Q'' \models B$, thus $Q'' \models B \wedge \varphi$. From $Q \Rightarrow Q' \Rightarrow Q''$ we have $Q \Rightarrow Q''$. Thus $Q \models \langle\!\langle\tau\rangle\!\rangle(B \wedge \varphi)$ as required.

The proof that $Q \models A$ implies $P \models A$ is symmetric, using the fact that $P \stackrel{.}{\approx} Q$ entails $Q \stackrel{.}{\approx} P$. $\qquad\square$

Lemma 17. If $P \stackrel{.}{\not\equiv} Q$ then there exists a distinguishing weak formula B for P and Q such that $\mathrm{supp}(B) \subseteq \mathrm{supp}(P)$.

Proof. The proof has been formalised in Isabelle. Since $P \stackrel{.}{\not\equiv} Q$ there is a distinguishing weak formula A for P and Q. Let Π_P be the set of name permutations that leave $\mathrm{supp}(P)$ invariant and choose $B = \bigwedge\{\pi \cdot A \,|\, \pi \in \Pi_P\}$. In the terminology of Pitts [Pit13] ch. 5, B is the conjunction of $\mathrm{hull}_{\mathrm{supp}(P)}A$; this set is supported by $\mathrm{supp}(P)$ (but not uniformly bounded). Because $\models$ is equivariant we get $P \models \pi \cdot A$ for all conjuncts $\pi \cdot A$ of B, and since $Q \not\models A = \mathrm{id} \cdot A$ we get $Q \not\models B$. $\qquad\square$

Theorem 12. $P \mathrel{\dot{\equiv}} Q \implies P \mathrel{\dot{\approx}} Q$

Proof. The proof has been formalised in Isabelle. We establish that $\mathrel{\dot{\equiv}}$ is a weak bisimulation. Obviously it is symmetric. So assume $P \mathrel{\dot{\equiv}} Q$, we need to prove the two requirements on a weak bisimulation.

(1) Weak static implication. The proof is by contradiction. Assume that $\mathrel{\dot{\equiv}}$ does not satisfy the weak static implication requirement. Then there exist P, Q, φ such that $P \mathrel{\dot{\equiv}} Q$ and $P \vdash \varphi$ and for all Q' such that $Q \Rightarrow Q'$ and $Q' \vdash \varphi$ there exists a distinguishing formula $A_{Q'}$ such that $P \models A_{Q'}$ and not $Q' \models A_{Q'}$. By Lemma 17 $\mathrm{supp}(A_{Q'}) \subseteq \mathrm{supp}(P)$, which means that the infinite conjunction A of all these $A_{Q'}$ is well formed. We thus have that $\langle\!\langle \tau \rangle\!\rangle (A \wedge \varphi)$ is a distinguishing formula for P and Q, contradicting $P \mathrel{\dot{\equiv}} Q$.

(2) Weak simulation. The proof is by contradiction. Assume that $\mathrel{\dot{\equiv}}$ does not satisfy the weak simulation requirement. Then there exist P, Q, P', α with $\mathrm{bn}(\alpha)\#Q$ such that $P \mathrel{\dot{\equiv}} Q$ and $P \xrightarrow{\alpha} P'$ and, letting $\mathcal{Q} = \{Q' \mid Q \stackrel{\hat{\alpha}}{\Rightarrow} Q'\}$, for all $Q' \in \mathcal{Q}$ it holds that $P' \mathrel{\dot{\not\equiv}} Q'$. Assume $\mathrm{bn}(\alpha)\#P$, otherwise just find an alpha-variant of the transition satisfying this. By $P' \mathrel{\dot{\not\equiv}} Q'$, for all $Q' \in \mathcal{Q}$ there exists a weak distinguishing formula for P' and Q'. The formula may depend on Q', and by Lemma 17 we can find such a distinguishing formula $B_{Q'}$ for P' and Q' with $\mathrm{supp}(B_{Q'}) \subseteq \mathrm{supp}(P')$. Let B be the conjunction of all $B_{Q'}$. We thus get that $P \models \langle\!\langle \alpha \rangle\!\rangle B$ but not $Q \models \langle\!\langle \alpha \rangle\!\rangle B$, contradicting $P \mathrel{\dot{\equiv}} Q$.

$\square$

We omit the proofs of the following results (which have also been formalised in Isabelle), as they are similar to the corresponding proofs for the full logic in Section 3.3. In addition, we need to verify that the constructions used below yield weak formulas. This is immediate from Definition 25.

Lemma 18. If $P \mathrel{\dot{\not\equiv}} Q$, write $B_{P,Q}$ for a distinguishing weak formula for P and Q such that $\mathrm{supp}(B_{P,Q}) \subseteq \mathrm{supp}(P)$. Then

$$D(P, Q) := \bigwedge_\pi \pi^{-1} \cdot B_{\pi \cdot P, \pi \cdot Q}$$

defines a distinguishing weak formula for P and Q with support included in $\mathrm{supp}(P)$. Moreover, the function D is equivariant.

Definition 27. A *characteristic weak formula* for P is a weak formula A such that for all Q, $P \mathrel{\dot{\approx}} Q$ iff $Q \models A$.

Lemma 19. Let D be defined as in Lemma 18. The formula $\mathrm{Char}(P) := \bigwedge_{P \mathrel{\dot{\not\equiv}} Q} D(P, Q)$ is a characteristic weak formula for P.

Lemma 20. Let $\mathrm{Char(P)}$ be defined as in Lemma 19. The function $\mathrm{Char} \colon P \mapsto \mathrm{Char}(P)$ is equivariant.

Theorem 13 (Weak Expressive Completeness)**.** Let S be a finitely supported set of states that is closed under weak bisimilarity, i.e., for all $P \in S$ and Q, $P \mathrel{\dot{\approx}} Q$ implies $Q \in S$. Then $P \in S$ iff $P \models \bigvee_{P' \in S} \mathrm{Char}(P')$.

7.4. **Disjunction elimination.** As defined in Section 4, disjunction is a derived logical operator, expressed through conjunction and negation. This is still true in the weak modal logic, but there is a twist in that neither general conjunctions nor negations may be applied to unguarded state predicates. The examples in Section 7.1 demonstrate why this restriction is necessary: negated or conjoined state predicates in formulas would mean that adequacy no longer holds. Interestingly, we can allow disjunctions of unguarded predicates while maintaining adequacy; in fact, adding disjunction would not increase the expressive power of the logic. In this section we demonstrate this. An uninterested reader may skip this section without loss of continuity.

The *extended* weak logic is as follows, where a simultaneous induction defines both extended weak formulas (ranged over by E) and preformulas (ranged over by B) corresponding to subformulas with unguarded state predicates.

Definition 28 (Extended weak formulas E and preformulas B).

$$E \quad ::= \quad \bigwedge_{i \in I} E_i \quad | \quad \neg E \quad | \quad \langle\!\langle \alpha \rangle\!\rangle E \quad | \quad \langle\!\langle \tau \rangle\!\rangle B$$

$$B \quad ::= \quad E \wedge B \quad | \quad \varphi \quad | \quad \bigvee_{i \in I} B_i$$

The last clause in the definition of preformulas is what distinguishes this logic from the logic in Definition 25. (Thus an extended weak formula is also an ordinary weak formula if it does not contain a disjunction of unguarded state predicates.) For instance, $\langle\!\langle \tau \rangle\!\rangle (\varphi_0 \vee \varphi_1)$ is an extended weak formula, as is

$$\langle\!\langle \tau \rangle\!\rangle (((\langle\!\langle \beta \rangle\!\rangle \top) \wedge \varphi_0) \vee ((\langle\!\langle \gamma \rangle\!\rangle \top) \wedge \varphi_1))$$

saying that it is possible to do a sequence of unobservable actions such that either continuing with β and satisfying φ_0 hold, or continuing with γ and satisfying φ_1 hold.

Lemma 21. $\langle\!\langle \tau \rangle\!\rangle \bigvee_{i \in I} B_i$ is a formula iff $\bigvee_{i \in I} \langle\!\langle \tau \rangle\!\rangle B_i$ is, and in this case

$$\langle\!\langle \tau \rangle\!\rangle \bigvee_{i \in I} B_i \stackrel{.}{\equiv} \bigvee_{i \in I} \langle\!\langle \tau \rangle\!\rangle B_i$$

Proof. The proof is by directly expanding definitions, noting that $\mathrm{supp}(\{\langle\!\langle \tau \rangle\!\rangle B_i \mid i \in I\})$ is equal to $\mathrm{supp}(\{B_i \mid i \in I\})$. □

Theorem 14. For any extended weak formula E there is an (ordinary) weak formula $\Delta(E)$ such that $E \stackrel{.}{\equiv} \Delta(E)$.

Proof. The idea is to push disjunctions in preformulas to top level using the fact that (finite) conjunction distributes over disjunction, and then use Lemma 21. Say that a preformula is in *normal form* if it either is $A \wedge \varphi$ where A here and in the following stands for an ordinary weak formula, or is a disjunction of normal forms. We let C range over normal preformulas, thus

$$C \quad ::= \quad A \wedge \varphi \quad | \quad \bigvee_{i \in I} C_i$$

The intuition is that in a normal preformula no conjunction can contain a disjunction of preformulas.

Define a function δ from normal preformulas to (ordinary) weak formulas by

$$\delta(A \wedge \varphi) \quad = \quad \langle\!\langle \tau \rangle\!\rangle (A \wedge \varphi)$$
$$\delta(\bigvee_{i \in I} C_i) \quad = \quad \bigvee_{i \in I} \delta(C_i)$$

By induction δ is equivariant, hence the disjunction in the above definition is finitely supported. Moreover, $\delta(C) \stackrel{.}{\equiv} \langle\!\langle\tau\rangle\!\rangle C$ by Lemma 21 and induction.

Next, define a binary function ε from pairs consisting of an (ordinary) weak formula and a normal preformula, to normal preformulas, by

$$\begin{aligned}
\varepsilon(A, A' \wedge \varphi) &= (A \wedge A') \wedge \varphi \\
\varepsilon(A, \textstyle\bigvee_{i \in I} C_i) &= \textstyle\bigvee_{i \in I} \varepsilon(A, C_i)
\end{aligned}$$

Also ε is equivariant, and $\varepsilon(A, C) \stackrel{.}{\equiv} A \wedge C$ by induction and distributivity of conjunction.

We now provide an explicit transformation Δ from extended weak formulas to (ordinary) weak formulas. We also provide a transformation Δ_{pre} that maps preformulas to normal preformulas. The transformations Δ and Δ_{pre} are defined by mutual recursion:

$$\begin{aligned}
\Delta(\textstyle\bigwedge_{i \in I} E_i) &= \textstyle\bigwedge_{i \in I} \Delta(E_i) \\
\Delta(\neg E) &= \neg\Delta(E) \\
\Delta(\langle\!\langle\alpha\rangle\!\rangle E) &= \langle\!\langle\alpha\rangle\!\rangle\Delta(E) \\
\Delta(\langle\!\langle\tau\rangle\!\rangle B) &= \delta(\Delta_{\mathrm{pre}}(B)) \\[6pt]
\Delta_{\mathrm{pre}}(E \wedge B) &= \varepsilon(\Delta(E), \Delta_{\mathrm{pre}}(B)) \\
\Delta_{\mathrm{pre}}(\varphi) &= \top \wedge \varphi \\
\Delta_{\mathrm{pre}}(\textstyle\bigvee_{i \in I} B_i) &= \textstyle\bigvee_{i \in I} \Delta_{\mathrm{pre}}(B_i)
\end{aligned}$$

By simultaneous induction over extended weak formulas and preformulas it is easy to prove that both Δ and Δ_{pre} are equivariant. Therefore, all conjunctions and disjunctions that appear on the right-hand side in the above definition are finitely supported.

We now prove, again by simultaneous induction over extended weak formulas E and preformulas B, that $\Delta(E) \stackrel{.}{\equiv} E$ and $\Delta_{\mathrm{pre}}(B) \stackrel{.}{\equiv} B$. The cases $\Delta(\bigwedge_{i \in I} E_i)$ and $\Delta(\neg E)$ and $\Delta(\langle\!\langle\alpha\rangle\!\rangle E)$ are immediate by induction. For the case $\Delta(\langle\!\langle\tau\rangle\!\rangle B)$ we have

$$\Delta(\langle\!\langle\tau\rangle\!\rangle B) = \delta(\Delta_{\mathrm{pre}}(B)) \stackrel{.}{\equiv} \langle\!\langle\tau\rangle\!\rangle\Delta_{\mathrm{pre}}(B) \stackrel{.}{\equiv} \langle\!\langle\tau\rangle\!\rangle B$$

by induction. Likewise, for the case $\Delta_{\mathrm{pre}}(E \wedge B)$ we have

$$\Delta_{\mathrm{pre}}(E \wedge B) = \varepsilon(\Delta(E), \Delta_{\mathrm{pre}}(B)) \stackrel{.}{\equiv} \Delta(E) \wedge \Delta_{\mathrm{pre}}(B) \stackrel{.}{\equiv} E \wedge B$$

by induction. The case $\Delta_{\mathrm{pre}}(\varphi)$ is trivial, and the case $\Delta_{\mathrm{pre}}(\bigvee_{i \in I} B_i)$ is again immediate by induction. $\qquad\square$

8. State predicates versus actions

Transition system formalisms differ in how much information is considered to reside in states and how much is considered to reside in actions. One extreme is Lamport's TLA [Lam02] where all information is in the states. On the other extreme are most process algebras such as the pi-calculus, where states contain no information apart from the outgoing transitions. Advanced process algebras such as psi-calculi use both state predicates and actions. Clearly, many ways are possible and the choice is more dependent on the traditions and modelling convenience in different areas than on hard theoretical results.

The question if state information can be encoded as actions and vice versa is old. Already in 1989 Jonsson et al. [JKP90] provided a translation from a labelled transition system into an (unlabelled) Kripke structure, in order to use a model checker for CTL. With nominal transition systems it is hard to see how such a translation could work. The actions can contain binding names and encoding that into the states would require a substantial extension of our definitions.

In the other direction, it has long been a folklore fact in process algebra that a state predicate can be represented as an action on a transition leading back to the same state. In this section we make this idea fully formal: for the purposes of checking bisimilarity this transformation is indeed sound, and there is a companion transformation on modal logic formulas. This is straightforward for strong bisimulation and modal logic, and a little more involved for weak modal formulas.

8.1. **Strong bisimulation and logic.** Our idea is that for any transition system $\mathbf{T}$ there is another transition system $\mathcal{S}(\mathbf{T})$ where state predicates are replaced by self-loops. To formulate this idea we again use the notation $\text{STATES}_{\mathbf{T}}$ to mean the states in the transition system $\mathbf{T}$, and similarly for actions, bn, transitions, bisimilarity, etc.

Definition 29. The function $\mathcal{S}$ from transition systems to transition systems is defined as follows:

- $\text{STATES}_{\mathcal{S}(\mathbf{T})} = \text{STATES}_{\mathbf{T}}$
- $\text{ACT}_{\mathcal{S}(\mathbf{T})} = \text{ACT}_{\mathbf{T}} \uplus \text{PRED}_{\mathbf{T}}$
- $\text{bn}_{\mathcal{S}(\mathbf{T})}(\alpha) = \text{bn}_{\mathbf{T}}(\alpha)$ if $\alpha \in \text{ACT}_{\mathbf{T}}$; $\text{bn}_{\mathcal{S}(\mathbf{T})}(\varphi) = \emptyset$ if $\varphi \in \text{PRED}_{\mathbf{T}}$
- $\text{PRED}_{\mathcal{S}(\mathbf{T})} = \vdash_{\mathcal{S}(\mathbf{T})} = \emptyset$
- $P \xrightarrow{\alpha}_{\mathcal{S}(\mathbf{T})} P'$ if $P \xrightarrow{\alpha}_{\mathbf{T}} P'$ (for $\alpha \in \text{ACT}_{\mathbf{T}}$); $P \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} P$ if $P \vdash_{\mathbf{T}} \varphi$ (for $\varphi \in \text{PRED}_{\mathbf{T}}$)

It is easy to see that if $\mathbf{T}$ is a transition system then so is $\mathcal{S}(\mathbf{T})$. In particular equivariance of $\rightarrow_{\mathcal{S}(\mathbf{T})}$ follows from equivariance of $\rightarrow_{\mathbf{T}}$ and $\vdash_{\mathbf{T}}$ and the fact that the union of equivariant relations is equivariant.[2]

Theorem 15. $P \mathrel{\dot\sim}_{\mathbf{T}} Q \implies P \mathrel{\dot\sim}_{\mathcal{S}(\mathbf{T})} Q$

Proof. The proof has been formalised in Isabelle. We show that $\mathrel{\dot\sim}_{\mathbf{T}}$ is a $\mathcal{S}(\mathbf{T})$-bisimulation. Obviously it is symmetric. Static implication is trivial since $\vdash_{\mathcal{S}(\mathbf{T})}$ is false everywhere. For simulation, assume $P \mathrel{\dot\sim}_{\mathbf{T}} Q$ and that P has a transition in $\mathcal{S}(\mathbf{T})$. If the transition has an action in $\mathbf{T}$, then P has the same transition in $\mathbf{T}$. If the action is a state predicate in $\mathbf{T}$ then the transition must be $P \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} P$ where $P \vdash_{\mathbf{T}} \varphi$, hence by static implication $Q \vdash_{\mathbf{T}} \varphi$, thus also $Q \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} Q$. In both cases we therefore have a simulating transition from Q in $\mathcal{S}(\mathbf{T})$. $\square$

Theorem 16. $P \mathrel{\dot\sim}_{\mathcal{S}(\mathbf{T})} Q \implies P \mathrel{\dot\sim}_{\mathbf{T}} Q$

Proof. The proof has been formalised in Isabelle. We show that $\mathrel{\dot\sim}_{\mathcal{S}(\mathbf{T})}$ is a $\mathbf{T}$-bisimulation. Obviously it is symmetric. Assume $P \mathrel{\dot\sim}_{\mathcal{S}(\mathbf{T})} Q$. For static implication, if $P \vdash_{\mathbf{T}} \varphi$ then P has a transition with action φ in $\mathcal{S}(\mathbf{T})$, thus Q has a transition with the same action, hence $Q \vdash_{\mathbf{T}} \varphi$. For simulation, any transition in $\mathbf{T}$ is also a transition in $\mathcal{S}(\mathbf{T})$. $\square$

---

[2]Nominal Isabelle does not admit empty types. Therefore, the type $\text{PRED}_{\mathcal{S}(\mathbf{T})}$ is given by the unit type in our Isabelle formalisation. We still define $\vdash_{\mathcal{S}(\mathbf{T})}$ to be false everywhere, so that this minor difference has no further significance.

Modal logic formulas over $\mathcal{S}(\mathbf{T})$ use predicates as actions. We extend $\mathcal{S}$ to formulas as follows.

Definition 30. The function $\mathcal{S}$ from formulas over the transition system $\mathbf{T}$ to formulas over the transition system $\mathcal{S}(\mathbf{T})$ is defined by

$$\mathcal{S}(\varphi) = \langle \varphi \rangle \top$$

and is homomorphic on the other cases in Definition 3.

Theorem 17. $P \models_{\mathbf{T}} A$ iff $P \models_{\mathcal{S}(\mathbf{T})} \mathcal{S}(A)$

Proof. The proof has been formalised in Isabelle. It is by induction over formulas A, for arbitrary P. If $A = \varphi$ then clearly $P \models_{\mathbf{T}} \varphi$ iff $P \models_{\mathcal{S}(\mathbf{T})} \langle \varphi \rangle \top$. For the case $A = \langle \alpha \rangle A'$ we assume without loss of generality $\mathrm{bn}(\alpha)\#P$, otherwise just find an alpha-variant of the formula where this holds. This and the other cases are then immediate by induction. $\square$

This provides an alternate proof of Theorem 16. If $P \dot{\sim}_{\mathcal{S}(\mathbf{T})} Q$ then by Theorem 1 for any A it holds that $P \models \mathcal{S}(A)$ iff $Q \models \mathcal{S}(A)$, and thus $P \models A$ iff $Q \models A$, therefore by Theorem 2 $P \dot{\sim}_{\mathbf{T}} Q$. The converse does not follow since $\mathcal{S}$ is not surjective on formulas.

8.2. **Weak bisimulation.** For weak bisimulation the corresponding proofs are only a little bit more complicated.

Theorem 18. $P \dot{\approx}_{\mathbf{T}} Q \implies P \dot{\approx}_{\mathcal{S}(\mathbf{T})} Q$

Proof. The proof has been formalised in Isabelle. We show that $\dot{\approx}_{\mathbf{T}}$ is a weak $\mathcal{S}(\mathbf{T})$-bisimulation. Obviously it is symmetric. Weak static implication is trivial since $\vdash_{\mathcal{S}(\mathbf{T})}$ is false everywhere. For weak simulation, assume $P \dot{\approx}_{\mathbf{T}} Q$. There are two cases.

(1) $P \xrightarrow{\alpha}_{\mathcal{S}(\mathbf{T})} P'$ where $\alpha \in \mathrm{ACT}_{\mathbf{T}}$. Then $P \xrightarrow{\alpha}_{\mathbf{T}} P'$ and by $P \dot{\approx}_{\mathbf{T}} Q$ we get $Q \xRightarrow{\hat{\alpha}}_{\mathbf{T}} Q'$, which implies $Q \xRightarrow{\hat{\alpha}}_{\mathcal{S}(\mathbf{T})} Q'$, with $P' \dot{\approx}_{\mathbf{T}} Q'$.

(2) $P \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} P$ where $\varphi \in \mathrm{PRED}_{\mathbf{T}}$ and $P \vdash_{\mathbf{T}} \varphi$. By weak static implication $Q \Rightarrow_{\mathbf{T}} Q'$ and $Q' \vdash_{\mathbf{T}} \varphi$ and $Q' \dot{\approx}_{\mathbf{T}} P$. Thus $Q \Rightarrow_{\mathcal{S}(\mathbf{T})} Q'$ and $Q' \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} Q'$, hence $Q \xRightarrow{\hat{\varphi}}_{\mathcal{S}(\mathbf{T})} Q'$. $\square$

The converse uses the following lemma, which is familiar in many process algebras and interesting in its own right.

Lemma 22. If $P \Rightarrow Q \Rightarrow R$ and $P \dot{\approx} R$ then $Q \dot{\approx} R$.

Proof. The proof has been formalised in Isabelle. We establish that $\{(Q, R), (R, Q)\} \cup \dot{\approx}$ is a weak bisimulation. Obviously it is symmetric.

For weak static implication, clearly if $R \vdash \varphi$ then $Q \Rightarrow R \vdash \varphi$; and if $Q \vdash \varphi$ then $P \Rightarrow Q \vdash \varphi$, so by $P \dot{\approx} R$, Lemma 15, and weak static implication we obtain R' with $R \Rightarrow R' \vdash \varphi$ as required.

For weak simulation, clearly if $R \xrightarrow{\alpha} R'$ then $Q \Rightarrow R \xrightarrow{\alpha} R'$, hence $Q \xRightarrow{\hat{\alpha}} R'$ (and obviously $R' \dot{\approx} R'$); and if $Q \xrightarrow{\alpha} Q'$ then $P \xRightarrow{\hat{\alpha}} Q'$, so again by $P \dot{\approx} R$ and Lemma 15 we obtain R' with $R \xRightarrow{\hat{\alpha}} R'$ and $Q' \dot{\approx} R'$. $\square$

Theorem 19. $P \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q \implies P \mathrel{\dot{\approx}}_{\mathbf{T}} Q$

Proof. The proof has been formalised in Isabelle. We show that $\mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})}$ is a weak **T**-bisimulation. Obviously it is symmetric.

For weak static implication, assume $P \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q$ and $P \vdash_{\mathbf{T}} \varphi$. Thus $P \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} P$. By weak simulation we obtain Q' with $Q \xRightarrow{\hat{\varphi}}_{\mathcal{S}(\mathbf{T})} Q'$ and $P \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q'$. By construction of $\mathcal{S}(\mathbf{T})$ this means that there is a Q'' such that $Q \Rightarrow_{\mathcal{S}(\mathbf{T})} Q'' \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} Q'' \Rightarrow_{\mathcal{S}(\mathbf{T})} Q'$. Thus $Q \Rightarrow_{\mathbf{T}} Q''$ and $Q'' \vdash_{\mathbf{T}} \varphi$. Moreover, by applying Lemma 22 to $Q \Rightarrow_{\mathcal{S}(\mathbf{T})} Q'' \Rightarrow_{\mathcal{S}(\mathbf{T})} Q'$ and $Q \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} P \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q'$ we get $P \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q''$ as required.

For weak simulation, assume $P \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q$ and $P \xrightarrow{\alpha}_{\mathbf{T}} P'$. This implies $P \xrightarrow{\alpha}_{\mathcal{S}(\mathbf{T})} P'$, and thus by weak simulation $Q \xRightarrow{\hat{\alpha}}_{\mathcal{S}(\mathbf{T})} Q'$, which implies $Q \xRightarrow{\hat{\alpha}}_{\mathbf{T}} Q'$, with $P' \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q'$ as required. $\qquad\square$

8.3. **Weak logic.** For the weak logic, the correspondence between state predicates and actions is less obvious, and there appear to be more than one alternative. The transformation $\mathcal{S}$ of formulas does not preserve the property of being a weak formula since

$$\mathcal{S}(\langle\!\langle\tau\rangle\!\rangle(A \wedge \varphi)) = \langle\!\langle\tau\rangle\!\rangle(\mathcal{S}(A) \wedge \langle\varphi\rangle\top)$$

and the subformula $\langle\varphi\rangle\top$ is not weak. For the transition system $\mathcal{S}(\mathbf{T})$ we thus get an alternative logic adequate for weak bisimilarity by taking the formulas $\{\mathcal{S}(A) \mid A \text{ is a weak formula}\}$, i.e., with the last clause of Definition 25 replaced by formulas of kind $\langle\!\langle\tau\rangle\!\rangle(A \wedge \langle\varphi\rangle\top)$.

Corollary 2. $P \mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})} Q$ iff for all weak formulas A: $P \models_{\mathcal{S}(\mathbf{T})} \mathcal{S}(A)$ iff $Q \models_{\mathcal{S}(\mathbf{T})} \mathcal{S}(A)$

Proof. The proof has been formalised in Isabelle. It is immediate by Theorems 11, 12 and 17–19. $\qquad\square$

For this corollary to hold it is critical that the transition system is an image of $\mathcal{S}$, i.e., that the transitions involving actions φ only occur in loops. If this is not the case, there may be formulas in the image of $\mathcal{S}$ that can distinguish between weakly bisimilar states. An example is the following:
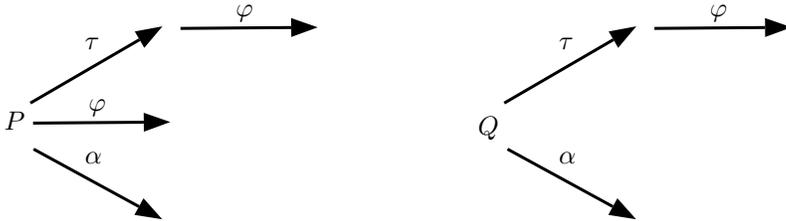
Here P and Q are weakly bisimilar. Let A be the weak formula $\langle\!\langle\tau\rangle\!\rangle(\langle\!\langle\alpha\rangle\!\rangle\top \wedge \varphi)$. Then $\mathcal{S}(A) = \langle\!\langle\tau\rangle\!\rangle(\langle\!\langle\alpha\rangle\!\rangle\top \wedge \langle\varphi\rangle\top)$, and $P \models \mathcal{S}(A)$ but not $Q \models \mathcal{S}(A)$.

Ideally, we would want an alternative transformation onto weak formulas, but it seems difficult to formulate this in a succinct way. Through expressive completeness (Section 7.3) we get the following:

Theorem 20. Let A be a weak formula over $\mathbf{T}$. Then there exists a weak formula A' over $\mathcal{S}(\mathbf{T})$ such that

$$P \models_{\mathbf{T}} A \quad \text{iff} \quad P \models_{\mathcal{S}(\mathbf{T})} A'$$

for all states P.

Proof. The proof has been formalised in Isabelle. Note that $\{Q \mid Q \models_{\mathbf{T}} A\}$ is finitely supported (with support included in $\operatorname{supp}(A)$), and moreover closed under $\mathrel{\dot{\approx}}_{\mathcal{S}(\mathbf{T})}$ by Theorems 11 and 19. It then follows from Theorem 13 that $A' = \bigvee_{Q \models_{\mathbf{T}} A} \operatorname{Char}(Q)$ has the desired properties. $\qquad\square$

Inherent in this proof is a construction of the weak formula A' from the weak formula A, but the construction depends implicitly on distinguishing formulas (cf. Theorem 2) and thus on the entire transition system.

We have failed to find a transformation of weak formulas that yields equivalent weak formulas over $\mathcal{S}(\mathbf{T})$ and is defined by induction over formulas, or, at the very least, is independent of the transition system. We have also failed to prove that such a transformation cannot exist, and thus have to leave this problem open. The following counterexamples shed light on the difficulties:

Counterexample. Define the transformation $\mathcal{S}$ on weak formulas by

$$\mathcal{S}(\langle\!\langle\tau\rangle\!\rangle(A \wedge \varphi)) = \langle\!\langle\varphi\rangle\!\rangle\mathcal{S}(A)$$

With this definition, the counterpart of Theorem 17 fails. A counterexample is $A = \neg\langle\!\langle\alpha\rangle\!\rangle\top$, $P \vdash_{\mathbf{T}} \varphi$ with $P \xrightarrow{\tau}_{\mathbf{T}} Q$ and $P \xrightarrow{\alpha}_{\mathbf{T}} Q$ for some $\alpha \neq \tau$, where Q has no outgoing transitions, cf. the diagrams below:
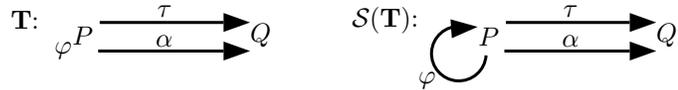
Since $P \overset{\varphi}{\Rightarrow}_{\mathcal{S}(\mathbf{T})} Q$ and Q has no $\langle\!\langle\alpha\rangle\!\rangle$ action, we have that

$$P \models_{\mathcal{S}(\mathbf{T})} \langle\!\langle\varphi\rangle\!\rangle\neg\langle\!\langle\alpha\rangle\!\rangle\top$$

The only state in $\mathbf{T}$ that satisfies φ also has an $\langle\!\langle\alpha\rangle\!\rangle$ action, thus it does *not* hold that

$$P \models_{\mathbf{T}} \langle\!\langle\tau\rangle\!\rangle(\neg\langle\!\langle\alpha\rangle\!\rangle\top \wedge \varphi)$$

Counterexample. Define the transformation $\mathcal{S}$ on weak formulas by

$$\mathcal{S}(\langle\!\langle\tau\rangle\!\rangle(A \wedge \varphi)) = \langle\!\langle\tau\rangle\!\rangle(\mathcal{S}(A) \wedge \langle\!\langle\varphi\rangle\!\rangle\top)$$

Then again the counterpart of Theorem 17 fails. A counterexample here is $A = \langle\!\langle\alpha\rangle\!\rangle\top$, with $P \xrightarrow{\tau}_{\mathbf{T}} Q$ and $Q \vdash_{\mathbf{T}} \varphi$, $P \xrightarrow{\alpha}_{\mathbf{T}} Q$ for some $\alpha \neq \tau$, cf. the diagrams below:
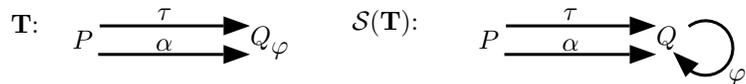
Here it holds that

$$P \models_{\mathcal{S}(\mathbf{T})} \langle\!\langle\tau\rangle\!\rangle(\langle\!\langle\alpha\rangle\!\rangle\top \wedge \langle\!\langle\varphi\rangle\!\rangle\top)$$

and it does *not* hold that

$$P \models_{\mathbf{T}} \langle\!\langle\tau\rangle\!\rangle(\langle\!\langle\alpha\rangle\!\rangle\top \wedge \varphi)$$

Finally, consider the *partial* transformation $\mathcal{S}'$ on weak formulas by

$$\mathcal{S}'(\langle\!\langle\tau\rangle\!\rangle((\langle\!\langle\tau\rangle\!\rangle A)\wedge\varphi)) = \langle\!\langle\varphi\rangle\!\rangle\mathcal{S}'(A)$$

where $\mathcal{S}'$ is homomorphic on the first three cases in Definition 25.

$\mathcal{S}'$ is not total since a formula $\langle\!\langle\tau\rangle\!\rangle(A\wedge\varphi)$ is in its domain only when $A = \langle\!\langle\tau\rangle\!\rangle A'$ for some A'. It is easy to see that $\mathcal{S}'$ is injective and surjective, i.e., every weak formula A on $\mathcal{S}(\mathbf{T})$ has a unique formula B on $\mathbf{T}$ such that $\mathcal{S}'(B) = A$. We write $\mathcal{S}'^{-1}$ for the inverse of $\mathcal{S}'$. Thus

$$\mathcal{S}'^{-1}(\langle\!\langle\varphi\rangle\!\rangle A) = \langle\!\langle\tau\rangle\!\rangle((\langle\!\langle\tau\rangle\!\rangle\mathcal{S}'^{-1}(A))\wedge\varphi)$$

and $\mathcal{S}'^{-1}$ is homomorphic on all other operators.

Theorem 21. $P \models_{\mathcal{S}(\mathbf{T})} A \quad$ iff $\quad P \models_{\mathbf{T}} \mathcal{S}'^{-1}(A)$

Proof. By induction over weak formulas on $\mathcal{S}(\mathbf{T})$. All cases are trivial by induction except for the case $\langle\!\langle\varphi\rangle\!\rangle A$ with $\varphi \in \mathrm{PRED}_\mathbf{T}$. Suppose $P \models_{\mathcal{S}(\mathbf{T})} \langle\!\langle\varphi\rangle\!\rangle A$. Then $P \Rightarrow_{\mathcal{S}(\mathbf{T})} P' \xrightarrow{\varphi}_{\mathcal{S}(\mathbf{T})} P'' \Rightarrow_{\mathcal{S}(\mathbf{T})} Q$ with $Q \models_{\mathcal{S}(\mathbf{T})} A$. By construction of $\mathcal{S}(\mathbf{T})$ we get $P' = P''$, and $P \Rightarrow_\mathbf{T} P' \Rightarrow_\mathbf{T} Q$ with $P' \vdash_\mathbf{T} \varphi$. By induction $Q \models_\mathbf{T} \mathcal{S}'^{-1}(A)$. This establishes that $P' \models_\mathbf{T} \langle\!\langle\tau\rangle\!\rangle\mathcal{S}'^{-1}(A)\wedge\varphi$ and thus $P \models_\mathbf{T} \langle\!\langle\tau\rangle\!\rangle(\langle\!\langle\tau\rangle\!\rangle\mathcal{S}'^{-1}(A)\wedge\varphi)$ as required. Conversely, if $P \models_\mathbf{T} \langle\!\langle\tau\rangle\!\rangle(\langle\!\langle\tau\rangle\!\rangle\mathcal{S}'^{-1}A\wedge\varphi)$ then $P \Rightarrow_\mathbf{T} P' \Rightarrow_\mathbf{T} Q$, and by construction of $\mathcal{S}(\mathbf{T})$ and induction we get $P \models_{\mathcal{S}(\mathbf{T})} \langle\!\langle\varphi\rangle\!\rangle A$. $\qquad\square$

An interesting consequence is that to express the distinguishing formulas guaranteed by Theorem 12, it is enough to consider formulas in $\mathrm{dom}(\mathcal{S}')$, i.e., in the last clause of Definition 25, it is enough to consider $A = \langle\!\langle\tau\rangle\!\rangle A'$. The reason is that if $P \not\approx_\mathbf{T} Q$ then by Theorem 19 also $P \not\approx_{\mathcal{S}(\mathbf{T})} Q$, which by Theorem 12 means there is a distinguishing formula B for P and Q in $\mathcal{S}(\mathbf{T})$, which by Theorem 21 means that $\mathcal{S}'^{-1}(B)$ is a distinguishing formula in $\mathbf{T}$.

In conclusion, the results in this section show that the transformation $\mathcal{S}$ indicates an alternative weak logic with formulas $\langle\!\langle\tau\rangle\!\rangle(A \wedge \langle\varphi\rangle\top)$, appropriate if φ only occurs in self-loops. The transformation $\mathcal{S}'$, on the other hand, indicates that a sublogic with formulas $\langle\!\langle\tau\rangle\!\rangle(\langle\!\langle\tau\rangle\!\rangle A\wedge\varphi)$ is expressively equivalent. It admits a smooth transformation of predicates into actions at the cost of a more complicated definition of a weak logic.

Through expressive completeness we obtain a transformation on our original weak logic, but it does not really shed light on the nature of state predicates and actions. Certainly, it means that any impossibility result for a transformation onto weak formulas must be qualified with a notion of independence from the transition system. This would require a proper definition of the set of transition systems that act as models for a given logic. One difficulty here is that the disjoint union of a set of transition systems is itself a transition system. By expressive completeness we get a transformation (dependent on this union); that same transformation thus applies to all the members in the set. Consider the set of all transition systems of cardinality at most κ; this set may have a cardinality higher than κ. Thus an impossibility result may actually depend on the cardinality limit in Definition 3.

9. APPLICATIONS

In this section we consider standard process calculi and their accompanying labelled bisimilarities, and investigate how to obtain an adequate modal logic using our framework.

In each of the final two examples, no HML has to our knowledge yet been proposed, and we immediately obtain one by instantiating the logic in the present paper.

Pi-calculus. The pi-calculus by Milner et al. [MPW92] has a labelled transition system with early input, defined using a structural operational semantics. This transition system satisfies the axioms of a nominal transition system, and the logic of Section 3 (with an empty set of state predicates) is adequate for early bisimilarity. Other variants of bisimilarity can be obtained by instead using the original transition system (with late input) and one of the variants of our logic described in Section 6.

The pi-calculus already has several notions of weak bisimulation, and Definition 23 corresponds to the early weak bisimulation. In the pi-calculus there are no state predicates, thus the weak static implication is unimportant. The weak logic of Definition 25 is adequate for early weak bisimulation.

Applied pi-calculus. The applied pi-calculus by Abadi and Fournet (2001) [AF01] comes equipped with a labelled transition system and a notion of weak labelled bisimulation. States contain a record of emitted messages; this record has a domain and can be used to equate open terms M and N modulo some rewrite system. The definition of bisimulation requires bisimilar processes to have the same domain and equate the same open terms, i.e., to be strongly statically equivalent. We model these requirements using state predicates "$x \in \text{dom}$" and "$M \equiv N$". Since satisfaction is invariant under silent transitions, weak and strong static implication coincide, and our weak HML is adequate for Abadi and Fournet's early weak labelled bisimilarity.

Spi-calculus. The spi-calculus by Abadi and Gordon (1999) [AG99] has a formulation as an environment-sensitive labelled transition system by Boreale et al. (2001) [BDNP01] equipped with state formulae (predicates) ϕ. Adding state predicates "$x \in \text{dom}$" to the state predicates makes our weak HML adequate with respect to Boreale's early weak bisimilarity.

Concurrent constraint pi calculus. The concurrent constraint pi calculus (CC-pi) by Buscemi and Montanari (2007) [BM07] extends the explicit fusion calculus [WG05] with a more general notion of constraint stores c. Using the labelled transition system of CC-pi and the associated bisimulation (Definition 2), we immediately get an adequate modal logic.

The reference equivalence for CC-pi is open bisimulation [BM08] (closely corresponding to hyperbisimulation in the fusion calculus [PV98]), which differs from labelled bisimulation in two ways: First, two equivalent processes must be equivalent under all store extensions. To encode this, we let the effects $\mathcal{F}$ be the set of constraint stores c different from 0, and let $c(P) = c \mid P$. Second, when simulating a labelled transition $P \xrightarrow{\alpha} P'$, the simulating process Q can use any transition $Q \xrightarrow{\beta} Q'$ with an equivalent label, as given by a state predicate $\alpha = \beta$. As an example, if $\alpha = \bar{a}\langle x \rangle$ is a free output label then $P \vdash \alpha = \beta$ iff $\beta = \bar{b}\langle y \rangle$ where $P \vdash a = b$ and $P \vdash x = y$. To encode this, we transform the labels of the transition system by replacing them with their equivalence classes, i.e., $P \xrightarrow{\alpha} P'$ becomes $P \xrightarrow{[\alpha]_P} P'$ where $\beta \in [\alpha]_P$ iff $P \vdash \beta = \alpha$. Hyperbisimilarity (Definition 13) on this transition system then corresponds to open bisimilarity, and the modal logic defined in Section 6 is adequate.

Psi-calculi. In psi-calculi by Bengtson et al. (2011) [BJPV11], the labelled transitions take the form $\Psi \rhd P \xrightarrow{\alpha} P'$, where the assertion environment Ψ is unchanged after the step.

We model this as a nominal transition system by letting the set of states be pairs (Ψ, P) of assertion environments and processes, and define the transition relation by $(\Psi, P) \xrightarrow{\alpha} (\Psi, P')$ if $\Psi \rhd P \xrightarrow{\alpha} P'$. The notion of bisimulation used with psi-calculi also uses an assertion environment and is required to be closed under environment extension, i.e., if $\Psi \rhd P \sim Q$, then $\Psi \otimes \Psi' \rhd P \sim Q$ for all Ψ'. We let the effects $\mathcal{F}$ be the set of assertions, and define $\Psi((\Psi', P)) = (\Psi \otimes \Psi', P)$. Hyperbisimilarity on this transition system then subsumes the standard psi-calculi bisimilarity, and the modal logic defined in Section 6 is adequate.

Weak bisimilarity for CC-pi and psi. Both CC-pi and psi-calculi has a special unobservable action τ, but until now only psi-calculi have a notion of weak labelled bisimulation (as remarked in Section 7.1), and neither has a weak HML. Through this paper they both gain both weak bisimulation and logic, although more work is needed to establish how compatible the bisimulation equivalence is with their respective syntactic constructs. A complication is that the natural formulation of bisimulation makes use of effects (store or assertion extensions) which are bisimulation requirements on neither predicates nor actions. In order to map them into our framework these would need to be cast as actions. Their interactions with the silent action could be an interesting topic for further research.

10. Related work

We here discuss other modal logics for process calculi, with a focus on how their constructors can be captured by finitely supported conjunction in our HML. This comparison is by necessity somewhat informal: formal correspondence fails to hold due to differences in the conjunction operator of the logic (finite, uniformly bounded or unbounded vs. bounded support).

HML for CCS. The first published HML is Hennessy and Milner (1980–1985) [HM80, Mil81, HM85]. They work with image-finite CCS processes, where finite (binary) conjunction suffices for adequacy, and define both strong and weak versions of the logic. The logic by Hennessy and Liu (1995) [HL95] for a value-passing calculus also uses binary conjunction, where image-finiteness is due to a late semantics and the logic contains quantification over data values. A similar idea and argument is in a logic for LOTOS by Calder et al. (2002) [CMS02], though that only considers stratified bisimilarity up to ω.

Hennessy and Liu's value-passing calculus is based on ordinary CCS. In this calculus, a receiving process $a(x).P$ can participate in a synchronisation on a, becoming an abstraction $(x)P$ where v is a bound variable. Dually, a sending process $\bar{a}\,v.Q$ becomes a concretion (v, Q) where v is a value. The abstraction and concretion above react as part of the synchronisation on a, yielding $P\{v/x\} \mid Q$. To capture the operations of abstractions and concretions in our framework, we add effects $\mathtt{id}$ and $?v$, with $?v((x)P) = P\{v/x\}$, and transitions $(v, P) \xrightarrow{!v} P$. Letting $L(a?, \_) = \{?v \mid v \in \text{values}\}$ and $L(\alpha, \_) = \{\mathtt{id}\}$ otherwise, late bisimilarity is $\{\mathtt{id}\}/L$-bisimilarity as defined in Section 6. We can then encode their universal quantifier $\forall x.A$ as $\bigwedge_v \langle ?v \rangle A\{v/x\}$, which has support $\text{supp}(A) \setminus \{x\}$, and their output modality $\langle c!x \rangle A$ as $\langle c! \rangle \bigvee_v \langle !v \rangle A\{v/x\}$, with support $\{c\} \cup (\text{supp}(A) \setminus \{x\})$.

An infinitary HML for CCS is discussed in Milner's book (1989) [Mil89], where also the process syntax contains infinite summation. There are no restrictions on the indexing sets and no discussion about how this can exhaust all names. The adequacy theorem is proved by stratifying bisimilarity and using transfinite induction over all ordinals, where

the successor step basically is the contraposition of the argument in Theorem 2, though without any consideration of finite support. A more rigorous treatment of the same ideas is by Abramsky (1991) [Abr91] where uniformly bounded conjunction is used throughout.

Koutavas et al. (2018) [KGH18] study a transactional CCS, extending the natural transition system in order to provide three equivalent adequate HMLs. Neither their processes nor their logics have a notion of α-renaming. The conjunction operator used in their logics is infinite, without an explicit bound on its cardinality. Without such a bound the set of formulas is not well-defined: let $\mathcal{F}$ be the set of all formulas, and consider the subset of formulas $\mathcal{S} := \left\{ \bigwedge_{A \in I} A \mid I \subseteq \mathcal{F} \right\}$. By Cantor's Theorem $|\mathcal{S}| > |\mathcal{F}|$, which is a contradiction.

Simpson (2004) [Sim04] gives a sequent calculus for proving HML properties of process calculi in GSOS format, using binary conjunction and assuming finite branching and a finite number of actions.

μ-calculus. Kozen's modal μ-calculus (1983) [Koz83] subsumes several other weak temporal logics including CTL* (Cranen et al. 2011) [CGR11], and can encode weak transitions using least fixed points. Dam (1996) [Dam96] gives a modal μ-calculus for the pi-calculus, treating bound names using abstractions and concretions, and provides a model checking algorithm. Bradford and Stevens (1999) [BS99] give a generic framework for parameterising the μ-calculus on data environments, state predicates, and action expressions. The logic defined in the present paper can encode the least fixpoint operator of μ-calculi by a disjunction of its finite unrollings, as seen in Section 5.2. We can immediately encode the atomic μ-calculus by Klin and Łełyk (2017) [KL17].

Pi-calculus. The first HML for the pi-calculus is by Milner et al. (1993) [MPW93], where infinite conjunction is used in the early semantics and conjunctions are restricted to use a finite set of free names. The adequacy proof is of the same structure as in this paper. The logic of Section 3, applied to the pi-calculus transition system from which bound input actions $x(y)$ have been removed, contains the logic $\mathcal{F}$ of Milner et al., or the equipotent logic $\mathcal{FM}$ if we take the set of name matchings $[a = b]$ as state predicates.

Koutavas and Hennessy (2012) [KH12] give a weak HML for a higher-order pi-calculus with both higher-order and first-order communication using an environment-sensitive LTS. Xu and Long (2015) [XL15] define a weak HML with countable conjunction for a purely higher-order pi-calculus. The adequacy proof uses stratification.

Ahn et al. (2017)[AHT17] give an innovative intuitionistic logic characterising open bisimilarity. In their setting, substitution effects are only employed inside the definition of the logical implication operator. Their soundness proof is standard; for the completeness proof the distinguishing formula for P wrt. Q is constructed in parallel with one for Q wrt. P.

There are several extensions of HML with spatial modalities. The one most closely related to our logic is by Berger et al. (2008) [BHY08]. They define an HML with both strong and weak action modalities, fixpoints, spatial conjunction and adjunction, and a scope extrusion modality, to study a typed value-passing pi-calculus with selection and recursion. The logic has three (may, must, and mixed) proof systems that are sound and relatively complete.

Spi Calculus. Frendrup et al. (2002) [FHNJ02] provide three Hennessy-Milner logics for the spi calculus [AG99]. All three logics use infinite quantification without any consideration of finite support. The transition system used is a variant of the one by Boreale et al. (2001) [BDNP01], where a state is a pair $\sigma \rhd P$ of a process P and its environment σ: a substitution that maps environment variables to public names and messages received from the process. This version of the spi calculus has expressions ξ, that are terms constructed from names and environment variables using encryption and decryption operators, and messages M, that only contain names and encryption. Substitution $\xi\sigma$ replaces environment variables in ξ with their values in σ, and evaluation $\mathbf{e}(\xi)$ is a partial function that attempts to perform the decryptions in ξ, yielding a message M if all decryptions are successful.

As usual for the spi calculus, the bisimulation (and logic) is defined in terms of the environment actions, rather than the process actions. In Frendrup's version of Boreale's environment-sensitive transition system, the transition labels are related to the process actions in the following way: when a process P receives message M on channel a, the label of a corresponding environment-sensitive transition $\sigma \rhd P \xrightarrow{a\,\xi} \sigma' \rhd P'$ describes how the environment σ computed the message $M = \mathbf{e}(\xi\sigma)$. For process output of message M on channel a, the corresponding environment-sensitive transition is simply $\sigma \rhd P \xrightarrow{\bar{a}} \sigma' \rhd P'$; the message M can be recovered from the updated environment σ'.

The logics of Frendrup et al. include a matching modality $[M = N]A$ that is defined using implication: $\sigma \rhd P \models [M = N]A$ iff $\mathbf{e}(M\sigma) = \mathbf{e}(N\sigma)$ implies $\sigma \rhd P \models A$. This is equipotent to having matching as a state predicate, since we can rewrite all non-trivial guards by $[M = N]A \iff A \vee \neg[M = N]\top$.

The logic of Section 3, applied to the nominal transition system with the environment labels τ, $\bar{a}$ and $a\,\xi$ above has the same modalities as the logic $\mathcal{F}$ of Frendrup et al. The logic $\mathcal{EM}$ by Frendrup et al. replaces the simple input modality by an early input modality $\langle \underline{a}\,(x) \rangle^E A$, which (after a minor manipulation of the input labels) can be encoded as the conjunction $\bigwedge_\xi \langle a\,\xi \rangle A\{\xi/x\}$ with support $\mathrm{supp}(A) \setminus \{x\}$. We do not consider their logic $\mathcal{LM}$ that uses a late input modality, since its application relies on sets that do not have finite support [FHNJ02, Theorem 6.12], which are not meaningful in nominal logic.

Frendrup et al. claim to "characterize early and late versions of the environment sensitive bisimilarity of [BDNP01]", but this claim only holds with some modification. First the definition of static equivalence [FNJ01, Definition 22] that is used in the adequacy proofs is strictly stronger than the one that appears in the published paper [FHNJ02, Definition 3.4]. Thus, the adequacy results [FHNJ02, Theorems 6.3, 6.4, and 6.14] are false as stated, but can be repaired by substituting the stronger notion of static equivalence. Then Frendrup's logics and bisimilarities become sound, but not complete, with respect to the bisimilarity of [BDNP01], since the latter uses the weaker notion of static equivalence (Definition 3.4). In contrast, we sketch a simple instance of our logic, for Boreale's labelled transition system, that is adequate for early bisimilarity in Section 9.

Applied Pi-calculus. A more recent work is a weak HML by Hüttel and Pedersen (2007) [HP07] for the applied pi-calculus by Abadi and Fournet (2001) [AF01], where completeness relies on an assumption of image-finiteness of the weak transitions. Similarly to the spi calculus, there is a requirement that terms M received by a process P can be computed from the current knowledge available to an observer of the process, which we here write $M \in \mathcal{S}(P)$.

The logic contains atomic formulae for term equality (indistinguishability) in the frame of a process, corresponding to our state predicates. However, Hüttel and Pedersen use a notion of equality (and thus static equivalence) that is stronger than the corresponding relation by Abadi and Fournet, and that is not well-defined modulo alpha-renaming [Ped06, p. 20]. Since our framework is based on nominal sets, we must identify alpha-equivalent processes, and instead use Abadi and Fournet's notion of term equality.

Hüttel and Pedersen's logic includes an early input modality and an existential quantifier. The early input modality $\langle \underline{a}\,(x) \rangle A$ can be straightforwardly encoded as the conjunction $\bigwedge_M \langle \underline{a}\, M \rangle A \{^M/_x\}$, with support $\{a\} \cup (\mathrm{supp}(A) \setminus \{x\})$. The definition of the existential quantifier takes the observer knowledge into account: P satisfies $\exists x.A$ if $x \# P$ and there is $M \in \mathcal{S}(P)$ such that $\{^M/_x\} \mid P$ satisfies A. The condition $M \in \mathcal{S}(P)$ makes the quantifier difficult to encode using effects, since there is no corresponding state predicate (for good reason: the main property modelled by cryptographic process calculi is that different cipher texts $\mathrm{E}(M,k)$ and $\mathrm{E}(N,k)$ are indistinguishable unless the key k is known). To treat the existential quantifier, we instead add an action (x) with $\mathrm{bn}((x)) = x$ and transitions $P \xrightarrow{(x)} \{^M/_x\} \mid P$ if $M \in \mathcal{S}(P)$ and $x \# P$. We can then encode $\exists x.A$ as $\langle (x) \rangle A$.

Fusion calculus. In an HML for the fusion calculus by Haugstad et al. (2006) [HTV06] the fusions (i.e., equality relations on names) are action labels φ. The corresponding modal operator $\langle \varphi \rangle A$ has the semantics that the formula A must be satisfied for all substitutive effects of φ (intuitively, substitutions that map each name to a fixed representative for its equivalence class). In order to represent fusion actions in the logics in this paper, we add substitution effects σ such that $\sigma(P') = P'\sigma$. The fusion modality $\langle \varphi \rangle A$ can then be encoded in our framework as $\langle \varphi \rangle \bigvee_\sigma \sigma @ A\sigma$, where the parameter σ of the disjunction ranges over the (finite set of) substitutive effects of φ. Their adequacy theorem uses the contradiction argument with infinite conjunction, with no argument about finiteness of names for the distinguishing formula.

Nominal transition systems. De Nicola and Loreti (2008) [DL08] define a general format for multiple-labelled transition systems with labels for name revelation and resource management, and an associated modal logic with name equality predicates, name quantification ($\exists$ and $И$), and a fixed-point modality. In contrast, we seek a small and expressive HML for general nominal transition systems. Indeed, the logic of De Nicola and Loreto can be seen as a special case of ours: their different transition systems can be merged into a single one, and we can encode their quantifiers and fixpoint operator as described in Section 4. Nominal SOS of Cimini et al. (2012) [CMRG12] is also a special case of our nominal transition systems. Aceto et al. (2017) [AFGP$^+$17] give conditions on rule formats for transitions from processes to residuals so that they generate nominal transition systems.

11. Conclusion

We have given a general account of transition systems and Hennessy-Milner Logic using nominal sets. The advantage of our approach is that it is more expressive than previous work. We allow infinite conjunctions that are not uniformly bounded, meaning that we can encode, e.g., quantifiers and the next-step operator. We have given ample examples of how the definition captures different variants of bisimilarity and how it relates to many different versions of HML in the literature.

We have formalised results of Sections 3 and 6–8. For this we use Nominal Isabelle [UK12], an implementation of nominal logic in Isabelle/HOL [NPW02]. This is a popular interactive proof assistant for higher-order logic with convenient specification mechanisms for, and automation to reason about, datatypes with binders. Our Isabelle theories are available from the Archive of Formal Proofs [WEP+16]. We here comment on some of the interesting aspects, and the difficulties involved for the results yet without a formal proof.

The main challenge, and perhaps most prominent contribution from the perspective of formalisation, is the definition of formulas (Definitions 3 and 17). Nominal Isabelle does not directly support infinitely branching datatypes. We therefore construct formulas from first principles in higher-order logic, by defining an inductive datatype of *raw* formulas (where alpha-equivalent raw formulas are *not* identified). The datatype constructor for conjunction recurses through sets of raw formulas of bounded cardinality, a feature made possible only by a recent re-implementation of Isabelle/HOL's datatype package [BHL+14].

We then define alpha-equivalence of raw formulas. For finitely branching datatypes, alpha-equivalence is based on a notion of free names. Here, the correct notion of free names in a raw formula is given by the support of the raw formula's alpha-equivalence class. To obtain formulas, we quotient raw formulas by alpha-equivalence, and finally carve out the subtype of all terms that can be constructed from finitely supported ones. We then prove important lemmas; for instance, a strong induction principle for formulas that allows the bound names in $\langle\alpha\rangle A$ to be chosen fresh for any finitely supported context. To our knowledge, this is the first mechanisation of infinitely branching nominal datatypes in a proof assistant. This also explains why we prefer the logic for weak bisimulation to be a sublogic (Definition 25): that way we do not have to re-do these results.

Armed with these definitions and induction principles, most of the proof mechanisation proceeded smoothly. The formal proofs follow closely the sketches provided in this paper, and often helped clarify points and correct minor errors. Our development currently amounts to approximately 8800 lines of Isar proof scripts. This is divided according to the following table:

Topic	Section	LOC
Basics	3	3062
L-bisimilarity	6	3215
Unobservables	7	1730
Predicates vs actions	8	817

Some parts of this paper lack formal proofs. Among them are the definitions and results in Section 5 on fixpoint operators. Here proof mechanisation would be very tedious, since Isabelle does not provide enough support for transfinite induction. The results of Sections 9 and 10, on applications and comparisons with related work, lack proofs for two reasons. First, the correspondences are not exact since we allow arbitrary finitely supported conjunctions, and second, any formal proof must begin by constructing a formal model of the related work under study, which would require a large effort including resolving any ambiguities in that work.

Finally, the result on disjunction elimination (Section 7.4, Theorem 14) is challenging for a different reason. The proof uses a case analysis and induction over derived operators. While we believe that our proof is correct, to be absolutely certain we need to prove that the ranges of the derived operators are disjoint, and derive induction principles. This turned

out to be much more difficult than first anticipated, and at the time of writing, we have not finished the proofs.

Acknowledgements

References

[Abr91]   Samson Abramsky. A domain equation for bisimulation. *Journal of Information and Computation*, 92(2):161–218, 1991.

[AF01]   Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *Proceedings of POPL '01*, pages 104–115. ACM, January 2001.

[AFGP+17] Luca Aceto, Ignacio Fábregas, Álvaro García-Pérez, Anna Ingólfsdóttir, and Yolanda Ortega-Mallén. Rule Formats for Nominal Process Calculi. In Roland Meyer and Uwe Nestmann, editors, *28th International Conference on Concurrency Theory (CONCUR 2017)*, volume 85 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:16, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[AG99]   Martín Abadi and Andrew D. Gordon. A calculus for cryptographic protocols: The Spi calculus. *Information and Computation*, 148(1):1–70, 1999.

[AHT17]   Ki Yung Ahn, Ross Horne, and Alwen Tiu. A Characterisation of Open Bisimilarity using an Intuitionistic Modal Logic. In Roland Meyer and Uwe Nestmann, editors, *28th International Conference on Concurrency Theory (CONCUR 2017)*, volume 85 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:17, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[BDNP01]   Michele Boreale, Rocco De Nicola, and Rosario Pugliese. Proof techniques for cryptographic processes. *SIAM Journal on Computing*, 31(3):947–986, 2001.

[BHL+14]   Jasmin Christian Blanchette, Johannes Hölzl, Andreas Lochbihler, Lorenz Panny, Andrei Popescu, and Dmitriy Traytel. Truly modular (co)datatypes for Isabelle/HOL. In Gerwin Klein and Ruben Gamboa, editors, *Proceedings of ITP 2014*, volume 8558 of *LNCS*, pages 93–110. Springer, 2014.

[BHY08]   Martin Berger, Kohei Honda, and Nobuko Yoshida. Completeness and logical full abstraction in modal logics for typed mobile processes. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming: 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, part II*, volume 5126 of *LNCS*, pages 99–111. Springer Berlin Heidelberg, 2008.

[BJPV11]   Jesper Bengtson, Magnus Johansson, Joachim Parrow, and Björn Victor. Psi-calculi: a framework for mobile processes with nominal data and logic. *Logical Methods in Computer Science*, 7(1), 2011.

[BM07]   Maria Grazia Buscemi and Ugo Montanari. CC-Pi: A constraint-based language for specifying service level agreements. In Rocco De Nicola, editor, *Proceedings of ESOP 2007*, volume 4421 of *LNCS*, pages 18–32. Springer, 2007.

[BM08]   Maria Grazia Buscemi and Ugo Montanari. Open bisimulation for the concurrent constraint pi-calculus. In Sophia Drossopoulou, editor, *Proceedings of ESOP 2008*, volume 4960 of *LNCS*, pages 254–268. Springer, 2008.

[BS99]   Julian C. Bradfield and Perdita Stevens. Observational mu-calculus. Technical Report RS-99-5, BRICS, 1999.

[CE82]   Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. In Dexter Kozen, editor, *Logic of Programs, Workshop*, volume 131 of *LNCS*, pages 52–71. Springer, 1982.

[CGR11]   Sjoerd Cranen, Jan Friso Groote, and Michel Reniers. A linear translation from CTL* to the first-order modal μ-calculus. *Theoretical Computer Science*, 412(28):3129 – 3139, 2011. Festschrift in Honour of Jan Bergstra.

[CMRG12] Matteo Cimini, Mohammad Reza Mousavi, Michel A. Reniers, and Murdoch J. Gabbay. Nominal SOS. *Electronic Notes in Theoretical Computer Science*, 286:103–116, September 2012. Proceedings of the 28th Conference on the Mathematical Foundations of Programming Semantics (MFPS XXVIII).

[CMS02]   Muffy Calder, Savi Maharaj, and Carron Shankland. A modal logic for full LOTOS based on symbolic transition systems. *The Computer Journal*, 45(1):55–61, 2002.

[Dam96]   Mads Dam. Model checking mobile processes. *Information and Computation*, 129(1):35 – 51, 1996.

[DL08]   Rocco De Nicola and Michele Loreti. Multiple-labelled transition systems for nominal calculi and their logics. *Mathematical Structures in Computer Science*, 18(1):107–143, 2008.

[Eme96]   E. Allen Emerson. Model checking and the Mu-calculus. In Neil Immerman and Phokion G. Kolaitis, editors, *Descriptive Complexity and Finite Models, Proceedings of a DIMACS Workshop*, volume 31 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 185–214. American Mathematical Society, 1996.

[FHNJ02] Ulrik Frendrup, Hans Hüttel, and Jesper Nyholm Jensen. Modal logics for cryptographic processes. *Electronic Notes in Theoretical Computer Science*, 68(2):124–141, 2002. Proceedings of EXPRESS'02, 9th International Workshop on Expressiveness in Concurrency.

[FNJ01]   Ulrik Frendrup and Jesper Nyholm Jensen. Bisimilarity in the spi-calculus. Master's thesis, Aalborg University, 2001. http://projekter.aau.dk/projekter/en/studentthesis/bisimilarity-in-the-spicalculus(f1636334-8677-4fdf-aaa8-66f9d8e0fbf7).html.

[HL95]   Matthew Hennessy and Xinxin Liu. A modal logic for message passing processes. *Acta Informatica*, 32(4):375–393, 1995.

[HM80]   Matthew Hennessy and Robin Milner. On observing nondeterminism and concurrency. In J. W. de Bakker and Jan van Leeuwen, editors, *Automata, Languages and Programming, 7th Colloquium*, volume 85 of *Lecture Notes in Computer Science*, pages 299–309. Springer, 1980.

[HM85]   Matthew Hennessy and Robin Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM*, 32(1):137–161, 1985.

[HP07]   Hans Hüttel and Michael D. Pedersen. A logical characterisation of static equivalence. *Electronic Notes in Theoretical Computer Science*, 173:139–157, 2007. Proceedings of the 23rd Conference on the Mathematical Foundations of Programming Semantics (MFPS XXIII).

[HTV06]   Arild Martin Møller Haugstad, Anders Franz Terkelsen, and Thomas Vindum. A modal logic for the fusion calculus. Unpublished, Aalborg University, http://vbn.aau.dk/ws/files/61067487/1149104946.pdf, 2006.

[JBPV10] Magnus Johansson, Jesper Bengtson, Joachim Parrow, and Björn Victor. Weak equivalences in psi-calculi. In *Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11-14 July 2010, Edinburgh, United Kingdom*, pages 322–331, 2010.

[JKP90]   Bengt Jonsson, Ahmed Hussain Khan, and Joachim Parrow. Implementing a model checking algorithm by adapting existing automated tools. In Joseph Sifakis, editor, *Automatic Verification Methods for Finite State Systems: International Workshop, Grenoble, France June 12–14, 1989 Proceedings*, volume 407 of *LNCS*, pages 179–188. Springer Berlin Heidelberg, Berlin, Heidelberg, 1990.

[KGH18]   Vasileios Koutavas, Maciej Gazda, and Matthew Hennessy. Distinguishing between communicating transactions. *Information and Computation*, 259:1 – 30, 2018.

[KH12]   Vasileios Koutavas and Matthew Hennessy. First-order reasoning for higher-order concurrency. *Computer Languages, Systems & Structures*, 38(3):242–277, 2012.

[KL17]   Bartek Klin and Mateusz Lelyk. Modal mu-Calculus with Atoms. In Valentin Goranko and Mads Dam, editors, *26th EACSL Annual Conference on Computer Science Logic (CSL 2017)*, volume 82 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30:1–30:21, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[Koz83]   Dexter Kozen. Results on the propositional mu-calculus. *Theoretical Computer Science*, 27(3):333 – 354, 1983. Special Issue Ninth International Colloquium on Automata, Languages and Programming (ICALP) Aarhus, Summer 1982.

[Kur22]   Casimir Kuratowski. Une méthode d'élimination des nombres transfinis des raisonnements mathématiques. *Fundamenta Mathematicae*, 3(1):76–108, 1922.

[Lam02]   Leslie Lamport. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, June 2002.

[Lar88]   Kim G. Larsen. Proof systems for Hennessy-Milner logic with recursion. In M. Dauchet and M. Nivat, editors, *Proceedings of CAAP '88*, volume 299 of *LNCS*, pages 215–230. Springer, 1988.

[Mil81]   Robin Milner. A modal characterisation of observable machine-behaviour. In *Proceedings of CAAP 1981*, volume 112 of *LNCS*, pages 25–34, 1981.

[Mil89]   Robin Milner. *Communication and Concurrency*. Prentice Hall, 1989.

[MPW92]   Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I. *Information and Computation*, 100(1):1–40, 1992.

[MPW93]   Robin Milner, Joachim Parrow, and David Walker. Modal logics for mobile processes. *Theoretical Computer Science*, 114(1):149 – 171, 1993.

[NPW02]   Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.

[PBE+15]   Joachim Parrow, Johannes Borgström, Lars-Henrik Eriksson, Ramunas Gutkovas, and Tjark Weber. Modal logics for nominal transition systems. In Luca Aceto and David de Frutos-Escrig, editors, *26th International Conference on Concurrency Theory, CONCUR 2015, Madrid, Spain, September 1.4, 2015*, volume 42 of *LIPIcs*, pages 198–211. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

[Ped06]   Michael Pedersen. Logics for the applied pi calculus. Master's thesis, Aalborg University, 2006. BRICS RS-06-19.

[Pit13]   Andrew M. Pitts. *Nominal Sets*. Cambridge University Press, 2013. Cambridge Books Online.

[Pit15]   Andy Pitts, 2015. Personal communication.

[PV98]   Joachim Parrow and Björn Victor. The fusion calculus: Expressiveness and symmetry in mobile processes. In *Proceedings of LICS 1998*, pages 176–185. IEEE Computer Society Press, 1998.

[PWBE17]   Joachim Parrow, Tjark Weber, Johannes Borgström, and Lars-Henrik Eriksson. Weak nominal modal logic. In *Formal Techniques for Distributed Objects, Components, and Systems - 37th IFIP WG 6.1 International Conference, FORTE 2017, Held as Part of the 12th International Federated Conference on Distributed Computing Techniques, DisCoTec 2017, Neuchâtel, Switzerland, June 19-22, 2017, Proceedings*, pages 179–193, 2017.

[San93]   Davide Sangiorgi. A theory of bisimulation for the π-calculus. In Eike Best, editor, *Proceedings of CONCUR '93*, volume 715 of *LNCS*, pages 127–142. Springer, 1993.

[Sim04]   Alex Simpson. Sequent calculi for process verification: Hennessy-Milner logic for an arbitrary GSOS. *The Journal of Logic and Algebraic Programming*, 60-61:287–322, 2004.

[Tar55]   Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.

[UK12]   Christian Urban and Cezary Kaliszyk. General bindings and alpha-equivalence in Nominal Isabelle. *Logical Methods in Computer Science*, 8(2), 2012.

[WEP+16]   Tjark Weber, Lars-Henrik Eriksson, Joachim Parrow, Johannes Borgström, and Ramunas Gutkovas. Modal logics for nominal transition systems. *Archive of Formal Proofs*, October 2016. http://isa-afp.org/entries/Modal_Logics_for_NTS.shtml, Formal proof development.

[WG05]   Lucian Wischik and Philippa Gardner. Explicit fusions. *Theoretical Computer Science*, 304(3):606–630, 2005.

[XL15]   Xian Xu and Huan Long. A logical characterization for linear higher-order processes. *Journal of Shanghai Jiaotong University (Science)*, 20(2):185–194, 2015.