

A Hybrid Linear Logic for Constrained Transition Systems

Joëlle Despeyroux¹ and Kaustuv Chaudhuri²

¹ INRIA and CNRS, I3S, Sophia-Antipolis, France

² INRIA, Saclay, France

Abstract

Linear implication can represent state transitions, but real transition systems operate under temporal, stochastic or probabilistic constraints that are not directly representable in ordinary linear logic. We propose a general modal extension of intuitionistic linear logic where logical truth is indexed by constraints and hybrid connectives combine constraint reasoning with logical reasoning. The logic has a focused cut-free sequent calculus that can be used to internalize the rules of particular constrained transition systems; we illustrate this with an adequate encoding of the synchronous stochastic pi-calculus.

1998 ACM Subject Classification F.4.1. Mathematical Logic: Modal and Temporal Logics; F.1.2. Modes of Computation: Parallelism and Concurrency

Keywords and phrases linear logic, hybrid logic, stochastic pi-calculus, focusing, adequacy

1 Introduction

To reason about state transition systems, we need a logic of state. Linear logic [18] is such a logic and has been successfully used to model such diverse systems as process calculi, references and concurrency in programming languages, security protocols, multi-set rewriting, and graph algorithms. Linear logic achieves this versatility by representing propositions as *resources* that are combined using \otimes , which can then be transformed using the linear implication (\multimap). However, linear implication is timeless: there is no way to correlate two concurrent transitions. If resources have lifetimes and state changes have temporal, probabilistic or stochastic *constraints*, then the logic will allow inferences that may not be realizable in the system being modelled. The need for formal reasoning in such constrained systems has led to the creation of specialized formalisms such as Computation Tree Logic (CTL)[14], Continuous Stochastic Logic (CSL) [2] or Probabilistic CTL (PCTL) [19]. These approaches pay a considerable encoding overhead for the states and transitions in exchange for the constraint reasoning not provided by linear logic. A prominent alternative to the logical approach is to use a suitably enriched process algebra such as the stochastic and probabilistic π -calculi or the κ -calculus [11]. Processes are animated by means of simulation and then compared with the observations. Process calculi do not however completely fill the need for *formal reasoning for constrained transition systems*.

We propose a simple yet general method to add constraint reasoning to linear logic. It is an old idea—*labelled deduction* [36] with *hybrid* connectives [5]—applied to a new domain. Precisely, we parameterize ordinary logical truth on a *constraint domain*: $A@w$ stands for the truth of A under constraint w . Only a basic monoidal structure is assumed about the constraints from a proof-theoretic standpoint. We then use the hybrid connectives of *satisfaction* (at) and *localization* (\Downarrow) to perform generic symbolic reasoning on the constraints at the propositional level. We call the result *hybrid linear logic* (HyLL); it has a generic cut-free (but cut admitting) sequent calculus that can be strengthened with a focusing restriction [1] to obtain a normal form for proofs. Any instance of HyLL that gives a semantic interpretation to the constraints enjoys these proof-theoretic properties.



© Joëlle Despeyroux and Kaustuv Chaudhuri;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Focusing allows us to treat HyLL as a *logical framework* for constrained transition systems. Logical frameworks with hybrid connectives have been considered before; hybrid *LF* (*HLF*), for example, is a generic mechanism to add many different kinds of resource-awareness, including linearity, to ordinary *LF* [32]. *HLF* follows the usual *LF* methodology of keeping the logic of the framework minimal: its proof objects are β -normal η -long natural deduction terms, but the equational theory of such terms is sensitive to permutative equivalences [37]. With a focused sequent calculus, we have more direct access to a canonical representation of proofs, so we can enrich the framework with any connectives that obey the focusing discipline. The representational adequacy of an encoding in terms of (partial) focused sequent derivations tends to be more straightforward than in a natural deduction formulation. We illustrate this by encoding the synchronous stochastic π -calculus ($S\pi$) in HyLL using rate functions as constraints.

In addition to the novel stochastic component, our encoding of $S\pi$ is a conceptual improvement over other encodings of π -calculi in linear logic. In particular, we perform a full propositional reflection of processes as in [23], but our encoding is first-order and adequate as in [6]. HyLL does not itself prescribe an operational semantics for the encoding of processes; thus, bisimilarity in continuous time Markov chains (*CTMC*) is not the same as logical equivalence in stochastic HyLL, unlike in *CSL* [12]. This is not a deficiency; rather, the *combination* of focused HyLL proofs and a proof search strategy tailored to a particular encoding is necessary to produce faithful symbolic executions. This exactly mirrors $S\pi$ where it is the simulation rather than the transitions in the process calculus that is shown to be faithful to the *CTMC* semantics [27].

This work has the following main contributions. First is the logic HyLL itself and its associated proof-theory, which has a very standard and well understood design. Second, we show how to obtain many different instances of HyLL for particular constraint domains because we only assume a basic monoidal structure for constraints. Third, we illustrate the use of focused sequent derivations to obtain adequate encodings by giving a novel adequate encoding of $S\pi$. Our encoding is, in fact, *fully adequate*, *i.e.*, partial focused proofs are in bijection with traces. The ability to encode $S\pi$ gives an indication of the versatility of HyLL.

This paper is organized as follows: in sec. 2, we present the inference (sequent calculus) rules for HyLL and describe the two main semantic instances: temporal and stochastic constraints. In sec. 3 we sketch the general focusing restriction on HyLL sequent proofs. In sec. 4 we give the encoding of $S\pi$ in stochastic HyLL, and show that the encoding is representationally adequate for focused proofs (theorems 18 and 20). We end with an overview of related (sec. 5) and future work (sec. 6). The full version of this paper is available as a technical report [9].

2 Hybrid Linear Logic

In this section we define HyLL, a conservative extension of intuitionistic first-order linear logic (*ILL*) [18] where the truth judgements are labelled by worlds representing constraints. Like in *ILL*, propositions are interpreted as *resources* which may be composed into a *state* using the usual linear connectives, and the linear implication (\multimap) denotes a transition between states. The world label of a judgement represents a constraint on states and state transitions; particular choices for the worlds produce particular instances of HyLL. The common component in all the instances of HyLL is the proof theory, which we fix once and for all. We impose the following minimal requirement on the kinds of constraints that HyLL can deal with.

► **Definition 1.** A *constraint domain* \mathcal{W} is a monoid structure $\langle W, \cdot, \iota \rangle$. The elements of W are called *worlds*, and the partial order $\leq : W \times W$ —defined as $u \leq w$ if there exists $v \in W$ such that $u \cdot v = w$ —is the *reachability relation* in \mathcal{W} .



The identity world ι is \leq -initial and is intended to represent the lack of any constraints. Thus, the ordinary ILL is embeddable into any instance of HyLL by setting all world labels to the identity. When needed to disambiguate, the instance of HyLL for the constraint domain \mathcal{W} will be written $\text{HyLL}(\mathcal{W})$. Two design choices are important to note. First, we only require the worlds to be monoids, not lattices, because this lets us give a sufficiently general system that it can be instantiated with rate functions as the constraint domain. Second, we do not assume that the monoid is commutative so that we can still choose to use lattices for the constraint domain.

Atomic propositions are written using minuscule letters (a, b, \dots) applied to a sequence of *terms* (s, t, \dots) , which are drawn from an untyped term language containing term variables (x, y, \dots) and function symbols (f, g, \dots) applied to a list of terms.. Non-atomic propositions are constructed from the connectives of first-order intuitionistic linear logic and the two hybrid connectives *satisfaction* (at) , which states that a proposition is true at a given world (w, u, v, \dots) , and *localization* (\Downarrow) , which binds a name for the (current) world the proposition is true at. The following grammar summarizes the syntax of HyLL propositions.

$$A, B, \dots ::= a \bar{r} \mid A \otimes B \mid \mathbf{1} \mid A \multimap B \mid A \& B \mid \top \mid A \oplus B \mid \mathbf{0} \mid !A \mid \forall x. A \mid \exists x. A \\ \mid (A \text{ at } w) \mid \Downarrow u. A \mid \forall u. A \mid \exists u. A$$

Note that in the propositions $\Downarrow u. A$, $\forall u. A$ and $\exists u. A$, the scope of the world variable u is all the worlds occurring in A . World variables cannot be used in terms, and neither can term variables occur in worlds; this restriction is important for the modular design of HyLL because it keeps purely logical truth separate from constraint truth. We let α range over variables of either kind. Note that the \Downarrow connective commutes with every propositional connective, including itself. That is, $\Downarrow u. (A * B)$ is equivalent to $(\Downarrow u. A) * (\Downarrow u. B)$ for all binary connectives $*$, and $\Downarrow u. * A$ is equivalent to $*(\Downarrow u. A)$ for every unary connective $*$, assuming the commutation will not cause an unsound capture of u . It is purely a matter of taste where to place the \Downarrow , and repetitions are harmless.

The unrestricted connectives $\wedge, \vee, \supset, \text{etc.}$ of intuitionistic (non-linear) logic can also be defined in terms of the linear connectives and the exponential $!$ using any of the available embeddings of intuitionistic logic into linear logic, such as Girard's embedding [18].

2.1 Sequent Calculus for HyLL

In this section, we give a sequent calculus presentation of HyLL and prove a cut-admissibility theorem. The sequent formulation in turn will lead to an analysis of the polarities of the connectives in order to get a focused sequent calculus that can be used to compile a logical theory into a system of derived inference rules with nice properties (sec. 3). For instance, if a given theory defines a transition system, then the derived rules of the focused calculus will exactly exhibit the same transitions. This is key to obtain the necessary representational adequacy theorems, as we shall see for the $S\pi$ -calculus example chosen in this paper (sec. 4.1).

We start with the judgements from linear logic [18] and enrich them with a modal situated truth. We present the syntax of hybrid linear logic in a sequent calculus style, using Martin-Löf's principle of separating judgements and logical connectives. Instead of the ordinary mathematical judgement "A is true", for a proposition A , judgements of HyLL are of the form "A is true at world w ", abbreviated as $A @ w$. We use sequents of the form $\Gamma ; \Delta \Longrightarrow C @ w$ where Γ and Δ are sets of judgements of the form $A @ w$, with Δ being moreover a *multiset*. Γ is called the *unrestricted context*: its hypotheses can be consumed any number of times. Δ is a *linear context*: every hypothesis in it must be consumed singly in the proof. Note that in a judgement $A @ w$ (as in a proposition $A \text{ at } w$), w can be any expression in \mathcal{W} , not only a variable. The notation $A[\tau/\alpha]$ stands for the replacement of all free occurrences of the variable α in A with the expression τ , avoiding capture. Note that the expressions in the rules are to be readen up to alpha-conversion.



2.1.0.1 Judgemental rules

$$\frac{}{\Gamma; a\vec{t}@u \Rightarrow a\vec{t}@u} \text{init} \quad \frac{\Gamma, A@u; \Delta, A@u \Rightarrow C@w}{\Gamma, A@u; \Delta \Rightarrow C@w} \text{copy}$$

2.1.0.2 Multiplicatives

$$\frac{\Gamma; \Delta \Rightarrow A@w \quad \Gamma; \Delta' \Rightarrow B@w}{\Gamma; \Delta, \Delta' \Rightarrow A \otimes B@w} \otimes R \quad \frac{\Gamma; \Delta, A@u, B@u \Rightarrow C@w}{\Gamma; \Delta, A \otimes B@u \Rightarrow C@w} \otimes L$$

$$\frac{}{\Gamma; \cdot \Rightarrow \mathbf{1}@w} \mathbf{1}R \quad \frac{\Gamma; \Delta \Rightarrow C@w}{\Gamma; \Delta, \mathbf{1}@u \Rightarrow C@w} \mathbf{1}L$$

$$\frac{\Gamma; \Delta, A@w \Rightarrow B@w}{\Gamma; \Delta \Rightarrow A \multimap B@w} \multimap R \quad \frac{\Gamma; \Delta \Rightarrow A@u \quad \Gamma; \Delta', B@u \Rightarrow C@w}{\Gamma; \Delta, \Delta', A \multimap B@u \Rightarrow C@w} \multimap L$$

2.1.0.3 Additives

$$\frac{}{\Gamma; \Delta \Rightarrow \top@w} \top R \quad \frac{}{\Gamma; \Delta, \mathbf{0}@u \Rightarrow C@w} \mathbf{0}L$$

$$\frac{\Gamma; \Delta \Rightarrow A@w \quad \Gamma; \Delta \Rightarrow B@w}{\Gamma; \Delta \Rightarrow A \& B@w} \&R \quad \frac{\Gamma; \Delta, A_i@u \Rightarrow C@w}{\Gamma; \Delta, \Delta', A_1 \& A_2@u \Rightarrow C@w} \&L_i$$

$$\frac{\Gamma; \Delta \Rightarrow A_i@w}{\Gamma; \Delta \Rightarrow A_1 \oplus A_2@w} \oplus R_i \quad \frac{\Gamma; \Delta, A@u \Rightarrow C@w \quad \Gamma; \Delta, B@u \Rightarrow C@w}{\Gamma; \Delta, A \oplus B@u \Rightarrow C@w} \oplus L$$

2.1.0.4 Quantifiers

$$\frac{\Gamma; \Delta \Rightarrow A@w}{\Gamma; \Delta \Rightarrow \forall \alpha. A@w} \forall R^\alpha \quad \frac{\Gamma; \Delta, [\tau/\alpha]A@u \Rightarrow C@w}{\Gamma; \Delta, \forall \alpha. A@u \Rightarrow C@w} \forall L$$

$$\frac{\Gamma; \Delta \Rightarrow [\tau/\alpha]A@w}{\Gamma; \Delta \Rightarrow \exists \alpha. A@w} \exists R \quad \frac{\Gamma; \Delta, A@u \Rightarrow C@w}{\Gamma; \Delta, \exists \alpha. A@u \Rightarrow C@w} \exists L^\alpha$$

For $\forall R^\alpha$ and $\exists L^\alpha$, α is assumed to be fresh with respect to the conclusion. For $\exists R$ and $\forall L$, τ stands for a term or world, as appropriate.

2.1.0.5 Exponentials

$$\frac{\Gamma; \cdot \Rightarrow A@w}{\Gamma; \cdot \Rightarrow !A@w} !R \quad \frac{\Gamma, A@u; \Delta \Rightarrow C@w}{\Gamma; \Delta, !A@u \Rightarrow C@w} !L$$

2.1.0.6 Hybrid connectives

$$\frac{\Gamma; \Delta \Rightarrow A@u}{\Gamma; \Delta \Rightarrow (A \text{ at } u)@v} \text{at}R \quad \frac{\Gamma; \Delta, A@u \Rightarrow C@w}{\Gamma; \Delta, (A \text{ at } u)@v \Rightarrow C@w} \text{at}L$$

$$\frac{\Gamma; \Delta \Rightarrow [w/u]A@w}{\Gamma; \Delta \Rightarrow \downarrow u. A@w} \downarrow R \quad \frac{\Gamma; \Delta, [v/u]A@v \Rightarrow C@w}{\Gamma; \Delta, \downarrow u. A@v \Rightarrow C@w} \downarrow L$$

■ **Figure 1** The sequent calculus for HyLL



The full collection of rules of the HyLL sequent calculus is in fig. 1. The rules for the linear connectives are borrowed from [8] where they are discussed at length, so we omit a more thorough discussion here. The rules for the first-order quantifiers are completely standard. A brief discussion of the hybrid rules follows. To introduce the *satisfaction* proposition ($A \text{ at } u$) (at any world v') on the right, the proposition A must be true in the world u . The proposition ($A \text{ at } u$) itself is then true at any world, not just in the world u . In other words, ($A \text{ at } u$) carries with it the world at which it is true. Therefore, suppose we know that ($A \text{ at } u$) is true (at any world v); then, we also know that $A @ u$. The other hybrid connective of *localisation*, \downarrow , is intended to be able to name the current world. That is, if $\downarrow u. A$ is true at world w , then the variable u stands for w in the body A . This interpretation is reflected in its introduction rule on the right $\downarrow R$. For left introduction, suppose we have a proof of $\downarrow u. A @ v$ for some world v . Then, we also know $[v/u]A @ v$.

There are only two structural rules: the init rule infers an atomic initial sequent, and the copy rule introduces a contracted copy of an unrestricted assumption into the linear context (reading from conclusion to premise). Weakening and contraction are admissible rules:

► **Theorem 2** (structural properties).

1. If $\Gamma ; \Delta \Longrightarrow C @ w$, then $\Gamma, \Gamma' ; \Delta \Longrightarrow C @ w$. (*weakening*)
2. If $\Gamma, A @ u, A @ u ; \Delta \Longrightarrow C @ w$, then $\Gamma, A @ u ; \Delta \Longrightarrow C @ w$. (*contraction*)

Proof. By straightforward structural induction on the given derivations. ◀

The most important structural properties are the admissibility of the identity and the cut principles. The identity theorem is the general case of the init rule and serves as a global syntactic completeness theorem for the logic. Dually, the cut theorem below establishes the syntactic soundness of the calculus; moreover there is no cut-free derivation of $\cdot ; \cdot \Longrightarrow \mathbf{0} @ w$, so the logic is also globally consistent.

► **Theorem 3** (identity). $\Gamma ; A @ w \Longrightarrow A @ w$.

Proof. By induction on the structure of A (see [9]). ◀

► **Theorem 4** (cut).

1. If $\Gamma ; \Delta \Longrightarrow A @ u$ and $\Gamma ; \Delta', A @ u \Longrightarrow C @ w$, then $\Gamma ; \Delta, \Delta' \Longrightarrow C @ w$.
2. If $\Gamma ; \cdot \Longrightarrow A @ u$ and $\Gamma, A @ u ; \Delta \Longrightarrow C @ w$, then $\Gamma ; \Delta \Longrightarrow C @ w$.

Proof. By lexicographic structural induction on the given derivations, with cuts of kind 2 additionally allowed to justify cuts of kind 1. The style of proof sometimes goes by the name of *structural cut-elimination* [8]. See [9] for the details. ◀

We can use the admissible cut rules to show that the following rules are invertible: $\otimes L$, $\mathbf{1} L$, $\oplus L$, $\mathbf{0} L$, $\exists L$, $\neg R$, $\& R$, $\top R$, and $\forall R$. In addition, the four hybrid rules, $\text{at} R$, $\text{at} L$, $\downarrow R$ and $\downarrow L$ are invertible. In fact, \downarrow and at commute freely with all non-hybrid connectives:

► **Theorem 5** (Invertibility). *The following rules are invertible:*

1. On the right: $\& R$, $\top R$, $\neg R$, $\forall R$, $\downarrow R$ and $\text{at} R$;
2. On the left: $\otimes L$, $\mathbf{1} L$, $\oplus L$, $\mathbf{0} L$, $\exists L$, $\downarrow L$ and $\text{at} L$.

► **Corollary 6** (consistency). *There is no proof of $\cdot ; \cdot \Longrightarrow \mathbf{0} @ w$.*

Proof. A straightforward consequence of thm. 4. ◀

HyLL is conservative with respect to ordinary intuitionistic logic: as long as no hybrid connectives are used, the proofs in HyLL are identical to those in ILL [8]. The proof (omitted) is by simple structural induction.



► **Theorem 7** (conservativity). *Call a proposition or multiset of propositions pure if it contains no instance of the hybrid connectives and no instance of quantification over a world variable, and let Γ , Δ and A be pure. Then, If $\Gamma ; \Delta \Longrightarrow_{\text{HyLL}} C@w$ is derivable, then so is $\Gamma ; \Delta \Longrightarrow_{\text{ILL}} C$.* ◀

In the rest of this paper we use the following derived connectives:

► **Definition 8** (modal connectives).

$$\begin{aligned} \Box A &\triangleq \downarrow u. \forall w. (A \text{ at } u \cdot w) & \Diamond A &\triangleq \downarrow u. \exists w. (A \text{ at } u \cdot w) \\ \rho_v A &\triangleq \downarrow u. (A \text{ at } u \cdot v) & \dagger A &\triangleq \forall u. (A \text{ at } u) \end{aligned}$$

The connective ρ represents a form of delay. Note its derived right rule:

$$\frac{\Gamma ; \Delta \vdash A@w \cdot v}{\Gamma ; \Delta \vdash \rho_v A@w} \rho R$$

The proposition $\rho_v A$ thus stands for an *intermediate state* in a transition to A . Informally it can be thought to be “ v before A ”; thus, $\forall v. \rho_v A$ represents *all* intermediate states in the path to A , and $\exists v. \rho_v A$ represents *some* such state. The modally unrestricted proposition $\dagger A$ represents a resource that is consumable in any world; it is mainly used to make transition rules applicable at all worlds.

It is worth remarking that the HyLL proof theory can be seen as at least as powerful as (the linear restriction of) intuitionistic S5 [36]:

► **Theorem 9** (HyLL is S5). *The following sequent is derivable: $\cdot ; \Diamond A@w \Longrightarrow \Box \Diamond A@w$.*

Proof. See [9]. ◀

Obviously HyLL is more expressive as it allows direct manipulation of the worlds using the hybrid connectives: for example, the ρ connective is not definable in S5.

2.2 Temporal Constraints

As a pedagogical example, consider the constraint domain $\mathcal{T} = \langle \mathbf{R}^+, +, 0 \rangle$ representing instants of time. This domain can be used to define the lifetime of resources, such as keys, sessions, or delegations of authority. Delay (defn. 8) in $\text{HyLL}(\mathcal{T})$ represents intervals of time; $\rho_d A$ means “ A will become available after delay d ”, similar to metric tense logic [31]. This domain is very permissive because addition is commutative, resulting in the equivalence of $\rho_u \rho_v A$ and $\rho_v \rho_u A$. The “forward-looking” connectives G and F of ordinary tense logic are precisely \Box and \Diamond of defn. 8.

In addition to the future connectives, the domain \mathcal{T} also admits past connectives if we add saturating subtraction (*i.e.*, $a - b = 0$ if $b \geq a$) to the language of worlds. We can then define the duals H and P of G and F as:

$$H A \triangleq \downarrow u. \forall w. (A \text{ at } u - w) \quad P A \triangleq \downarrow u. \exists w. (A \text{ at } u - w)$$

While this domain does not have any branching structure like CTL, it is expressive enough for many common idioms because of the branching structure of derivations involving \oplus . CTL reachability (“in some path in some future”), for instance, is the same as our \Diamond ; similarly CTL steadiness (“in some path for all futures”) is the same as \Box . CTL stability (“in all paths in all futures”), however, has no direct correspondance in HyLL. Note that model checking cannot cope with temporal expressions involving the “in all paths” notion anyway. Thus approaches using ordinary temporal logics and model checking, like BIOCHAM, for example, cannot deal with those expressions either.

On the other hand, the availability of linear reasoning makes certain kinds of reasoning in HyLL much more natural than in ordinary temporal logics. One important example is of *oscillation* between



states in systems with kinetic feedback. In a temporal specification language such as BIOCHAM [7], only finite oscillations are representable using a nested syntax, while in HyLL we use a simple bi-implication; for example, the oscillation between A and B with delay d is represented by the rule $\dagger(A \multimap \rho_d B) \& (B \multimap \rho_d A)$ (or $\dagger(A \multimap \diamond B) \& (B \multimap \diamond A)$ if the oscillation is aperiodic). If HyLL(\mathcal{T}) were extended with constrained implication and conjunction in the style of CILL [35] or η [13], then we can define localized versions of \square and \diamond , such as “ A is true everywhere/somewhere in an interval”. They would also allow us to define the “until” and “since” operators of linear temporal logic [20].

2.3 Probabilistic Constraints

The material in this section requires some background in probability and measure theory, and can be skipped at a first reading, without significant loss of continuity.

Transitions in practice rarely have precise delays. Phenomenological and experimental evidence is used to construct a probabilistic model of the transition system where the delays are specified as probability distributions of continuous variables.

2.3.1 General Case

Let us recall some basic definitions in probability theory. The probability of X being in A is defined by: $\mu_X(A) \triangleq \int_{x \in A} \mu_X dx$. For example: $\text{Prob}(X \leq x) = \mu_X(x) = \int_{-\infty}^x \mu_X dt$.

► **Fact 1** (see [16, 34]). If X and Y are independent random variables in \mathbf{R} , with distribution μ_X and μ_Y , respectively, then the distribution μ_{X+Y} of the random variable $X + Y$ is given by

$$\mu_{X+Y}(A) = \mu_X * \mu_Y(A) = \int_{\{x+y \in A\}} \mu_X(dx) \otimes \mu_Y(dy)$$

for all Borel¹ subset A of \mathbf{R} .

The space of probability distributions of random variables, together with the convolution operator $*$ and the Dirac mass at 0 (δ_0)² as neutral element, forms a monoid. More precisely:

► **Definition 10.** The *probabilities domain* \mathcal{P} is the monoid $\langle \mathbf{M}_1(\mathbf{R}), *, \delta_0 \rangle$ where $\mathbf{M}_1(\mathbf{R})$ is the set of the Borel probability measures over \mathbf{R} and δ_0 is Dirac mass at 0. The instance HyLL(\mathcal{P}) will sometimes be called “probabilistic hybrid linear logic”.

An element $w = \mu_X(x)$ of a world \mathcal{P} thus represents the probability of X to have its value in the interval $[-\infty, x]$. $\bar{w} = 1 - \mu_X(x)$ represents the probability of X to have its value greater than x . $A @ \bar{w}$ therefore means “ A is true with probability greater than x ”.

2.3.2 Markov Processes

The meaning of the random variable X depends on the intended application. In the applications in the area of systems biology, X can represent the concentration of a product, while in economics, X could be the duration of an activity, for example. The standard model of stochastic transition systems is continuous time Markov chains (CTMCs) where the delays of transitions between states are distributed according to the Markov assumption of memorylessness (Markov processes) with the further condition that their state-space are countable sets [34].

¹ The set of Borel events is the set of the Lebesgue measurable functions.

² The Dirac mass at 0 δ_0 is not a real function but a generalized function. It is just a measure.



► **Fact 2** (see [15, 34]). Given a continuous-time Markov process $(X_t, t \geq 0)$ taking values in a measurable space (E, \mathcal{E}) , the family $(P(t), t \geq 0)$ of linear operators on the set of bounded Borel functions $\mathcal{B}(E)$ defined by: for all $f \in \mathcal{B}(E)$ and for all $x \in E$,

$$(P(t)f)(x) = \mathbf{E}[f(X_t) \mid X_0 = x],$$

where \mathbf{E} is the expectation function, is a semigroup for the convolution: for all $s, t \geq 0$,

$$P(t + s) = P(t) * P(s)$$

with neutral element $P(0)$, the identity operator. When the process X is a Feller process (see [34], chapter 3, section 2), $(P(t), t \geq 0)$ is a Feller semigroup, i.e. strongly continuous and conservative.

For a given continuous-time Markov process $(X_t, t \geq 0)$, the associated monoid is based on the set $(P(t), t \geq 0)$ defined above. More precisely, we can define the Markov domain as follows:

► **Definition 11.** For a given continuous-time Markov process $(X_t, t \geq 0)$, taking values in a measurable space (E, \mathcal{E}) , the *Markov domain* \mathcal{M} is the monoid $\langle (P(t), t \geq 0), *, P(0) \rangle$ where $(P(t), t \geq 0)$ is the sub-Markov semigroup of linear operators on the set of bounded Borel functions $\mathcal{B}(E)$ defined by for all $f \in \mathcal{B}(E)$ and for all $x \in E$, $(P(t)f)(x) = \mathbf{E}[f(X_t) \mid X_0 = x]$. The instance $\text{HyLL}(\mathcal{M})$ will sometimes be called “Markov hybrid linear logic”.

In the above definition, f can be any function in $\mathcal{B}(E)$, and we have

$$(P(t)f)(x) = \mathbf{E}[f(X_t) \mid X_0 = x].$$

An element $w = P(t)$ of \mathcal{M} represents a function which associates to any function f (where $f \in \mathcal{B}(E)$) and to any initial value x for the variable X_t , the expectation of $f(X_t)$, knowing that $X_0 = x$. We can choose $f = \mathbf{1}_A$: the indicator (i.e. characteristic) function of a set A . In this case (recalling that $\mathbf{E}(\mathbf{1}_A) = P(A)$),

$$(P(t)f)(x) = (P(t)\mathbf{1}_A)(x) = \mathbf{E}[\mathbf{1}_A(X_t) \mid X_0 = x] = P\{X_t \in A \mid X_0 = x\}$$

For example, for $A = [-\infty, y]$: $(P(t)\mathbf{1}_A)(x) = P\{X_t \leq y \mid X_0 = x\} = \mathbf{F}_{X_t|X_0=x}(y)$ where \mathbf{F} is the cumulative distribution function of X_t . Other interesting examples for f are the square function $\text{sq}(y) = y^2$ and the identity function $\text{id}(y) = y$. Using these functions, we can define the variance of X_t :

$$(P(t) \text{sq})(x) - (P(t) \text{id})^2(x) = \mathbf{E}(X_t^2 \mid X_0 = x) - (\mathbf{E}(X_t \mid X_0 = x))^2 = \mathbf{Var}(X_t \mid X_0 = x)$$

In principle, using suitable functions f , we should be able to define *any descriptors of the probability distribution of our variable X_t* .

The meaning of $A @ w$ varies depending on the choice for the function f . For example, we have seen that in the case of $f = \mathbf{1}_{[-\infty, y]}$, $P(t)f(x) = \mathbf{F}_{X_t|X_0=x}(y)$. In this case, $A @ w$ means “ A is true with probability less than y at time t ”. For $\bar{w} = 1 - P(t)$, $A @ \bar{w}$ means “ A is true with probability greater than y at time t ”.

The cumulative distributions of the continuous time Markov chains (CTMCs) used in $S\pi$ are exponential [29]. More precisely: $\text{Prob}(X_t \leq x + rt \mid X_0 = x) = \mathbf{F}_{X_t|X_0=x}(x + rt) = 1 - e^{-rt}$, where r (rates) are functions depending on the time t . In this case, we can define the worlds w by particularizing the general case of probabilities domains (defn. 10 - $w = \mu_X(y)$), to the case where the cumulative distributions of all variables X are exponential with the above meaning. The worlds w will therefore be defined by $w = \mu_X(x + rt) = \text{Prob}(X_t \leq x + rt \mid X_0 = x) = \mathbf{F}_{X_t|X_0=x}(x + rt) = 1 - e^{-rt}$.



3 Focusing

As HyLL is intended to represent transition systems adequately, it is crucial that HyLL derivations in the image of an encoding have corresponding transitions. However, transition systems are generally specified as rewrite algebras over an underlying congruence relation. These congruences have to be encoded propositionally in HyLL, so a HyLL derivation will generally require several inference rules to implement a single transition; moreover, several trivially different reorderings of these “micro” inferences would correspond to the same transition. It is therefore futile to attempt to define an operational semantics directly on HyLL inferences.

We restrict the syntax to focused derivations [1], which ignores many irrelevant rule permutations in a sequent proof and divides the proof into clear *phases* that define the grain of atomicity. The logical connectives are divided into two classes, *negative* and *positive*, and rule permutations for connectives of like polarity are confined to *phases*. A *focused derivation* is one in which the positive and negative rules are applied in alternate maximal phases in the following way: in the *active* phase, all negative rules are applied (in irrelevant order) until no further negative rule can apply; the phase then switches and one positive proposition is selected for *focus*; this focused proposition is decomposed under focus (*i.e.*, the focus persists unto its sub-propositions) until it becomes negative, and the phase switches again.

As noted before, the logical rules of the hybrid connectives at and \downarrow are invertible, so they can be considered to have both polarities. It would be valid to decide a polarity for each occurrence of each hybrid connective independently; however, as they are mainly intended for book-keeping during logical reasoning, we define the polarity of these connectives in the following *parasitic* form: if its immediate subformula is positive (resp. negative) connective, then it is itself positive (resp. negative). These connectives therefore become invisible to focusing. This choice of polarity can be seen as a particular instance of a general scheme that divides the \downarrow and at connectives into two polarized forms each. To complete the picture, we also assign a polarity for the atomic propositions; this restricts the shape of focusing phases further [10]. The full syntax of positive (P, Q, \dots) and negative (M, N, \dots) propositions is as follows:

$$\begin{aligned} P, Q, \dots &::= p \uparrow | P \otimes Q | \mathbf{1} | P \oplus Q | \mathbf{0} | !N | \exists \alpha. P | \downarrow u. P | (P \text{ at } w) | \downarrow N \\ N, M, \dots &::= n \uparrow | N \& N | \top | P \multimap N | \forall \alpha. N | \downarrow u. N | (N \text{ at } w) | \uparrow P \end{aligned}$$

The two syntactic classes refer to each other via the new connectives \uparrow and \downarrow . Sequents in the focusing calculus are of the following forms.

$$\left. \begin{array}{l} \Gamma ; \Delta ; \Omega \Longrightarrow \cdot ; P @ w \\ \Gamma ; \Delta ; \Omega \Longrightarrow N @ w ; \cdot \end{array} \right\} \text{active} \qquad \left. \begin{array}{l} \Gamma ; \Delta ; [N @ u] \Longrightarrow P @ w \\ \Gamma ; \Delta \Longrightarrow [P @ w] \end{array} \right\} \text{focused}$$

In each case, Γ and Δ contain only negative propositions (*i.e.*, of the form $N @ u$) and Ω only positive propositions (*i.e.*, of the form $P @ u$). The full collection of inference rules are in fig. 2. The sequent form $\Gamma ; \Delta ; \cdot \Longrightarrow \cdot ; P @ w$ is called a *neutral sequent*; from such a sequent, a left or right focused sequent is produced with the rules lf , cplf or rf . Focused logical rules are applied (non-deterministically) and focus persists unto the subformulas of the focused proposition as long as they are of the same polarity; when the polarity switches, the result is an active sequent, where the propositions in the innermost zones are decomposed in an irrelevant order until once again a neutral sequent results.

Soundness of the focusing calculus with respect to the ordinary sequent calculus is immediate by simple structural induction. In each case, if we forget the additional structure in the focused derivations, then we obtain simply an unfocused proof. We omit the obvious theorem. Completeness, on the other hand, is a hard result. We omit the proof because focusing is by now well known for linear logic, with a number of distinct proofs via focused cut-elimination (see *e.g.* the detailed proof in [10]). The hybrid connectives pose no problems because they allow all cut-permutations.



3.0.2.1 Focused logical rules

$$\begin{array}{c}
\frac{}{\Gamma; [n \vec{t}@w] \Rightarrow \Downarrow n \vec{t}@w} \text{ li} \quad \frac{\Gamma; \Delta; P@u \Rightarrow \cdot; Q@w}{\Gamma; \Delta; [\uparrow P@u] \Rightarrow Q@w} \uparrow L \quad \frac{\Gamma; \Delta; [N_i@u] \Rightarrow Q@w}{\Gamma; \Delta; [N_1 \& N_2@u] \Rightarrow Q@w} \&L_i \\
\frac{\Gamma; \Delta \Rightarrow [P@u] \quad \Gamma; \Xi; [N@u] \Rightarrow Q@w}{\Gamma; \Delta, \Xi; [P \multimap N@u] \Rightarrow Q@w} \multimap L \quad \frac{\Gamma; \Delta; [\tau/\alpha]N@u \Rightarrow Q@w}{\Gamma; \Delta; [\forall \alpha. N@u] \Rightarrow Q@w} \forall L \\
\frac{\Gamma; \Delta; [v/u]N@v \Rightarrow Q@w}{\Gamma; \Delta; [\downarrow u. N@v] \Rightarrow Q@w} \downarrow LF \quad \frac{\Gamma; \Delta; [N@u] \Rightarrow Q@w}{\Gamma; \Delta; [(N \text{ at } u)@v] \Rightarrow Q@w} \text{ at } LF \\
\frac{}{\Gamma; \uparrow p \vec{t}@w \Rightarrow [p \vec{t}@w]} \text{ ri} \quad \frac{\Gamma; \Delta; \cdot \Rightarrow N@w; \cdot}{\Gamma; \Delta \Rightarrow [\Downarrow N@w]} \Downarrow R \quad \frac{\Gamma; \Delta \Rightarrow [P@w] \quad \Gamma; \Xi \Rightarrow [Q@w]}{\Gamma; \Delta, \Xi \Rightarrow [P \otimes Q@w]} \otimes R \\
\frac{}{\Gamma; \cdot \Rightarrow [1@w]} \mathbf{1}R \quad \frac{\Gamma; \Delta \Rightarrow [P_i@w]}{\Gamma; \Delta \Rightarrow [P_1 \oplus P_2@w]} \oplus R_i \quad \frac{\Gamma; \cdot; \cdot \Rightarrow N@w; \cdot}{\Gamma; \cdot \Rightarrow [!N]@w} !R \\
\frac{\Gamma; \Delta \Rightarrow [[\tau/\alpha]P@w]}{\Gamma; \Delta \Rightarrow [\exists \alpha. P@w]} \exists R \quad \frac{\Gamma; \Delta \Rightarrow [[w/u]P@w]}{\Gamma; \Delta \Rightarrow [\downarrow u. P@w]} \downarrow RF \quad \frac{\Gamma; \Delta \Rightarrow [P@u]}{\Gamma; \Delta \Rightarrow [(P \text{ at } u)@w]} \text{ at } RF
\end{array}$$

3.0.2.2 Active logical rules

(R of the form $\cdot; Q@w$ or $N@w; \cdot$, and L of the form $\Gamma; \Delta; \Omega$)

$$\begin{array}{c}
\frac{L, P@u, Q@u \Rightarrow R}{L, P \otimes Q@u \Rightarrow R} \otimes L \quad \frac{L \Rightarrow R}{L, 1@u \Rightarrow R} \mathbf{1}L \quad \frac{L, P@u \Rightarrow R \quad L, Q@u \Rightarrow R}{L, P \oplus Q@u \Rightarrow R} \oplus L \quad \frac{}{L, 0@u \Rightarrow R} \mathbf{0}L \\
\frac{L, [v/u]P@v \Rightarrow R}{L, \downarrow u. P@v \Rightarrow R} \downarrow LA \quad \frac{L, P@u \Rightarrow R}{L, (P \text{ at } u)@v \Rightarrow R} \text{ at } LA \quad \frac{L, P@u \Rightarrow R}{L, \exists \alpha. P@u \Rightarrow R} \exists L^\alpha \\
\frac{\Gamma, N@u; \Delta; \Omega \Rightarrow R}{\Gamma; \Delta; \Omega, !N@u \Rightarrow R} !L \quad \frac{\Gamma; \Delta, N@w; \Omega \Rightarrow R}{\Gamma; \Delta; \Omega, \Downarrow N@w \Rightarrow R} \Downarrow L \quad \frac{\Gamma; \Delta, \uparrow p \vec{t}; \Omega \Rightarrow R}{\Gamma; \Delta; \Omega, p \vec{t}@w \Rightarrow R} \text{ lp} \\
\frac{L \Rightarrow M@w; \cdot \quad L \Rightarrow N@w; \cdot}{L \Rightarrow M \& N@w; \cdot} \&R \quad \frac{}{L \Rightarrow \top@w; \cdot} \top R \quad \frac{L, P@w \Rightarrow N@w; \cdot}{L \Rightarrow P \multimap N@w; \cdot} \multimap R \\
\frac{L \Rightarrow [w/u]N@w; \cdot}{L \Rightarrow \downarrow u. N@w; \cdot} \downarrow RA \quad \frac{L \Rightarrow N@u}{L \Rightarrow (N \text{ at } u)@w} \text{ at } RA \quad \frac{L \Rightarrow N@u; \cdot}{L \Rightarrow \forall \alpha. N@u; \cdot} \forall R^\alpha \\
\frac{L \Rightarrow \cdot; P@w}{L \Rightarrow \uparrow P@w; \cdot} \uparrow R \quad \frac{L \Rightarrow \cdot; \Downarrow n \vec{t}@w}{L \Rightarrow n \vec{t}@w; \cdot} \text{ rp}
\end{array}$$

3.0.2.3 Focusing decisions

(L of the form $\Gamma; \Delta$)

$$\frac{\Gamma; \Delta; [N@u] \Rightarrow Q@w \quad N \text{ not } \uparrow p \vec{t}}{\Gamma; \Delta, N@u; \cdot \Rightarrow \cdot; Q@w} \text{ lf} \quad \frac{\Gamma, N@u; \Delta; [N@u] \Rightarrow Q@w}{\Gamma, N@u; \Delta; \cdot \Rightarrow \cdot; Q@w} \text{ cplf} \\
\frac{\Gamma; \Delta \Rightarrow [P@w] \quad P \text{ not } \downarrow n \vec{t}}{\Gamma; \Delta; \cdot \Rightarrow \cdot; P@w} \text{ rf}$$

■ Figure 2 Focusing rules for HyLL.



© Joëlle Despeyroux and Kaustuv Chaudhuri;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

► **Theorem 12** (focusing completeness). *Let Γ^- and $C^-@w$ be negative polarizations of Γ and $C@w$ (that is, adding \uparrow and \downarrow to make C and each proposition in Γ negative) and Δ^+ be a positive polarization of Δ . If $\Gamma; \Delta \Longrightarrow C@w$, then $\cdot; \cdot; !\Gamma^-, \Delta^+ \Longrightarrow C^-@w; \cdot$.*

4 Encoding the Synchronous Stochastic π -calculus

In this section, we shall illustrate the use of $\text{HyLL}(\mathcal{P})$ as a logical framework for constrained transition systems by encoding the syntax and the operational semantics of the synchronous stochastic π -calculus ($S\pi$), which extends the ordinary π -calculus by assigning to every channel and internal action an *inherent* rate of synchronization. In $S\pi$, each rate characterises an exponential distribution [29], such that the probability of a reaction with rate r is given by $\text{Prob}(X_t \leq x + rt \mid X_0 = x) = \mathbf{F}_{X_t \mid X_0 = x}(x + rt) = 1 - e^{-rt}$, where the rates r are functions depending on the time t . We have seen in sec. 2.3 that, in this case, we can use a particular instance of $\text{HyLL}(\mathcal{P})$, where the worlds w are $\mu_X(x + rt)$, representing the probability of X to have its value less or equal than $x + rt$. Note that the distributions have the same shape for any variables X_t ; They only depend on rates $r(t)$ and time t . We shall use this fact to encode $S\pi$ in $\text{HyLL}(\mathcal{P})$: a $S\pi$ reaction with rate r will be encoded by a transition of probability $w = \mu_X(x + rt) = \text{Prob}(X_t \leq x + rt \mid X_0 = x) = \mathbf{F}_{X_t \mid X_0 = x}(x + rt) = 1 - e^{-rt}$. In the rest of this section, worlds w of this shape, defined by a rate r , will simply be written r (see defn. 13 and defn. 15).

$\text{HyLL}(\mathcal{P})$ can therefore be seen as a formal language for expressing $S\pi$ executions (traces). For the rest of this section we shall use r, s, t, \dots instead of u, v, w, \dots to highlight the fact that the worlds represent rates, with the understanding that \cdot is convolution (fact 1) and ι is Θ . We don't directly use rates because the syntax and transitions of $S\pi$ are given generically for a π -calculus with labelled actions, and it is only the interpretation of the labels that involves probabilities.

We first summarize the syntax of $S\pi$, which is a minor variant of a number of similar presentations such as [29]. For hygienic reasons we divide entities into the syntactic categories of *processes* (P, Q, \dots) and *sums* (M, N, \dots), defined as follows. We also include environments of recursive definitions (E) for constants.

$$\begin{aligned} (\text{Processes}) \quad P, Q, \dots &::= \nu_r P \mid P \mid Q \mid 0 \mid X_n x_1 \cdots x_n \mid M \\ (\text{Sums}) \quad M, N, \dots &::= !x(y). P \mid ?x. P \mid \tau_r. P \mid M + N \\ (\text{Environments}) \quad E &::= E, X_n \triangleq P \mid \cdot \end{aligned}$$

$P \mid Q$ is the parallel composition of P and Q , with unit 0. The restriction $\nu_r P$ abstracts over a free channel x in the process $P x$. We write the process using higher-order abstract syntax [26], *i.e.*, P in $\nu_r P$ is (syntactically) a function from channels to processes. This style lets us avoid cumbersome binding rules in the interactions because we reuse the well-understood binding structure of the λ -calculus. A similar approach was taken in the earliest encoding of (ordinary) π -calculus in (unfocused) linear logic [23], and is also present in the encoding in CLF [6].

A sum is a non-empty choice (+) over terms with *action prefixes*: the output action $!x(y)$ sends y along channel x , the input action $?x$ reads a value from x (which is applied to its continuation process), and the internal action τ_r has no observable I/O behaviour. Replication of processes happens via guarded recursive definitions [24]; in [33] it is argued that they are more practical for programming than the replication operator $!$. In a definition $X_n \triangleq P$, X_n denotes a (higher-order) defined constant of arity n ; given channels x_1, \dots, x_n , the process $X_n x_1 \cdots x_n$ is synonymous with $P x_1 \cdots x_n$. The constant X_n may occur on the right hand side of any definition in E , including in its body P , as long as it is prefixed by an action; this prevents infinite recursion without progress.

Interactions are of the form $E \vdash P \xrightarrow{r} Q$ denoting a transition from the process P to the process Q , in a global environment E , by performing an action at rate r . Each channel x is associated with an inherent rate specific to the channel, and internal actions τ_r have rate r . The restriction $\nu_r P$ defines the rate of the abstracted channel as r .



<i>Interactions</i>	
$\frac{}{!x(y). P + M \mid ?x. Q + M' \xrightarrow{\text{rate}(x)} P \mid Q} \text{ SYN}$	$\frac{}{\tau_r. P \xrightarrow{r} P} \text{ INT}$
$\frac{P \xrightarrow{r} P'}{P \mid Q \xrightarrow{r} P' \mid Q} \text{ PAR}$	$\frac{\forall x_s. (P x \xrightarrow{r} Q x)}{v_s P \xrightarrow{r} v_s Q} \text{ RES}$
.....	
<i>Congruence</i>	
$\frac{}{P \mid 0 \equiv P}$	$\frac{}{P \mid Q \equiv Q \mid P}$
$\frac{}{P \mid (Q \mid R) \equiv (P \mid Q) \mid R}$	$\frac{}{v_r 0 \equiv 0}$
$X_n \triangleq P \in E$	
$E \vdash X_n x_1 \cdots x_n \equiv P x_1 \cdots x_n$	
$\frac{}{v_r(\lambda x. v_s(\lambda y. P)) \equiv v_s(\lambda y. v_r(\lambda x. P))}$	$\frac{\forall x_r. (P x \equiv Q x)}{v_r P \equiv v_r Q}$
$v_r(\lambda x. P \mid Q(x)) \equiv P \mid v_r Q$	
$\frac{P \equiv P'}{P \mid Q \equiv P' \mid Q}$	$\frac{P \equiv P'}{!x(m). P \equiv !x(m). P'}$
$\frac{\forall n. (P n \equiv Q n)}{?x. P \equiv ?x. Q}$	$\frac{P \equiv P'}{\tau_r. P \equiv \tau_r. P'}$
$M \equiv M'$	
$\frac{}{M + N \equiv N + M}$	$\frac{}{M + (N + K) \equiv (M + N) + K}$
$M + N \equiv M' + N$	
$M + N \equiv M$	

■ **Figure 3** Interactions and congruence in $S\pi$. The environment E is elided in most rules.

The full set of interactions and congruences are in fig. 3. We generally omit the global environment E in the rules as it never changes. It is possible to use the congruences to compute a normal form for processes that are a parallel composition of sums and each reaction selects two suitable sums to synchronise on a channel until there are no further reactions possible; this refinement of the operational semantics is used in $S\pi$ simulators such as SPiM [28].

► **Definition 13** (syntax encoding).

1. The encoding of the process P as a positive proposition, written $\llbracket P \rrbracket_p$, is as follows (sel is a positive atom and rt a negative atom).

$$\begin{aligned} \llbracket P \mid Q \rrbracket_p &= \llbracket P \rrbracket_p \otimes \llbracket Q \rrbracket_p & \llbracket v_r P \rrbracket_p &= \exists x. !(rt \ x \ \text{at} \ r) \otimes \llbracket P x \rrbracket_p \\ \llbracket 0 \rrbracket_p &= \mathbf{1} & \llbracket X_n x_1 \cdots x_n \rrbracket_p &= X_n x_1 \cdots x_n \\ \llbracket M \rrbracket_p &= \Downarrow(\text{sel} \ \neg \llbracket M \rrbracket_s) \end{aligned}$$

2. The encoding of the sum M as a negative proposition, written $\llbracket M \rrbracket_s$, is as follows (out, in and tau are positive atoms).

$$\begin{aligned} \llbracket M + N \rrbracket_s &= \llbracket M \rrbracket_s \ \& \ \llbracket N \rrbracket_s & \llbracket !x(m). P \rrbracket_s &= \Uparrow(\text{out} \ x \ m \ \otimes \ \llbracket P \rrbracket_p) \\ \llbracket ?x. P \rrbracket_s &= \forall n. \Uparrow(\text{in} \ x \ n \ \otimes \ \llbracket P n \rrbracket_p) & \llbracket \tau_r. P \rrbracket_s &= \Uparrow(\text{tau} \ r \ \otimes \ \llbracket P \rrbracket_p) \end{aligned}$$

3. The encoding of the definitions E as a context, written $\llbracket E \rrbracket_e$, is as follows.

$$\begin{aligned} \llbracket E, X_n \triangleq P \rrbracket_e &= \llbracket E \rrbracket_e, \dagger \forall x_1, \dots, x_n. X_n x_1 \cdots x_n \ \circ\text{-}\circ \ \llbracket P x_1 \cdots x_n \rrbracket_p \\ \llbracket \cdot \rrbracket_e &= \cdot \end{aligned}$$

where $P \circ\text{-}\circ Q$ is defined as $(P \neg \Uparrow Q) \ \& \ (Q \neg \Uparrow P)$.

The encoding of processes is positive, so they will be decomposed in the active phase when they occur on the left of the sequent arrow, leaving a collection of sums. The encoding of restrictions will introduce a fresh unrestricted assumption about the rate of the restricted channel. Each sum encoded as a processes undergoes a polarity switch because \neg is negative; the antecedent of this implication is a *guard* sel. This pattern of guarded switching of polarities prevents unsound congruences such as



$!x(m). !y(n). P \equiv !y(n). !x(m). P$ that do not hold for the synchronous π calculus.³ This guard also *locks* the sums in the context: the $S\pi$ interaction rules INT and SYN discard the non-interacting terms of the sum, so the environment will contain the requisite number of sel s only when an interaction is in progress. The action prefixes themselves are also synchronous, which causes another polarity switch. Each action releases a token of its respective kind— out , in or tau —into the context. These tokens must be consumed by the interaction before the next interaction occurs. For each action, the (encoding of the) continuation process is also released into the context.

The proof of the following congruence lemma is omitted. Because the encoding is (essentially) a $\otimes/\&$ structure, there are no distributive laws in linear logic that would break the process/sum structure.

► **Theorem 14** (congruence).

$E \vdash P \equiv Q$ iff both $\llbracket E \rrbracket_e @ \iota ; \cdot ; \llbracket P \rrbracket_p @ \iota \Longrightarrow \cdot ; \llbracket Q \rrbracket_p @ \iota$ and $\llbracket E \rrbracket_e @ \iota ; \cdot ; \llbracket Q \rrbracket_p @ \iota \Longrightarrow \cdot ; \llbracket P \rrbracket_p @ \iota$.

Now we encode the interactions. Because processes were lifted into propositions, we can be parsimonious with our encoding of interactions by limiting ourselves to the atomic interactions SYN and INT (below); the PAR , RES and CONG interactions will be ambiently implemented by the logic. Because there are no concurrent interactions—only one interaction can trigger at a time in a trace—the interaction rules must obey a locking discipline. We represent this lock as the proposition act that is consumed at the start of an interaction and produced again at the end. This lock also carries the net rate of the prefix of the trace so far: that is, an interaction $P \xrightarrow{r} Q$ will update the lock from $\text{act} @ s$ to $\text{act} @ s \cdot r$. The encoding of individual atomic interactions must also remove the in , out and tau tokens introduced in context by the interacting processes.

► **Definition 15** (interaction).

Let $\text{inter} \triangleq \dagger(\text{act} \multimap \uparrow\text{int} \ \& \ \uparrow\text{syn})$ where act is a positive atom and int and syn are as follows:

$$\begin{aligned} \text{int} &\triangleq (\text{sel} \text{ at } \iota) \otimes \Downarrow \forall r. ((\text{tau } r \text{ at } \iota) \multimap \rho_r \uparrow\text{act}) \\ \text{syn} &\triangleq (\text{sel} \otimes \text{sel} \text{ at } \iota) \otimes \Downarrow \forall x, r, m. ((\text{out } x \ m \otimes \text{in } x \ m \text{ at } \iota) \multimap \Downarrow (\text{rt } x \text{ at } r) \multimap \rho_r \uparrow\text{act}). \end{aligned}$$

The number of interactions that are allowed depend on the number of instances of inter in the linear context: each focus on inter implements a single interaction. If we are interested in all finite traces, we will add inter to the unrestricted context so it may be reused as many times as needed.

4.1 Representational Adequacy.

Adequacy consists of two components: completeness and soundness. Completeness is the property that every $S\pi$ execution is obtainable as a HyLL derivation using this encoding, and is the comparatively simpler direction (see thm. 18). Soundness is the reverse property, and is false for unfocused HyLL as such. However, it *does* hold for focused proofs (see thm. 20). In both cases, we reason about the following canonical sequents of HyLL.

► **Definition 16.** The *canonical context* of P , written $\wr P \wr$, is given by:

$$\begin{aligned} \wr X_n x_1 \cdots x_n \wr &= \uparrow X_n x_1 \cdots x_n & \wr P \wr Q \wr &= \wr P \wr, \wr Q \wr & \wr 0 \wr &= \cdot & \wr \nu_r P \wr &= \wr P a \wr \\ \wr M \wr &= \text{sel} \multimap \llbracket M \rrbracket_s \end{aligned}$$

For $\wr \nu_r P \wr$, the right hand side uses a *fresh* channel a that is not free in the rest of the sequent it occurs in.

³ Note: $(x \multimap a \otimes (x \multimap b \otimes c)) \multimap (x \multimap b \otimes (x \multimap a \otimes c))$ is not provable in linear logic.



Suppose $L = \text{rt}x@r, \text{inter}@t$ and $R = (\llbracket S \rrbracket_p \text{ at } t) \otimes \text{act}@t$. (All judgements $@t$ omitted.)

$$\begin{array}{c}
\frac{L; \{Q\}, \{Ra\}, \uparrow \text{act}@s \cdot r; \cdot \Longrightarrow \cdot; R}{5} \\
\frac{L; \{Q\}, \uparrow \text{out } xa, \uparrow \text{in } xa, \{Ra\}, \\ \forall x, r, m. ((\text{out } xm \otimes \text{in } xm \text{ at } t) \multimap \Downarrow (\text{rt}x \text{ at } r) \multimap \rho_r \text{ act})@s \quad ; \cdot \Longrightarrow \cdot; R}{4} \\
\frac{L; \uparrow \text{out } xa, \{Q\}, \text{sel} \multimap \forall y. \uparrow (\text{in } xy \otimes \llbracket Ry \rrbracket_p), \\ \uparrow \text{sel}, \forall x, r, m. ((\text{out } xm \otimes \text{in } xm \text{ at } t) \multimap \Downarrow (\text{rt}x \text{ at } r) \multimap \rho_r \text{ act})@s \quad ; \cdot \Longrightarrow \cdot; R}{3} \\
\frac{L; \text{sel} \multimap \uparrow (\text{out } xa \otimes \llbracket Q \rrbracket_p), \text{sel} \multimap \forall y. (\text{in } xy \otimes \llbracket Ry \rrbracket_p), \\ \uparrow \text{sel}, \uparrow \text{sel}, \forall x, r, m. ((\text{out } xm \otimes \text{in } xm \text{ at } t) \multimap \Downarrow (\text{rt}x \text{ at } r) \multimap \rho_r \text{ act})@s \quad ; \cdot \Longrightarrow \cdot; R}{2} \\
\frac{L; \uparrow \text{act}@s, \text{sel} \multimap \uparrow (\text{out } xa \otimes \llbracket Q \rrbracket_p), \text{sel} \multimap \forall y. \uparrow (\text{in } xy \otimes \llbracket Ry \rrbracket_p); [\text{inter}] \Longrightarrow R}{1} \\
\frac{L; \uparrow \text{act}@s, \text{sel} \multimap \uparrow (\text{out } xa \otimes \llbracket Q \rrbracket_p), \text{sel} \multimap \forall y. \uparrow (\text{in } xy \otimes \llbracket Ry \rrbracket_p); \cdot \Longrightarrow \cdot; R}{1} \\
\frac{L; \uparrow \text{act}@s, \{!x(a). Q \mid ?x. R\}; \cdot \Longrightarrow \cdot; R}{1} \\
\hline
\text{Steps} \\
\begin{array}{lll}
1: \text{focus on } \text{inter} \in L & 3: \text{sel for output + full phases} & 5: \text{cleanup} \\
2: \text{select syn from } \text{inter}, \text{active rules} & 4: \text{sel for input + full phases} &
\end{array}
\end{array}$$

■ **Figure 4** Example interaction in the $S\pi$ -encoding.

As an illustration, take $P \triangleq !x(a). Q \mid ?x. R$. We have:

$$\{P\} = \text{sel} \multimap \uparrow (\text{out } xa \otimes \llbracket Q \rrbracket_p), \text{sel} \multimap \forall y. \uparrow (\text{in } xy \otimes \llbracket Ry \rrbracket_p)$$

Obviously, the canonical context is what would be emitted to the linear zone at the end of the active phase if $\llbracket P \rrbracket_p$ were to be present in the left active zone.

► **Definition 17.** A neutral sequent is *canonical* iff it has the shape

$$\llbracket E \rrbracket_e, \text{rates}, \text{inter}@t; \uparrow \text{act}@s, \{P_1 \mid \dots \mid P_k\}@t; \cdot \Longrightarrow \cdot; (\llbracket Q \rrbracket_p \text{ at } t) \otimes \text{act}@t$$

where *rates* contains elements of the form $\text{rt } x@r$ defining the rate of the channel x as r , and all free channels in $\llbracket E \rrbracket_e, \{P_1 \mid \dots \mid P_k \mid Q\}$ have a single such entry in *rates*.

Figure 4 contains an example of a derivation for a canonical sequent involving P . Focusing on any (encoding of a) sum in $\{P\}@t$ will fail because there is no sel in the context, so only inter can be given focus; this will consume the act and release two copies of $(\text{sel} \text{ at } t)$ and the continuation into the context. Focusing on the latter will fail now (because $\text{out } xm$ and $\text{in } xm$ (for some m) are not yet available), so the only applicable foci are the two sums that can now be “unlocked” using the sel s. The output and input can be unlocked in an irrelevant order, producing two tokens $\text{in } xa$ and $\text{out } xa$. Note in particular that the witness a was chosen for the universal quantifier in the encoding of $?x. Q$ because the subsequent consumption of these two tokens requires the messages to be identical. (Any other choice will not lead to a successful proof.) After both tokens are consumed, we get the final form $\text{act}@s \cdot r$, where r is the inherent rate of x (found from the *rates* component of the unrestricted zone). This sequent is canonical and contains $\{Q \mid Ra\}$.

Our encoding therefore represents every $S\pi$ action in terms of “micro” actions in the following rigid order: one micro action to determine what kind of action (internal or synchronization), one micro action per sum to select the term(s) that will interact, and finally one micro action to establish the contract of the action. Thus we see that focusing is crucial to maintain the semantic interpretation of (neutral) sequents. In an unfocused calculus, several of these steps could have partial overlaps, making such a semantic interpretation inordinately complicated. We do not know of any encoding



of the π calculus that can provide such interpretations in unfocused sequents without changing the underlying logic. In CLF [6] the logic is extended with explicit monadic staging, and this enables a form of adequacy [6]; however, the encoding is considerably more complex because processes and sums cannot be fully lifted and must instead be specified in terms of a lifting computation. Adequacy is then obtained via a permutative equivalence over the lifting operation. Other encodings of π calculi in linear logic, such as [17] and [3], concentrate on the easier asynchronous fragment and lack adequacy proofs anyhow.

► **Theorem 18** (completeness). *If $E \vdash P \xrightarrow{r} Q$, then the following canonical sequent is derivable.*

$$\llbracket E \rrbracket_e, \text{rates}, \text{inter}@l; \uparrow \text{act}@s, \lambda P\}@l; \cdot \Longrightarrow \cdot; (\llbracket Q \rrbracket_p \text{ at } l) \otimes \text{act}@s \cdot r.$$

Proof. By structural induction of the derivation of $E \vdash P \xrightarrow{r} Q$. Every interaction rule of $S\pi$ is implementable as an admissible inference rule for canonical sequents. For **cong**, we appeal to thm. 14. ◀

Completeness is a testament to the expressivity of the logic – all executions of $S\pi$ are also expressible in HyLL. However, we also require the opposite (soundness) direction: that every canonical sequent encodes a possible $S\pi$ trace. The proof hinges on the following canonicity lemma.

► **Lemma 19** (canonical derivations). *In a derivation for a canonical sequent, the derived inference rules for **inter** are of one of the two following forms (conclusions and premises canonical).*

$$\frac{\llbracket E \rrbracket_e, \text{rates}, \text{inter}@l; \uparrow \text{act}@s, \lambda P\}@l; \cdot \Longrightarrow \cdot; (\llbracket P \rrbracket_p \text{ at } l) \otimes \text{act}@s}{\llbracket E \rrbracket_e, \text{rates}, \text{inter}@l; \uparrow \text{act}@s \cdot r, \lambda Q\}@l; \cdot \Longrightarrow \cdot; (\llbracket R \rrbracket_p \text{ at } l) \otimes \text{act}@t}}{\llbracket E \rrbracket_e, \text{rates}, \text{inter}@l; \uparrow \text{act}@s, \lambda P\}@l; \cdot \Longrightarrow \cdot; (\llbracket R \rrbracket_p \text{ at } l) \otimes \text{act}@t}}$$

where: either $E \vdash P \xrightarrow{r} Q$, or $E \vdash P \equiv Q$ with $r = l$.

Proof. This is a formal statement of the phenomenon observed earlier in the example (fig. 4): $\llbracket R \rrbracket_p \otimes \text{act}$ cannot be focused on the right unless $P \equiv R$, in which case the derivation ends with no more foci on **inter**. If not, the only elements available for focus are **inter** and one of the congruence rules $\llbracket E \rrbracket_e$ in the unrestricted context. In the former case, the derived rule consumes the $\uparrow \text{act}@s$, and by the time **act** is produced again, its world has advanced to $s \cdot r$. In the latter case, the definition of a top level X_n in $\lambda P\}@l$ is (un)folded (without advancing the world). The proof proceeds by induction on the structure of P . ◀

Lemma 19 is a strong statement about HyLL derivations using this encoding: every partial derivation using the derived inference rules represents a prefix of an $S\pi$ trace. This is sometimes referred to as *full adequacy*, to distinguish it from adequacy proofs that require complete derivations [25]. The structure of focused derivations is crucial because it allows us to close branches early (using **init**). It is impossible to perform a similar analysis on unfocused proofs for this encoding; both the encoding and the framework will need further features to implement a form of staging [6, Chapter 3].

► **Corollary 20** (soundness).

If $\llbracket E \rrbracket_e, \text{rates}, \text{inter}@l; \uparrow \text{act}@l, \lambda P\}@l; \cdot \Longrightarrow \cdot; (\llbracket Q \rrbracket_p \text{ at } l) \otimes \text{act}@r$ is derivable, then $E \vdash P \xrightarrow{r}^ Q$.*

Proof. Directly from lem. 19. ◀



4.2 Stochastic Correctness with respect to simulation

So far the $\text{HyLL}(\mathcal{P})$ encoding of $S\pi$ represents any $S\pi$ trace symbolically. However, not every symbolic trace of an $S\pi$ process can be produced according to the operational semantics of $S\pi$, which is traditionally given by a simulator. This is the main difference between HyLL (and $S\pi$) and the approach of CSL [2], where the truth of a proposition is evaluated against a CTMC, which is why equivalence in CSL is identical to CTMC bisimulation [12]. In this section we sketch how the execution could be used directly on the canonical sequents to produce only correct traces (proofs). The proposal in this section should be seen by analogy to the execution model of $S\pi$ simulators such as SPiM [27], although we do not use the Gillespie algorithm.

The main problem of simulation is determining which of several competing enabled actions in a canonical sequent to select as the “next” action from the *race condition* of the actions enabled in the sequent. Because of the focusing restriction, these enabled actions are easy to compute. Each element of $\langle P \rangle$ is of the form $\text{se1} \multimap \llbracket M \rrbracket_s$, so the enabled actions in that element are given precisely by the topmost occurrences of \uparrow in $\llbracket M \rrbracket_s$. Because none of the sums can have any restricted channels (they have all been removed in the active decomposition of the process earlier), the rates of all the channels will be found in the *rates* component of the canonical sequent.

The effective rate of a channel x is related to its inherent rate by scaling by a factor proportional to the *activity* on the channel, as defined in [27]. Note that this definition is on the *rate constants* of exponential distributions, not the rates themselves. The distribution of the minimum of a list of random variables with exponential distribution is itself an exponential distribution whose rate constant is the sum of those of the individual variables. Each individual transition on a channel is then weighted by the contribution of its rate to this sum. The choice of the transition to select is just the ordinary logical non-determinism. Note that the rounds of the algorithm do not have an associated *delay* element as in [27]; instead, we compute (symbolically) a distribution over the delays of a sequence of actions.

Because stochastic correctness is not necessary for the main adequacy result in the previous subsection, we leave the details of simulation to future work.

5 Related Work

Logically, the HyLL sequent calculus is a variant of labelled deduction, a very broad topic not elaborated on here. The combination of linear logic with labelled deduction isn’t new to this work. In the η -logic [13] the constraint domain is intervals of time, and the rules of the logic generate constraint inequalities as a side-effect; however its sole aim is the representation of proof-carrying authentication, and it does not deal with genericity or focusing. The main feature of η not in HyLL is a separate constraint context that gives new constrained propositions. HyLL is also related to the Hybrid Logical Framework (HLF) [32] which captures linear logic itself as a labelled form of intuitionistic logic. Encoding constrained π calculi directly in HLF would be an interesting exercise: we would combine the encoding of linear logic with the constraints of the process calculus. Because HLF is a very weak logic with a proof theory based on natural deduction, it is not clear whether (and in what forms) an adequacy result in HyLL can be transferred to HLF.

Temporal logics such as CSL and PCTL [19] are popular for logical reasoning on temporal properties of transition systems with probabilities. In such logics, truth is defined in terms of correctness with respect to a constrained forcing relation on the constraint algebra. In CSL and PCTL states are formal entities (names) labeled with atomic propositions. Formulae are interpreted on algebraic structures that are discrete (in PCTL) or continuous (in CSL) time Markov chains. Transitions between states are viewed as couples of states labeled with a probability (the probability of the transition), which is defined as a function from $S \times S$ into $[0, 1]$, where S is the set of states.



While such logics have been very successful in practice with efficient tools, the proof theory of these logics is very complex. Indeed, such modal logics generally cannot be formulated in the sequent calculus, and therefore lack cut-elimination and focusing. In contrast, HyLL has a very traditional proof theoretic pedigree, but lacks such a close correspondence between logical and algebraic equivalence. Probably the most well known and relevant stochastic formalism not already discussed is that of stochastic Petri-nets [22], which have a number of sophisticated model checking tools, including the PRISM framework [21]. Recent advances in proof theory suggest that the benefits of model checking can be obtained without sacrificing proofs and proof search [4].

6 Conclusion and Future Work

We have presented HyLL, a hybrid extension of intuitionistic linear logic with a simple notion of situated truth, a traditional sequent calculus with cut-elimination and focusing, and a modular and instantiable constraint system (set of worlds) that can be directly manipulated using hybrid connectives. We have proposed three instances of HyLL (i.e. three particular instances of the set of worlds): one modelling temporal constraints and the others modelling probabilistic or stochastic (continuous time Markov processes) constraints. We have shown how to obtain representationally adequate encodings of constrained transition systems, such as the synchronous stochastic π -calculus in a suitable instance of HyLL.

Several instantiations of HyLL besides the ones in this paper seem interesting. For example, we can already use disjunction (\oplus) to explain disjunctive states, but it is also possible to obtain a more extensional branching by treating the worlds as points in an arbitrary partially-ordered set instead of a monoid. Another possibility is to consider lists of worlds instead of individual worlds – this would allow defining periodic availability of a resource, such as one being produced by an oscillating process. The most interesting domain is that of discrete probabilities: here the underlying semantics is given by discrete time Markov chains instead of CTMCs, which are often better suited for symbolic evaluation [38].

The logic we have provided so far is a logical framework well suited *to represent* constrained transition systems. The design of a logical framework *for* (i.e. to reason about) constrained transition systems is left for future work -and might be envisioned by using a two-levels logical framework such as the Abella system.

An important open question is whether a general logic such as HyLL can serve as a framework for specialized logics such as CSL and PCTL. A related question is what benefit linearity truly provides for such logics – linearity is obviously crucial for encoding process calculi that are inherently stateful, but CSL requires no such notion of single consumption of resources.

In the κ -calculus, reactions in a biological system are modeled as reductions on graphs with certain state annotations. It appears (though this has not been formalized) that the κ -calculus can be embedded in HyLL even more naturally than $S\pi$, because a solution—a multiset of chemical products—is simply a tensor of all the internal states of the binding sites together with the formed bonds. One important innovation of κ is the ability to extract semantically meaningful “stories” from simulations. We believe that HyLL provides a natural formal language for such stories.

We became interested in the problem of encoding stochastic reasoning in a resource aware logic because we were looking for the logical essence of biochemical reactions. What we envision for the domain of “biological computation” is a resource-aware stochastic or probabilistic λ -calculus that has HyLL propositions as (behavioral) types.



Acknowledgements

This work was partially supported by the Information Society Technologies programme of the European Commission, Future and Emerging Technologies under the IST-2005-015905 MOBIUS project, and by the European TYPES project.

References

- 1 Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *J. of Logic and Computation*, 2(3):297–347, 1992.
- 2 A. Aziz, K. Sanwal, V. Singhal, and R. Brayton. Model checking continuous time Markov chains. *ACM Transactions on Computational Logic*, 1(1):162–170, 2000.
- 3 David Baelde. Logique linéaire et algèbre de processus. Technical report, INRIA Futurs, LIX and ENS, 2005.
- 4 David Baelde, Andrew Gacek, Dale Miller, Gopalan Nadathur, and Alwen Tiu. The Bedwyr system for model checking over syntactic expressions. In F. Pfenning, editor, *21th Conf. on Automated Deduction (CADE)*, number 4603 in LNAI, pages 391–397, New York, 2007. Springer.
- 5 Torben Braüner and Valeria de Paiva. Intuitionistic hybrid logic. *Journal of Applied Logic*, 4:231–255, 2006.
- 6 Iliano Cervesato, Frank Pfenning, David Walker, and Kevin Watkins. A concurrent logical framework II: Examples and applications. Technical Report CMU-CS-02-102, Carnegie Mellon University, 2003. Revised, May 2003.
- 7 Nathalie Chabrier-Rivier, François Fages, and Sylvain Soliman. The biochemical abstract machine BIOCHAM. In *International Workshop on Computational Methods in Systems Biology (CMSB-2)*, LNCS. Springer, 2004.
- 8 Bor-Yuh Evan Chang, Kaustuv Chaudhuri, and Frank Pfenning. A judgmental analysis of linear logic. Technical Report CMU-CS-03-131R, Carnegie Mellon University, December 2003.
- 9 Kaustuv Chaudhuri and Joëlle Despeyroux. A logic for constrained process calculi with applications to molecular biology. Technical report, INRIA, 2010. Available at <http://hal.inria.fr/inria-00402942>.
- 10 Kaustuv Chaudhuri, Frank Pfenning, and Greg Price. A logical characterization of forward and backward chaining in the inverse method. *J. of Automated Reasoning*, 40(2-3):133–177, March 2008.
- 11 Vincent Danos and Cosimo Laneve. Formal molecular biology. *Theor. Comput. Sci.*, 325(1):69–110, 2004.
- 12 Josée Desharmais and Prakash Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *Journal of Logic and Algebraic Programming*, 56:99–115, 2003.
- 13 Henry DeYoung, Deepak Garg, and Frank Pfenning. An authorization logic with explicit time. In *Computer Security Foundations Symposium (CSF-21)*, pages 133–145. IEEE Computer Society, 2008.
- 14 E. Allen Emerson. Temporal and modal logic. In *TCS*, pages 995–1072. Elsevier, 1995.
- 15 Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes; Characterization and Convergence*. Wiley series in Probability and Mathematical Statistics. Wiley-interscience, 1986.
- 16 Dominique Foata and Aimé Fuchs. *Calcul des probabilités, cours exercices et problèmes corrigés*. Dunod, 2e édition, 2003.
- 17 Deepak Garg and Frank Pfenning. Type-directed concurrency. In Martín Abadi and Luca de Alfaro, editors, *16th International Conference on Concurrency Theory (CONCUR)*, volume 3653 of LNCS, pages 6–20. Springer, 2005.
- 18 Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- 19 H. Hansson and B. Jonsson. A logic for reasoning about time and probability. *Formal Aspects of Computing*, 6, 1994.



- 20 Johan Anthony Willem Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, University of California, Los Angeles, 1968.
- 21 M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic symbolic model checking using PRISM: a hybrid approach. *International Journal of Software Tools for Technology Transfer*, 6(2), 2004.
- 22 M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modelling with Generalised Stochastic Petri Nets*. Wiley Series in Parallel Computing. Wiley and Sons, 1995.
- 23 Dale Miller. The π -calculus as a theory in linear logic: Preliminary results. In E. Lamma and P. Mello, editors, *3rd Workshop on Extensions to Logic Programming*, number 660 in LNCS, pages 242–265, Bologna, Italy, 1993. Springer.
- 24 Robin Milner. *Communicating and Mobile Systems : The π -Calculus*. Cambridge University Press, New York, NY, USA, 1999.
- 25 Vivek Nigam and Dale Miller. Focusing in linear meta-logic. In *Proceedings of IJCAR: International Joint Conference on Automated Reasoning*, volume 5195 of *LNAI*, pages 507–522. Springer, 2008.
- 26 Frank Pfenning and Conal Elliott. Higher-order abstract syntax. In *Proceedings of the ACM-SIGPLAN Conference on Programming Language Design and Implementation*, pages 199–208. ACM Press, June 1988.
- 27 Andrew Phillips and Luca Cardelli. A correct abstract machine for the stochastic pi-calculus. *Concurrent Models in Molecular Biology*, August 2004.
- 28 Andrew Phillips and Luca Cardelli. A correct abstract machine for the stochastic pi-calculus. In *Proceedings of BioConcur'04*, ENTCS, 2004.
- 29 Andrew Phillips, Luca Cardelli, and Giuseppe Castagna. A graphical representation for biological processes in the stochastic pi-calculus. *Transactions on Computational Systems Biology VII*, pages 123–152, 2006.
- 30 Corrado Priami. Stochastic spi-calculus. *The computer journal*, 38 (7):578–589, 1995.
- 31 Arthur N. Prior. *Time and Modality*. Oxford: Clarendon Press, 1957.
- 32 Jason Reed. Hybridizing a logical framework. In *International Workshop on Hybrid Logic (HyLo)*, Seattle, USA, August 2006.
- 33 A. Regev, W. Silverman, and E. Shapiro. Representation and simulation of biochemical processes using the π -calculus and process algebra. In L. Hunter R. B. Altman, A. K. Dunker and T. E. Klein, editors, *Pacific Symposium on Biocomputing*, volume 6, pages 459–470, Singapore, 2001. World Scientific Press.
- 34 L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes and Martingales*, volume 1: Foundations. Cambridge Mathematical Library, 2nd edition, 2000.
- 35 Uluç Saranlı and Frank Pfenning. Using constrained intuitionistic linear logic for hybrid robotic planning problems. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3705–3710. IEEE, 2007.
- 36 Alex Simpson. *The Proof Theory and Semantics of Intuitionistic Modal Logic*. PhD thesis, University of Edinburgh, 1994.
- 37 Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A concurrent logical framework I: Judgments and properties. Technical Report CMU-CS-02-101, Carnegie Mellon University, 2003. Revised, May 2003.
- 38 Peng Wu, Catuscia Palamidessi, and Huimin Lin. Symbolic bisimulations for probabilistic systems. In *QEST'07*, pages 179–188. IEEE Computer Society, 2007.

