

Lang-n-Send Extended: Sending Regular Expressions to Monitors

Matteo Cimini

University of Massachusetts Lowell
Lowell, MA, USA

matteo_cimini@uml.edu

In prior work, Cimini has presented LANG-N-SEND, a π -calculus with language definitions.

In this paper, we present an extension of this calculus called LANG-N-SENDtm. First, we revise LANG-N-SEND to work with transition system specifications rather than its language specifications. This revision allows the use of negative premises in deduction rules. Next, we extend LANG-N-SEND with monitors and with the ability of sending and receiving regular expressions, which then can be used in the context of larger regular expressions to monitor the execution of programs.

We present a reduction semantics for LANG-N-SENDtm, and we offer examples that demonstrate the scenarios that our calculus captures.

1 Introduction

As the field of software language engineering advances [16], it is increasingly easier for programmers to quickly define and deploy their own programming languages. Cimini has presented in [14] a π -calculus called LANG-N-SEND that accommodates “language-oriented” concurrent scenarios. Processes of LANG-N-SEND can define languages with a syntax for structural operational semantics (SOS), and use these languages to execute programs. Processes can also send and receive fragments of operational semantics through channels. An example of the type of scenarios that LANG-N-SEND enables is the following, which is a simplified version of the first example in [14]. A server stores a library of operational semantics definitions. This server can send these fragments of semantics through a channel upon requests. At a certain point, a client defines a program and a language with which it intends to execute the program. The client notices that the program that it is about to execute is safety-critical. Therefore, the client would like to execute the program in the context of the disrupt operator [6]. The client, then, requests the fragment of operational semantics that defines the disrupt operator from the server, and receives it through a channel. The client augments its language with the disrupt operator, and executes the program using this new language. In particular, it uses the disrupt operator to also specify the error code to be executed in case of a disruption.

The crux of LANG-N-SEND consists of two operators: Program executions $(\mathcal{L}, trace) \triangleright_x program$ and `isInTrace`, used as follows:

$$(\mathcal{L}, trace) \triangleright_x program \parallel x(trace).isInTrace(a, trace) \Rightarrow P ; Q$$

where \mathcal{L} is a language definition and x is a channel. $(\mathcal{L}, trace) \triangleright_x program$ uses the operational semantics of \mathcal{L} to prove a transition from *program*. The evaluation of *program* proceeds this way one step at a time. Each transition is labelled, and LANG-N-SEND accumulates the execution trace in *trace*. When *program* terminates, the final trace is sent over the channel x . The process on the right of the parallel operator receives the trace, and analyzes it with `isInTrace(a, trace) $\Rightarrow P ; Q$` . This process

checks that the label a is one of the labels in $trace$, and continues as P in such a case. Otherwise, it continues as Q .

This paper addresses two limitations of LANG-N-SEND.

First Addition: from Higher-order Logic Programs to Transition System Specifications [14] provides a syntax for language definitions \mathcal{L} . This syntax has been specifically devised to represent operational semantics. The semantics of \mathcal{L} is based on higher-order logic programming as realized with hereditary Harrop formulae [22]: \mathcal{L} is compiled into a higher-order logic program \mathcal{P} , and LANG-N-SEND computes the steps of *program* using the provability relation \models of higher-order logic programming [22], i.e., $\mathcal{P} \models (\longrightarrow \text{label program program}')$, for some *program'* and *label*.

However, higher-order logic programs of [22], and therefore \models , do not contemplate the use of negation. This prevents LANG-N-SEND from defining languages with operators that use negative premises.

In this paper, we revise LANG-N-SEND to adopt transition system specifications (TSSs) [8], a well-known and widely used formalism for SOS specifications. Our motivation for adopting TSSs is that they include negative premises with a well-established semantics [19]. This is a consequential addition: It has been shown that negative premises are actually necessary to express some operators such as the priority operator [3]. That is, SOS with negative premises is strictly more expressive than SOS without them. Other examples of operators that use negative premises include timed operators [20], and some formulations of sequential composition, to name a few.

Our salient challenge here consists in incorporating TSSs into LANG-N-SEND. Fortunately, the design of LANG-N-SEND is rather modular insofar this aspect is concerned, and we simply can use the definitions of TSSs from the literature [8, 18, 19] in lieu of the syntax for language definitions \mathcal{L} .

Second Addition: Online Monitors and Communication of Regular Expressions Monitoring is a runtime verification technique that is based on executing a program and observing its behavior. Its goal is to establish whether such execution satisfies or violates a correctness property (see [9] for a survey on the subject). There are two types of monitoring: *offline* monitoring and *online* monitoring. Offline monitoring executes the program and records its execution trace. The trace is then analyzed after the execution terminates. Conversely, online monitoring performs its analysis alongside the execution of the program. That is, an online monitor acts after each step of the execution, and analyzes the trace that has been generated up to that point.

LANG-N-SEND includes offline monitoring with `isInTrace`. This operation only checks whether an action has occurred. However, this is an inexpressive form of trace analysis. For example, suppose that we were to check whether a process performs valid actions on a file. We would need to see whether read/write operations are performed *after* opening the file. We would also need to check whether the file is being closed at the end. Simply checking whether an action has occurred with `isInTrace` is not enough, as we need to express more refined temporal properties.

Clearly, `isInTrace` is insufficient for most scenarios. Therefore, this paper extends LANG-N-SEND with a more powerful way of analyzing traces. Specifically, we augment LANG-N-SEND with regular expressions, and we add the ability of checking whether a trace satisfies or violates a regular expression. Our salient challenge here consists in equipping LANG-N-SEND with appropriate linguistic features for the monitoring of regular expressions. We do that as follows.

LANG-N-SEND does *not* feature online monitoring. We then model it with an extended form for program executions:

$$(\mathcal{T}, trace) >_x \text{ program with monitors } m_1 m_2 \cdots m_n$$

where \mathcal{T} is a TSS and m_1, m_2, \dots , and m_n are online monitors. Each m_i carries the regular expression to be checked during the execution of *program*, as well as a process to be executed in case such regular expression is violated.

We also have extended the offline monitoring capabilities of LANG-N-SEND by replacing `isInTrace` with the following operation:

$$\text{verifyThis}(\text{trace}, \text{regex}) ? P : Q$$

This operation checks that the regular expression *regex* validates *trace*. If the check is successful then we continue with *P*, otherwise we continue with *Q*. (Our `verifyThis` operator is not limited to checking acceptance for regular expressions, though we postpone this discussion to Section 3.)

LANG-N-SEND is tailored to express dynamic scenarios where language fragments are sent and received, and where processes are instructed to execute programs received from other processes. In this dynamic context, it is natural to also receive, from external processes, the regular expressions to monitor. We have therefore extended LANG-N-SEND with the ability of sending and receiving regular expressions through channels.

We call this new calculus LANG-N-SENDtm (as in “plus monitoring”). Notice that, in this paper, our goal is to address the operations described above with suitable linguistic features. We argue in Section 6 that a calculus that is smaller than the one that we show, and that can encode our operations, may exist, but we leave exploring that research venue as future work.

Contributions We present a reduction semantics for LANG-N-SENDtm in Section 4. To demonstrate the type of scenarios that LANG-N-SENDtm captures, we provide the following examples in Section 5:

- **Negative premise (Example 1).** A server receives the semantics of the parallel operator from another process, which decides whether parallel processes are allowed to spend idle time or whether they must run with maximal progress.
- **Offline monitoring (Example 2).** A server receives programs from clients, executes them to the end, and checks that the programs have used files correctly (open before read/write operations, and close at the end). The server does so by checking that the final trace is accepted by an appropriate regular expression.
- **Online monitoring and sending/receiving of regular expressions (Example 3).** This example refines Example 2. Programs can perform a privileged action on files as long as they respect a correct sequence of actions. This sequence of actions changes every day, and is provided by an external process. The server receives this sequence as a regular expression, and uses it to install an online monitor for the execution of programs.

We believe that LANG-N-SENDtm provides a suitable formalism to express these and similar scenarios. The paper is organized as follows. Section 2 provides the definition of TSSs from the literature. Section 3 presents the syntax of LANG-N-SENDtm. Section 4 presents a reduction semantics for LANG-N-SENDtm. Section 5 demonstrates our calculus with the examples described above. Section 6 offers a discussion of selected aspects such as deadlocks, implementation ideas, and ideas for simplifying our calculus. Section 7 discusses related work, and Section 8 concludes the paper.

2 Preliminaries: Transition System Specifications

We recall the definitions for transition system specifications from [8, 19].

Definition 1 (Signatures and Terms). A signature Σ is a pair (F, ar) where F is a set of function symbols, and the function $ar : F \rightarrow \mathbb{N}$ determines the arity of the functions in F . Given a signature $\Sigma = (F, ar)$, $T(\Sigma)$ is the set of terms of the signature Σ , and is defined as the minimal set satisfying the following: (We use the symbol t for terms).

- $V \subseteq T(\Sigma)$, where V is a set of variables,
- if $t_1, \dots, t_n \in T(\Sigma)$, $f \in F$, and $ar(f) = n$ then $f(t_1, \dots, t_n) \in T(\Sigma)$.

We define Σ_\emptyset as the *empty signature* with $\Sigma_\emptyset \triangleq (\{\}, \{\})$, that is, both F and ar are empty sets.

Definition 2 (Transition System Specifications (TSS)). A transition system specification \mathcal{T} is a triple (Σ, L, D) , where Σ is a signature, L is a set of labels, and D is a set of deduction rules. We use the symbol λ for labels. Deduction rules are formed with formulae in the way that we describe below. A positive formula is of the form $t \xrightarrow{\lambda} t'$, and a negative formula is of the form $t \not\xrightarrow{\lambda}$. A formula f is either a positive formula or a negative formula. Deduction rules are of the form (H, f) , where H is a set of formulae called premises of the rule, and f is a positive formula called conclusion of the rule. We write a deduction rule (H, f) as $\frac{H}{f}$.

The notion of derivability of formulae for TSSs with negative premises is from [18]. As this definition is standard and we do not use any of its machinery, we do not redefine it, but we write $(\Sigma, L, D) \vdash f$ when the formula f is derived from the TSS (Σ, L, D) according to the semantics of [18].

The following definitions from [19] define the componentwise union of two TSSs.

Definition 3 (Union of Signatures). Given two signatures $\Sigma_1 = (F_1, ar_1)$ and $\Sigma_2 = (F_2, ar_2)$ such that $f \in F_1 \cap F_2 \Rightarrow ar_1(f) = ar_2(f)$, we have

$$\Sigma_1 \oplus \Sigma_2 = (F_1 \cup F_2, ar'), \text{ with } ar'(f) = \begin{cases} ar_1(f), & f \in F_1 \\ ar_2(f), & \text{otherwise} \end{cases}$$

Definition 4 (Union of TSSs). Given two TSSs (Σ_1, L_1, D_1) and (Σ_2, L_2, D_2) such that $\Sigma_1 \oplus \Sigma_2$ is defined, we define $(\Sigma_1, L_1, D_1) \oplus (\Sigma_2, L_2, D_2) = (\Sigma_1 \oplus \Sigma_2, L_1 \cup L_2, D_1 \cup D_2)$.

As an example, we define a TSS for a subset of CCS with inaction nil , a unary operator for each action a of a finite set Act , and the parallel operator \parallel . (As usual, Act also contains complement actions which can be denoted as \bar{a} for any action a .) We call this subset *partialCCS*. The set of variables V of the TSS of *partialCCS* ranges over p, q , and so on. We define *partialCCS* as follows.

$D \triangleq$

$$\left\{ \begin{array}{l} a.p \xrightarrow{a} p, \quad \frac{p \xrightarrow{a} p'}{p \parallel q \xrightarrow{a} p' \parallel q}, \quad \frac{q \xrightarrow{a} q'}{p \parallel q \xrightarrow{a} p \parallel q'}, \\ \\ \frac{p \xrightarrow{\tau} p'}{p \parallel q \xrightarrow{\tau} p' \parallel q}, \quad \frac{q \xrightarrow{\tau} q'}{p \parallel q \xrightarrow{\tau} p \parallel q'}, \\ \\ \frac{p \xrightarrow{a} p' \quad q \xrightarrow{\bar{a}} q'}{p \parallel q \xrightarrow{\tau} p' \parallel q'}, \quad \frac{p \xrightarrow{\bar{a}} p' \quad q \xrightarrow{a} q'}{p \parallel q \xrightarrow{\tau} p' \parallel q'} \end{array} \right\}$$

partialCCS $\triangleq ((\{nil, \parallel\} \cup Act, ar), Act \cup \{\tau\}, D)$, where ar assigns the arity 0 to nil , the arity 2 to \parallel , and the arity 1 to every element of Act .

3 Syntax of LANG-N-SEND^{+m}

The syntax of LANG-N-SEND^{+m} is defined as follows. We assume a set of channels x, y, z , and so on. We assume that this set and the sets F s, L s, and V s of TSSs (see Definition 1 and 2) are pairwise disjoint. (Recall that \mathcal{T} denotes a TSS, and t is a term. We also use the notation $\tilde{\cdot}$ for finite sequences.)

Language Builder	ℓ	$::=$	$\mathcal{T} \mid \ell \text{ union } \ell$
Regular Expression	\mathcal{E}	$::=$	$\lambda \mid \varepsilon \mid \mathcal{E} \cdot \mathcal{E} \mid \mathcal{E} \mid \mathcal{E} \mid \mathcal{E}^*$
Trace as Reg. Exp.	\mathcal{E}_{tr}	$::=$	$\lambda \mid \varepsilon \mid \mathcal{E}_{tr} \cdot \mathcal{E}_{tr}$
Transmittable	e	$::=$	x
(language builders)			$\mid \mathcal{T} \mid e \text{ union } e$
(reg. exp.)			$\mid \lambda \mid \varepsilon \mid e \cdot e \mid e \mid e \mid e^*$
(terms)			$\mid t$
Monitor	m	$::=$	$e \Rightarrow P$
Process	P, Q	$::=$	$\mathbf{0} \mid x(y).P \mid \bar{x}(e).P \mid P \parallel Q \mid P + Q \mid \nu x.P \mid !P$
(online monitoring)			$\mid (e, \mathcal{E}_{tr}) >_x e \text{ with monitors } \tilde{m}$
(offline monitoring)			$\mid \text{verifyThis}(e, e) ? P : Q$
(checking labels)			$\mid \text{labels}(\tilde{\lambda}, e) ? P : Q$

Language builder expressions ℓ evaluate to TSSs \mathcal{T} . We can combine two TSSs with union, which performs the union operation that we have seen in Section 2.

LANG-N-SEND^{+m} executes programs and keeps track of their execution trace. We analyze these traces with regular expressions over the set of labels L as the alphabet. The grammar of regular expressions \mathcal{E} is standard (with λ s as atomic symbols). To recall: ε is the empty string, we use an explicit concatenation operator \cdot (though literature often uses juxtaposition), \mid is the alternation operator, and \mathcal{E}^* is the Kleene closure of \mathcal{E} . As usual, the semantics of a regular expression \mathcal{E} is a set of strings. We denote this set with $\llbracket \mathcal{E} \rrbracket$. The semantics $\llbracket \mathcal{E} \rrbracket$ of regular expressions is standard and we omit it here.

Traces are finite strings of labels $\lambda_1 \lambda_2 \dots \lambda_n$. We represent traces with regular expressions of the form $\lambda_1 \cdot \lambda_2 \dots \cdot \lambda_n$, i.e., a concatenation of labels. Therefore, we have traces \mathcal{E}_{tr} as a special case of regular expressions. The semantics $\llbracket \mathcal{E}_{tr} \rrbracket$ of a trace \mathcal{E}_{tr} is a singleton set with one string.

LANG-N-SEND^{+m} can send and receive *transmittable expressions* e through channels. Transmittable expressions are channels, language builder expressions, regular expressions, and terms. Similarly to the π -calculus, channel names x, y , and so on, are binding variables for channels. Additionally, LANG-N-SEND^{+m} uses channel names as binding variables for language builder expressions, regular expressions, and terms, as well. Transmittable expressions can be language builder expressions that contain variables such as $\mathcal{T} \text{ union } x$, where x will be substituted after a communication takes place. Similarly, transmittable expressions can be regular expressions that contain variables such as $\mathcal{E} \cdot x$, where x will be substituted later. Notice that LANG-N-SEND^{+m} processes are such that expressions like $\mathcal{T} \text{ union } x$ and $\mathcal{E} \cdot x$ will have x already substituted when we reach the moment where these expressions are used.

LANG-N-SEND^{+m} contains the processes of the π -calculus, except that the output prefix sends transmittable expressions. Furthermore, LANG-N-SEND^{+m} contains the following processes.

$(e_1, \mathcal{E}_{tr}) >_x e_2 \text{ with monitors } \tilde{m}$ is a *program execution* with online monitors \tilde{m} . This process evaluates e_1 to a TSS \mathcal{T} . Here, e_2 is a term t when this process is activated. (We offer some remarks on type errors in Section 6.) The term t is the program to be executed. This process executes the program t according to the semantics of \mathcal{T} . To do so, we use the derivability of formulae of TSSs to derive a transition from t . Program executions evaluate t one step at a time. Each of these transitions has a label,

and we concatenate these labels in \mathcal{E}_{tr} . We assume that every program execution starts with the empty string ε . Therefore, \mathcal{E}_{tr} is the trace of the execution up to a certain point. After each transition, we check that the current trace satisfies all the monitors \tilde{m} . Each monitor m contains a regular expression and a process P . If the regular expression does not validate the trace then the whole program execution is discarded and P is executed instead. If there are multiple monitors that are not satisfied, LANG-N-SENDtm non-deterministically executes the process of one of them. (We purposely under-specify this part. Actual implementations may fix a selection method, for example based on the order in which monitors appear.)

When the execution of t terminates, the trace \mathcal{E}_{tr} is sent over the channel x .

LANG-N-SENDtm accommodates offline monitors, as well, which analyze the trace of the whole execution after t terminates. We do so in the following way. As just described, the trace \mathcal{E}_{tr} can be received over the channel x . Afterwards, it can be used with the process $\text{verifyThis}(\mathcal{E}_{tr}, \mathcal{E}) ? P : Q$, which behaves as P if the regular expression \mathcal{E} validates the trace \mathcal{E}_{tr} , and behaves as Q otherwise. More specifically, our verifyThis operator works in a slightly more general form: $\text{verifyThis}(e_1, e_2) ? P : Q$, where e_1 and e_2 are regular expressions \mathcal{E}_1 and \mathcal{E}_2 when this process is activated. Notice that \mathcal{E}_1 is not necessarily some trace \mathcal{E}_{tr} . verifyThis checks whether \mathcal{E}_2 subsumes \mathcal{E}_1 , i.e., $\llbracket \mathcal{E}_1 \rrbracket \subseteq \llbracket \mathcal{E}_2 \rrbracket$ ¹. We offer this general form as a convenience to programmers. For example, a process may already be planning to run an online monitor with \mathcal{E}_2 , and may receive \mathcal{E}_1 from another process with instructions to monitor it, as well. This process can execute $\text{verifyThis}(\mathcal{E}_1, \mathcal{E}_2) ? P : Q$ and program P to run a monitor with \mathcal{E}_2 only, as it subsumes \mathcal{E}_1 , as in

$$\begin{aligned} & \text{verifyThis}(\mathcal{E}_1, \mathcal{E}_2) ? (\mathcal{T}, \varepsilon) \rangle_x t \text{ with monitors } \mathcal{E}_2 \Rightarrow \overline{\text{monitor-fail}}.0 \\ & \quad : \\ & \quad (\mathcal{T}, \varepsilon) \rangle_x t \text{ with monitors } \mathcal{E}_1 \Rightarrow \overline{\text{monitor-fail}}.0 \\ & \quad \quad \mathcal{E}_2 \Rightarrow \overline{\text{monitor-fail}}.0 \end{aligned}$$

Notice that when verifyThis is used with a trace as in $\text{verifyThis}(\mathcal{E}_{tr}, \mathcal{E}) ? P : Q$, then $\llbracket \mathcal{E}_{tr} \rrbracket$ is a singleton set with a string and $\llbracket \mathcal{E}_{tr} \rrbracket \subseteq \llbracket \mathcal{E} \rrbracket$ holds whenever that string is in $\llbracket \mathcal{E} \rrbracket$.

A process $\text{labels}(\tilde{\lambda}, e) ? P : Q$, where e evaluates to a TSS \mathcal{T} , checks whether the set of labels of \mathcal{T} is a subset of the labels $\tilde{\lambda}$. In such a case, the process behaves as P , otherwise it behaves as Q . This operation is useful to check, before executing programs, that a TSS works with the expected actions.

4 A Reduction Semantics for LANG-N-SENDtm

Figure 1 shows the reduction semantics of LANG-N-SENDtm in two parts. The first part of Figure 1, that is above the horizontal line, contains the standard definition of the structural congruence \equiv of the π -calculus, and includes the reduction rules of the π -calculus that are also part of the semantics of LANG-N-SENDtm [23]. The second part of Figure 1, that is below the horizontal line, contains the rest of the reduction semantics.

The main reduction relation is \longrightarrow . As in [14], this relation makes use of two auxiliary relations: $\longrightarrow_{\text{lan}}$ evaluates language builder expressions ℓ into TSSs, and $\longrightarrow_{\text{exe}}$ handles program executions.

Rule (COMM) realizes the communication of transmittable expressions. Substitution $P\{e/y\}$ substitutes the free occurrences of y in P with e . This substitution is capture-avoiding, its definition is straightforward, and therefore we do not show it. Notice that LANG-N-SENDtm adopts a call-by-name

¹The inclusion problem is decidable for regular expressions [21].

Reduction Semantics

$$\boxed{P \equiv P, P \longrightarrow P, \ell \longrightarrow_{\text{lan}} \ell, P \longrightarrow_{\text{exe}} P}$$

$$\begin{array}{c} P \parallel \mathbf{0} \equiv P \quad P \parallel Q \equiv Q \parallel P \quad (P \parallel Q) \parallel R \equiv P \parallel (Q \parallel R) \\ P + \mathbf{0} \equiv P \quad P + Q \equiv Q + P \quad (P + Q) + R \equiv P + (Q + R) \quad !P \equiv P \parallel !P \\ \nu x. \mathbf{0} \equiv \mathbf{0} \quad \nu x. \nu y. P \equiv \nu y. \nu x. P \quad \nu x. (P \parallel Q) \equiv \nu x. P \parallel Q, \text{ if } x \text{ is not a free name of } Q \\ \frac{P_1 \longrightarrow P'_1}{P_1 + P_2 \longrightarrow P'_1} \quad \frac{P_1 \longrightarrow P'_1}{P_1 \parallel P_2 \longrightarrow P'_1 \parallel P_2} \quad \frac{P \longrightarrow P'}{\nu x. P \longrightarrow \nu x. P'} \quad \frac{P \equiv P' \quad P' \longrightarrow Q' \quad Q' \equiv Q}{P \longrightarrow Q} \end{array}$$

(COMM)

$$x(y).P \parallel \bar{x}(e).Q \longrightarrow P\{e/y\} \parallel Q$$

(EXEC)

$$\frac{(\mathcal{T}, \mathcal{E}_{\text{tr}}) \triangleright_x t \text{ with monitors } \tilde{m} \longrightarrow_{\text{exe}} P}{(\mathcal{T}, \mathcal{E}_{\text{tr}}) \triangleright_x t \text{ with monitors } \tilde{m} \longrightarrow P}$$

(EXEC-CTX)

$$\frac{\ell \longrightarrow_{\text{lan}} \ell'}{(\ell, \mathcal{E}_{\text{tr}}) \triangleright_x t \text{ with monitors } \tilde{m} \longrightarrow (\ell', \mathcal{E}_{\text{tr}}) \triangleright_x t \text{ with monitors } \tilde{m}}$$

(VERIFY-SUCCESS)

$$\frac{[[\mathcal{E}_1]] \subseteq [[\mathcal{E}_2]]}{\text{verifyThis}(\mathcal{E}_1, \mathcal{E}_2) ? P : Q \longrightarrow P}$$

(VERIFY-FAIL)

$$\frac{[[\mathcal{E}_1]] \not\subseteq [[\mathcal{E}_2]]}{\text{verifyThis}(\mathcal{E}_1, \mathcal{E}_2) ? P : Q \longrightarrow Q}$$

(LABELS-SUCCESS)

$$\frac{\mathcal{T} = (\Sigma, L, D) \quad \tilde{\lambda} = \lambda_1, \dots, \lambda_n \quad L \subseteq \{\lambda_1, \dots, \lambda_n\}}{\text{lb}(\tilde{\lambda}, \mathcal{T}) ? P : Q \longrightarrow P}$$

(LABELS-FAIL)

$$\frac{\mathcal{T} = (\Sigma, L, D) \quad \tilde{\lambda} = \lambda_1, \dots, \lambda_n \quad L \not\subseteq \{\lambda_1, \dots, \lambda_n\}}{\text{lb}(\tilde{\lambda}, \mathcal{T}) ? P : Q \longrightarrow Q}$$

(LABELS-CTX)

$$\frac{\ell \longrightarrow_{\text{lan}} \ell'}{\text{lb}(\tilde{\lambda}, \ell) ? P : Q \longrightarrow \text{lb}(\tilde{\lambda}, \ell') ? P : Q}$$

(UNION)

$$\mathcal{T}_1 \text{ union } \mathcal{T}_2 \longrightarrow_{\text{lan}} \mathcal{T}_1 \oplus \mathcal{T}_2$$

(UNION-CTX1)

$$\frac{\ell_1 \longrightarrow_{\text{lan}} \ell'_1}{\ell_1 \text{ union } \ell_2 \longrightarrow_{\text{lan}} \ell'_1 \text{ union } \ell_2}$$

(UNION-CTX2)

$$\frac{\ell_2 \longrightarrow_{\text{lan}} \ell'_2}{\mathcal{T} \text{ union } \ell_2 \longrightarrow_{\text{lan}} \mathcal{T} \text{ union } \ell'_2}$$

(PROGRAM-STEP)

$$\frac{\mathcal{T} \vdash t \xrightarrow{\lambda} t' \quad \tilde{m} \equiv \mathcal{E}_1 \Rightarrow P_1 \cdots \mathcal{E}_n \Rightarrow P_n \quad [[\mathcal{E}_{\text{tr}} \cdot \lambda]] = \{s\} \quad s \in [[\mathcal{E}_i]] \text{ for all } 1 \leq i \leq n}{(\mathcal{T}, \mathcal{E}_{\text{tr}}) \triangleright_x t \text{ with monitors } \tilde{m} \longrightarrow_{\text{exe}} (\mathcal{T}, \mathcal{E}_{\text{tr}} \cdot \lambda) \triangleright_x t' \text{ with monitors } \tilde{m}}$$

(MONITOR-FAIL)

$$\frac{\mathcal{T} \vdash t \xrightarrow{\lambda} t' \quad \tilde{m} \equiv \mathcal{E}_1 \Rightarrow P_1 \cdots \mathcal{E}_n \Rightarrow P_n \quad [[\mathcal{E}_{\text{tr}} \cdot \lambda]] = \{s\} \quad s \notin [[\mathcal{E}_i]] \text{ for some } 1 \leq i \leq n}{(\mathcal{T}, \mathcal{E}_{\text{tr}}) \triangleright_x t \text{ with monitors } \tilde{m} \longrightarrow_{\text{exe}} P_i}$$

(PROGRAM-END)

$$\frac{\mathcal{T} \not\vdash t \xrightarrow{\lambda} t'}{(\mathcal{T}, \mathcal{E}_{\text{tr}}) \triangleright_x t \text{ with monitors } \tilde{m} \longrightarrow_{\text{exe}} !\bar{x}(\mathcal{E}_{\text{tr}}).\mathbf{0}}$$

Figure 1: Reduction semantics of LANG-N-SENDtm. In this figure, lb is short for the labels operator.

$$\begin{aligned}
\text{timeManagementProvider} &\triangleq \\
&!\text{whatTask}(y).(\overline{\text{getTimeManagement}}\langle\text{parallel}\rangle + \overline{\text{getTimeManagement}}\langle\text{parallel-max-progress}\rangle) \\
\text{server} &\triangleq \\
&!(\text{task}_1(x).\overline{\text{whatTask}}\langle\text{task}_1\rangle.\text{getTimeManagement}(l).(\text{almostTPA union } l, \varepsilon)\rangle_x \text{tpa_program}_1 \\
&+ \\
&\text{task}_2(x).\overline{\text{whatTask}}\langle\text{task}_2\rangle.\text{getTimeManagement}(l).(\text{almostTPA union } l, \varepsilon)\rangle_x \text{tpa_program}_2) \\
\text{system} &\triangleq (\text{server} \parallel \text{timeManagementProvider} \parallel \text{client}_1 \parallel \text{client}_2 \dots \parallel \text{client}_n)
\end{aligned}$$

Figure 2: Server decides idle time vs maximal progress (negative premises).

style for transmitting language fragments. In particular, an output prefix $\bar{x}\langle\mathcal{T}_1 \text{ union } \mathcal{T}_2\rangle.P$ transmits the whole expression $\mathcal{T}_1 \text{ union } \mathcal{T}_2$ without evaluating it, as it will be evaluated when it is used. (The evaluation strategy does not affect the type of scenarios that LANG-N-SENDtm strives to capture. We chose call-by-name to uniformly use (COMM) for all transmittable expressions and simplify our calculus.)

Rule (EXEC) handles program executions and simply defers to $\longrightarrow_{\text{exe}}$ -transitions. Rule (EXEC-CTX) evaluates ℓ when it is not a TSS yet.

Rule (VERIFY-SUCCESS) checks whether \mathcal{E}_2 subsumes \mathcal{E}_1 with $\llbracket\mathcal{E}_1\rrbracket \subseteq \llbracket\mathcal{E}_2\rrbracket$. In that case, the process takes a transition to P . Rule (VERIFY-FAIL) fires whenever \mathcal{E}_2 does not subsume \mathcal{E}_1 , and executes Q .

Rule (LABELS-SUCCESS) checks whether the labels of the TSS that is given as second argument are from the set of labels given as first argument. In that case, the process takes a transition to P . Rule (LABELS-FAIL) fires whenever that is not the case, and executes Q . Rule (LABELS-CTX) evaluates ℓ when it is not a TSS yet.

Rule (UNION) performs the union of two TSSs with the operation \oplus defined in Section 2. Rules (UNION-CTX1) and (UNION-CTX2) evaluate the first and second argument of `union`, respectively.

Rule (PROGRAM-STEP) handles program executions $(\mathcal{T}, \mathcal{E}_{tr})\rangle_x t$ with monitors \tilde{m} . We use the derivability relation \vdash of TSSs to check that a formula $t \xrightarrow{\lambda} t'$ is provable for some t' and some label λ . We then check that all the regular expressions of the monitors \tilde{m} validate the trace up to that point, which is \mathcal{E}_{tr} with λ appended. To do so, we first compute the string s of $\mathcal{E}_{tr} \cdot \lambda$ with $\llbracket\mathcal{E}_{tr} \cdot \lambda\rrbracket = \{s\}$. (Recall that the semantics of $\mathcal{E}_{tr} \cdot \lambda$ is a set with one string.) Then we check that s belongs to the semantics of each \mathcal{E}_i of the monitors (with $s \in \llbracket\mathcal{E}_i\rrbracket$).

Rule (MONITOR-FAIL) is similar to (PROGRAM-STEP) except that it fires when there exists a regular expression \mathcal{E}_i that does not validate the current trace. In this case the transition takes a step to the corresponding process P_i specified by the failing monitor. Notice that this transition is non-deterministic when there are multiple regular expressions \mathcal{E}_i that fail.

Rule (PROGRAM-END) detects that a step is not provable for t . Then, the execution of t is terminated. We spawn a replicated output prefix that sends the trace over the channel x . The reason for replicating this output is that there may be multiple processes that are interested in analyzing the trace, as we shall see in our second example of Section 5.

5 Examples

Example 1 (Negative Premises) Our first example makes use of the newly-added feature to use negative premises in the context of processes that communicate languages. In this example, we have a server

that decides whether parallel processes are allowed to spend idle time or whether they must run with maximal progress. We define the TSS of a subset of Hennessy and Regan's Process Algebra for Timed Systems (TPA) [20]. This subset of TPA contains inaction nil , unary operators $a.P$ for each of the actions a of a finite set Act , and the parallel operator \parallel . The transitions of TPA are labeled with actions of Act , the silent action τ , and the label σ for the passing of idle time. The transition $P \xrightarrow{\sigma} P$ means that the process P spends idle time.

We define *almostTPA* to be the subset of TPA just described. However, we omit the rule for the passing of idle time for the parallel operator. We first define the set of rules D_{tpa^-} . We then define *almostTPA* as an extension of *partialCSS* of Section 2. (Recall that Σ_\emptyset is the empty signature defined in Section 2.)

$$D_{tpa^-} = \{a.P \xrightarrow{\sigma} a.P, nil \xrightarrow{\sigma} nil\}.$$

$$almostTPA \triangleq partialCSS \oplus (\Sigma_\emptyset, \{\sigma\}, D_{tpa^-}).$$

We can complete *almostTPA* by including a way for time to pass in the context of the parallel operation. For example, we can add either of the following rules.

$$\begin{array}{c} \text{(PAR-IDLE)} \\ \frac{p \xrightarrow{\sigma} p' \quad q \xrightarrow{\sigma} q'}{p \parallel q \xrightarrow{\sigma} p' \parallel q'} \end{array} \qquad \begin{array}{c} \text{(PAR-MAX)} \\ \frac{p \xrightarrow{\sigma} p' \quad q \xrightarrow{\sigma} q' \quad p \parallel q \not\xrightarrow{\tau}}{p \parallel q \xrightarrow{\sigma} p' \parallel q'} \end{array}$$

(PAR-IDLE) lets the two processes spend idle time, if both processes can. Conversely, (PAR-MAX) implements *maximal progress* and allows idle time to pass only so long that the two processes cannot communicate. (TPA uses (PAR-MAX) in [20].)

We define these two rules in the context of empty TSSs, so that we can conveniently use our union operator to add them to *almostTPA*.

$$parallel \triangleq (\Sigma_\emptyset, \{\}, \{(\text{PAR-IDLE})\})$$

$$parallel-max-progress \triangleq (\Sigma_\emptyset, \{\}, \{(\text{PAR-MAX})\})$$

Figure 2 shows our example. *server* is a server that offers two services, $task_1$ and $task_2$. Upon a request from a client, *server* executes the program $tpa_program_1$ for $task_1$, and $tpa_program_2$ for $task_2$. These are programs of our subset of TPA. *server* has limited computational resources, and executing programs in maximal progress mode is computationally expensive. Therefore, *server* communicates with another process called *timeManagementProvider* through the channel *whatTask*, and sends the name of the service that has been requested. *timeManagementProvider* non-deterministically decides whether *server* should use maximal progress or not (perhaps based on the urgency of the task, as well as other factors). *timeManagementProvider* sends *parallel* or *parallel-max-progress* through the channel *getTimeManagement*. In other words, *timeManagementProvider* decides the semantics of the parallel operator, insofar idle time is concerned, that *server* must use. Then, *server* completes *almostTPA* with this fragment of TSS before executing the program.

Example 2 (Offline Monitoring) Figure 3 shows an example with offline monitoring. Here, *server* is a server that manages files. Clients send programs to *server*. Clients also send the TSSs with which *server* must execute these programs. *server* is capable of receiving TSSs and executing programs with them. However, the only actions that *server* supports are the following actions on files: open, read, write, and close. In other words, clients can define any TSS they wish, and any SOS operator they

$$\begin{aligned}
allowedLabels &\triangleq \text{open, read, write, close} \\
fileProtocol &\triangleq \text{open} \cdot (\text{read} \mid \text{write})^* \cdot \text{close} \\
server &\triangleq !(getProgram(l, w, id, x). \\
&\quad \text{labels}(allowedLabels, l) ? \\
&\quad \quad ((l, \varepsilon) >_x w \parallel x(tr).verifyThis(tr, fileProtocol^*) ? \mathbf{0} : \overline{flagClient}(id)) \\
&\quad ; \\
&\quad \overline{invalid-language}) \\
onlyOneWrite &\triangleq (\text{open} \mid \text{read} \mid \text{close})^* \cdot \text{write} \cdot (\text{open} \mid \text{read} \mid \text{close})^* \\
client_1 &\triangleq \nu x. (\overline{getProgram}(tss, tss_program, id, x) \parallel x(tr).verifyThis(tr, onlyOneWrite) ? P : \mathbf{0}) \\
system &\triangleq (server \parallel client_1 \parallel client_2 \dots \parallel client_n)
\end{aligned}$$

Figure 3: Server checks for the correct use of files (offline monitoring).

wish. Whichever operators they define, however, must compute transitions to open, read, write, and close files only, as these are the only actions that *server* recognizes. When the program terminates, the trace is sent over a channel and is available to both server and clients.

Our example models the scenario in which both client and server are running an offline monitor to analyze the trace of an execution. We describe both sides below.

server receives the language l and the program w through the channel *getProgram*. (To shorten our notation, *getProgram* sends and receives multiple arguments in polyadic style, though this is shorthand for a sequence of unary input and output prefixes.) *server* also receives the id of the client (as a channel name), and a channel x where to send the trace once the execution of w has finished. After receiving these arguments, *server* checks that the set of labels of l is formed with the allowed labels. If this check fails, the server signals an error through the channel *invalid-language*. Otherwise, the server executes w . As there are no online monitors, we simply write $(l, \varepsilon) >_x w$. The server is interested in analyzing the trace of this execution, and so it receives the trace at x and runs an offline monitor with *verifyThis*. The server checks that w has used files correctly, i.e., it has opened a file before read/write operations, and it has closed the file afterwards. The correct use of a file is expressed with the regular expression *fileProtocol*. As w may have used files multiple times, the server uses *verifyThis* to check that the trace is accepted by *fileProtocol*^{*} (with Kleene star). If this check succeeds then the server ends. Otherwise, the server flags the client as an unreliable programmer using the channel *flagClient*.

One of the clients, *client*₁, is also interested in analyzing the trace of an execution. *client*₁ verifies that its program has performed exactly one writing operation. This is expressed with the regular expression *onlyOneWrite*. *client*₁ sends a TSS tss and a program $tss_program$ (whose details are irrelevant) to the server. It also sends its id and a private channel x . Then, it receives the trace at x , and runs an offline monitor with *verifyThis* to check that the trace is accepted by *onlyOneWrite*. If this check succeeds then *client*₁ continues as P . Otherwise, it terminates.

Example 3 (Online Monitoring and Sending/Receiving of Regular Expressions) Figure 4 shows an example of online monitoring in LANG-N-SENDtm, and also illustrates the sending/receiving of regular expressions over channels. This example refines our previous example. Here, *server* additionally admits a privileged action on files, *delete*, which deletes a file. However, programs can perform a *delete*-transition only if they know the “password of the day” provided by the process *passwordMan-*

$$\begin{aligned}
\text{allowedLabels} &\triangleq \text{open, read, write, close, sudo, 0, 1, 2, \dots, 9, delete} \\
\text{ordinary} &\triangleq \text{open} \mid \text{read} \mid \text{write} \mid \text{close} \\
\text{new} &\triangleq \text{sudo} \mid 0 \mid 1 \mid 2 \mid \dots \mid 9 \mid \text{delete} \\
\\
\text{passwordManager} &\triangleq !(\overline{\text{getPasswordOfTheDay}}\langle 3 \cdot 4 \cdot 5 \cdot 6 \rangle) \\
\text{server} &\triangleq !(\text{getProgram}(tss, prg, id). \\
&\quad \text{labels}(\text{allowedLabels}, tss) ? \\
&\quad \text{getPasswordOfTheDay}(rexp). \\
&\quad \text{vx}(((tss, \epsilon) >_x \text{prg with monitors} \\
&\quad \quad (\text{new}^* \cdot \text{fileProtocol} \cdot \text{new}^*)^* \Rightarrow \overline{\text{flagClient}}\langle id \rangle) \\
&\quad \quad (\text{ordinary}^* \cdot (\text{sudo} \cdot \text{rexp} \cdot \text{delete}) \cdot \text{ordinary}^*)^* \Rightarrow \overline{\text{flagClient}}\langle id \rangle) \\
&\quad \quad \parallel x(\text{tr}).\overline{\text{end}}) \\
&\quad : \\
&\quad \overline{\text{invalid-language}}) \\
\\
\text{system} &\triangleq (\text{server} \parallel \text{passwordManager} \parallel \text{client}_1 \parallel \text{client}_2 \dots \parallel \text{client}_n)
\end{aligned}$$

Figure 4: Server receives a regular expression to check valid access to `delete` (online monitoring).

ager. Passwords are numeric. The password of the example in Figure 4 is 3456. Programs must first announce their intention to use the privileged action with a `sudo` action. Then, they must perform the actions that correspond to the digits of the password. In other words, *server* also admits actions 0, 1, 2, ..., 9, where, for example, the action 3 can be interpreted as “sent 3” or “pressed 3”. Programs can perform `delete` after having performed this sequence of actions. In our example, the correct sequence of actions for using `delete` is `sudo, 3, 4, 5, 6, and delete`, in this order.

server receives the language l , the program w , and the client id. (Clients are not interested about the trace in this example, and so they do not send the channel x of the previous example). The server checks that the set of labels of l is formed with the allowed labels. Then, the server receives a regular expression e through the channel `getPasswordOfTheDay`. This represents the fragment of a trace that corresponds to the correct sequence of actions that enables `delete`. The regular expression so received is substituted in lieu of e , as we shall discuss shortly. At this point, the server creates a private channel x and executes the program w giving x as the channel where to receive the final trace. The server also specifies two online monitors for this program execution. The first monitor performs the check on file operations that we have seen in the previous example, except that the check is performed at each step of the execution. Furthermore, the regular expression of the previous example is slightly modified to take into account the new actions of *server*, which may occur before and after `fileProtocol`.

The second monitor checks that `delete` is used properly. The regular expression of this monitor is $(\text{ordinary}^* \cdot (\text{sudo} \cdot 3 \cdot 4 \cdot 5 \cdot 6 \cdot \text{delete}) \cdot \text{ordinary}^*)^*$ after e has been substituted, that is, we check that the correct sequence appears within the other actions. If this check fails then the client is flagged for knowing the wrong password, or not using the correct protocol.

Notice that `sudo · sudo · 3 · 4 · 5 · 6 · delete`, as well as other acceptable sequences, are invalid. We believe that the example sufficiently demonstrates our approach even though our regular expressions could be more refined.

Finally, *server* can detect that all online monitors succeed throughout the execution of w with the input prefix $x(e).\overline{\text{end}}$. This process signals successful termination through the channel `end`.

6 Discussion

Type Errors and Deadlocks The syntax of $\text{LANG-N-SEND}^{\text{tm}}$ does not rule out erroneous uses of its operators. For example, (1) $(\mathcal{E}, \mathcal{E}_{tr}) \rangle_x t$ with monitors \tilde{m} and (2) $\text{verifyThis}(\mathcal{E}_{tr}, \mathcal{T}) ? P : Q$ are processes of $\text{LANG-N-SEND}^{\text{tm}}$. The former attempts to execute a program but a regular expression \mathcal{E} is given in lieu of a TSS. The latter checks the trace \mathcal{E}_{tr} against a TSS rather than a regular expression. Processes with these and similar type errors can deadlock. For example, none of the reduction rules among (EXEC), (EXEC-CTX), (PROGRAM-RUN), (MONITOR-FAIL) and (PROGRAM-END) of Figure 1 applies to (1) because \mathcal{E} is not a TSS nor a valid language builder expression. Similarly, neither (VERIFY-SUCCESS) nor (VERIFY-FAIL) applies to (2) because \mathcal{T} is not a regular expression.

As future work, we would like to design a type system that rejects this kind of type errors. Such a type system would rule out the type of deadlocks just described. Notice, however, that even if we eliminated deadlocks that are caused by type errors, $\text{LANG-N-SEND}^{\text{tm}}$ processes could deadlock anyway in as much the same way that π -calculus process can deadlock. Another source of deadlocks comes from the fact that the union \oplus is sometimes undefined (see Definition 4 of Section 2).

Implementation Aspects Although we have not implemented $\text{LANG-N-SEND}^{\text{tm}}$, we discuss some of the aspects that must be taken into account when implementing $\text{LANG-N-SEND}^{\text{tm}}$.

An important aspect is the implementation of program executions. An implementation must be able to take a term and a TSS, and compute the transitions of the term based on the TSS. This may be done by implementing the execution of inference rule systems, including the unification of formulae with the conclusions of inference rules, as well as all the other aspects of inference systems. A good alternative is to outsource this task and make the implementation interact with an external tool. For example, the work in [11, 13] presents a functional language with language definitions that are based on logic programming, and executes programs using λ -prolog (see [12] for implementation details). $\text{LANG-N-SEND}^{\text{tm}}$ would need to execute TSSs, instead, and tools such as [25] could be used to derive the transitions of terms and provide them to the implementation.

Similarly, implementors can either implement their own checkers for regular expressions or interact with one of the several tools for regular expressions.

Passing languages through channels is an interesting aspect of the implementation, as well. Implementors may use the datatypes that are available in the programming language being used to implement $\text{LANG-N-SEND}^{\text{tm}}$, and use them to represent TSSs. They then can use the network primitives of the programming language to send/receive values of this datatype, that is, to send/receive fragments of TSSs.

Comparison with $\text{HO}\pi$ and Session Types The higher-order π -calculus ($\text{HO}\pi$) can send/receive processes and execute them, which reminds of the core capability of our calculus. It is therefore natural to ask whether $\text{HO}\pi$ could be used in lieu of $\text{LANG-N-SEND}^{\text{tm}}$. If all the TSSs that are intended to be used in a $\text{LANG-N-SEND}^{\text{tm}}$ process are the $\text{HO}\pi$ or subsets of it then $\text{HO}\pi$ captures the sending/receiving/executing part of $\text{LANG-N-SEND}^{\text{tm}}$. Afterwards, $\text{HO}\pi$ would need to record the traces of the process being executed, which may not be difficult to do, and would need to represent regular expressions and solve their acceptance problem, which may be harder to do and may take some serious encoding effort. If the TSS that we intend to run is, say, an arbitrary language with an operator f , then we first need to prove that $\text{HO}\pi$ operations faithfully encode f . We see two issues here: First, $\text{HO}\pi$ may not be able to express every operator that can be defined in a TSS with negative premises. Second, translating these operators may be possible only when we know their semantics beforehand. However, new operators and

new deduction rules are passed around in LANG-N-SEND^{+m} , and via non-deterministic communications. Therefore, we may only know at run-time, right before executing a program, what TSS has been built.

The monitoring capabilities of LANG-N-SEND^{+m} remind of session types, as well. In this latter approach, a process is assigned a type that can express patterns of communications. Session types can express that the trace of the executed process should respect some regular expression. The main difference between monitoring and session types is that the latter are checked statically (compile-time), and reject processes that violate a specified protocol before running them. Monitors, instead, execute the programs to detect violations. It would be interesting to adopt session types rather than monitors in LANG-N-SEND^{+m} . However, we see great challenges with such endeavor as we would need to 1) synthesize a session type system for an arbitrary TSS defined by the user, and 2) translate regular expressions as session types of this type system. The task of 1) is also made harder by the fact that, as we pointed out above, TSSs can be non-deterministically built at run-time.

Ideas for Simplifying LANG-N-SEND^{+m} We have designed LANG-N-SEND^{+m} with linguistic features in mind that one-to-one match the operations that we intended to add. There may be room for simplifying LANG-N-SEND^{+m} . For example, $\text{labels}(\tilde{\lambda}, \mathcal{S}) ? P : Q$ can be seen as a special case of verifyThis where the set of labels of \mathcal{S} is translated into a concatenation of labels, like a trace, and where $\tilde{\lambda}$ is translated into a regular expression that accepts all the permutations of the labels of $\tilde{\lambda}$ and their substrings. In turn, offline monitoring could be performed with online monitoring, in principle, for example by sending the trace at the end of an execution to an online monitor that executes a program that replays the trace, while checking for a property. This suggests that there may be a calculus with online monitoring only, and perhaps auxiliary operations, where labels and offline monitoring can be encoded as macros. We leave working out the details of such a calculus as future work.

7 Related Work

[14] is a direct related work of LANG-N-SEND^{+m} . We have discussed the differences between this paper and [14] in Section 1 (Introduction). Also, [14] transmits language fragments in call-by-value style while LANG-N-SEND^{+m} does so in call-by-name style. [14] employs a union operation on languages but this operation is not standard, and it has been specifically devised to apply to the syntax for languages of [14]. Instead, we use the standard \oplus operator on TSSs. The examples in this paper showcase the added expressiveness of LANG-N-SEND^{+m} over the prior work done in [14].

There are several works on runtime monitoring (see [9] for a survey). Our paper does *not* offer a new monitoring technique. On the contrary, we have taken an existing approach, i.e., monitoring with regular expressions [7, 10], and have integrated it into a calculus with processes that communicate programs, traces, and languages.

Temporal logics such as LTL, and linear fragments of HML, HML with recursion, and the modal μ -calculus (see [1]) can be used in lieu of regular expressions to state properties on traces. Our first draft of LANG-N-SEND^{+m} had LTL formulae in their finite traces interpretation LTL_f [26] in lieu of regular expressions. However, regular expressions are more expressive than LTL_f [26], and their formalism is more widely known and used, so we simply chose to use that instead. We could use an expressive logic but the goal of this paper is not to use the most powerful logic. Rather, we wanted to demonstrate the type of scenarios that LANG-N-SEND^{+m} enables with a sufficiently expressive formalism that is also easy to read. In this light, we believe that regular expressions may be a suitable choice. As future work, we do plan to integrate more expressive logics in LANG-N-SEND^{+m} and make more sophisticated examples.

Works such as [2, 17] provide general frameworks for monitoring that is based on operational semantics. Some of these works have also been implemented in the `detectEr` tool chain [4, 5]. It would be interesting to integrate these frameworks in `LANG-N-SENDtm` in future work.

We have offered some discussion about the higher-order π -calculus and session types in the previous section. Generally speaking, the realm of process calculi is tremendously vast and diverse: Process calculi have been augmented with sophisticated operators, and have been applied to a plethora of domains. We are not aware, however, of process calculi where processes send fragments of TSSs or regular expressions through channels.

8 Conclusion

We have presented `LANG-N-SENDtm`, an extension of `LANG-N-SEND` of [14]. As `LANG-N-SEND`, our calculus is tailored to model language-oriented scenarios where processes send and receive language fragments. `LANG-N-SENDtm` also addresses two limitations of [14]. We use transition system specifications rather than `LANG-N-SEND`'s specification syntax that is based on higher-order logic programming. This allows `LANG-N-SENDtm` to define SOS specifications with negative premises. Furthermore, we have added monitoring capabilities based on regular expressions. Processes of `LANG-N-SENDtm` can also send and receive regular expressions.

We have presented a reduction semantics for `LANG-N-SENDtm`, and we have provided examples that demonstrate the type of programming scenarios that `LANG-N-SENDtm` captures. We believe that the three examples in [14], together with the examples in this paper, provide a good idea of the potential uses of `LANG-N-SENDtm`.

As future work, we would like to design a type system for `LANG-N-SENDtm`. We also would like to extend `LANG-N-SENDtm`. We plan to integrate methods for the automated analysis of language definitions such as that proposed in [15]. We also plan to add more operations on TSSs, such as removing rules, and renaming operators. We observe that the difference between (PAR-IDLE) and (PAR-MAX) is the single premise $P \parallel Q \xrightarrow{\tau}$. It would be interesting to make `LANG-N-SENDtm` more fine-grained in its capabilities to communicate fragments of TSSs. We plan to add the ability of sending/receiving premises which then can be added to rules. With such an addition, our first example could simply work with TPA with (PAR-IDLE), and add the negative premise above on the fly to make it become (PAR-MAX).

TSSs do not include syntax for binding (and neither does [14]). We plan to integrate the nominal transition systems of Parrow et al. [24] in our calculus, which can accommodate binders in SOS specifications. With such an addition, we would like to make examples with the π -calculus and its variants as TSSs that can be sent/received.

We also would like to investigate a suitable notion of bisimilarity equivalence for `LANG-N-SENDtm`.

Acknowledgements We are thankful to the anonymous reviewers for their suggestions, which helped improve our paper.

References

- [1] Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir & Karoliina Lehtinen (2019): *Adventures in monitorability: from branching to linear time and back again*. *Proc. ACM Program. Lang.* 3(POPL), pp. 52:1–52:29, doi:10.1145/3290365.

- [2] Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir & Karoliina Lehtinen (2019): *An Operational Guide to Monitorability*. In: *Software Engineering and Formal Methods - 17th International Conference, SEFM 2019, Oslo, Norway, September 18-20, 2019, Proceedings*, pp. 433–453, doi:10.1007/978-3-030-30446-1_23.
- [3] Luca Aceto & Anna Ingólfssdóttir (2008): *On the expressibility of priority*. *Information Processing Letters* 109(1), pp. 83–85, doi:10.1016/j.ipl.2008.09.002.
- [4] Duncan P. Attard, Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir & Karoliina Lehtinen (2021): *Better Late Than Never or: Verifying Asynchronous Components at Runtime*. In Kirstin Peters & Tim A. C. Willemse, editors: *Formal Techniques for Distributed Objects, Components, and Systems - 41st IFIP WG 6.1 International Conference, FORTE 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Proceedings, Lecture Notes in Computer Science 12719*, Springer, pp. 207–225, doi:10.1007/978-3-030-78089-0_14.
- [5] Duncan P. Attard & Adrian Francalanza (2016): *A Monitoring Tool for a Branching-Time Logic*. In Yliès Falcone & César Sánchez, editors: *Runtime Verification - 16th International Conference, RV 2016, Madrid, Spain, September 23-30, 2016, Proceedings, Lecture Notes in Computer Science 10012*, Springer, pp. 473–481, doi:10.1007/978-3-319-46982-9_31.
- [6] Jos C. M. Baeten & Jan A. Bergstra (2000): *Mode transfer in process algebra*. *Computing Science Reports* 00-01, Technische Universiteit Eindhoven. Available at <https://research.tue.nl/en/publications/mode-transfer-in-process-algebra>.
- [7] Howard Barringer, Allen Goldberg, Klaus Havelund & Koushik Sen (2004): *Rule-Based Runtime Verification*. In Bernhard Steffen & Giorgio Levi, editors: *Verification, Model Checking, and Abstract Interpretation*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 44–57, doi:10.1007/978-3-540-24622-0_5.
- [8] Roland Bol & Jan F. Groote (1996): *The Meaning of Negative Premises in Transition System Specifications*. *Journal of the ACM* 43(5), pp. 863–914, doi:10.1145/234752.234756.
- [9] Ian Cassar, Adrian Francalanza, Luca Aceto & Anna Ingólfssdóttir (2017): *A Survey of Runtime Monitoring Instrumentation Techniques*. In: *Proceedings Second International Workshop on Pre- and Post-Deployment Verification Techniques, PrePost@iFM 2017, Torino, Italy, 19th September 2017*, pp. 15–28, doi:10.4204/EPTCS.254.2.
- [10] Feng Chen & Grigore Roşu (2005): *Java-MOP: A Monitoring Oriented Programming Environment for Java*. In: *Proceedings of the 11th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS’05*, Springer-Verlag, Berlin, Heidelberg, pp. 546–550, doi:10.1007/978-3-540-31980-1_36.
- [11] Matteo Cimini (2018): *Languages as First-Class Citizens (Vision Paper)*. In: *Proceedings of the 11th ACM SIGPLAN International Conference on Software Language Engineering, SLE 2018*, Association for Computing Machinery, New York, NY, USA, pp. 65–69, doi:10.1145/3276604.3276983.
- [12] Matteo Cimini (2020): *On the Effectiveness of Higher-Order Logic Programming in Language-Oriented Programming*. In: *Functional and Logic Programming - 15th International Symposium, FLOPS 2020, Akita, Japan, September 14-16, 2020, Proceedings*, pp. 106–123, doi:10.1007/978-3-030-59025-3_7.
- [13] Matteo Cimini (2021): *A Calculus for Multi-language Operational Semantics*. In: *Software Verification - 13th International Conference, VSTTE 2021, New Haven, CT, USA, October 18-19, 2021, and 14th International Workshop, NSV 2021, Los Angeles, CA, USA, July 18-19, 2021, Revised Selected Papers*, pp. 25–42, doi:10.1007/978-3-030-95561-8_3.
- [14] Matteo Cimini (2022): *Lang-n-Send: Processes That Send Languages*. In: *Proceedings of the 13th International Workshop on Programming Language Approaches to Concurrency- and Communication-cEntric Software (PLACES 2022)*, 356, Open Publishing Association, pp. 46–56, doi:10.4204/eptcs.356.5.
- [15] Matteo Cimini, Dale Miller & Jeremy G. Siek (2020): *Extrinsically typed operational semantics for functional languages*. In: *Proceedings of the 13th ACM SIGPLAN International Conference on Software Lan-*

- guage Engineering, SLE 2020, Virtual Event, USA, November 16-17, 2020, pp. 108–125, doi:10.1145/3426425.3426936.
- [16] Sebastian Erdweg, Tijs van der Storm, Markus Völter, Meinte Boersma, Remi Bosman, William R. Cook, Albert Gerritsen, Angelo Hulshout, Steven Kelly, Alex Loh, Gabriël D. P. Konat, Pedro J. Molina, Martin Palatnik, Risto Pohjonen, Eugen Schindler, Klemens Schindler, Riccardo Solmi, Vlad A. Vergu, Eelco Visser, Kevin van der Vlist, Guido H. Wachsmuth & Jimi van der Woning (2013): *The State of the Art in Language Workbenches*. In Martin Erwig, Richard F. Paige & Eric Wyk, editors: *Software Language Engineering, Lecture Notes in Computer Science 8225*, Springer, pp. 197–217, doi:10.1007/978-3-319-02654-1_11.
- [17] Adrian Francalanza (2016): *A Theory of Monitors - (Extended Abstract)*. In Bart Jacobs & Christof Löding, editors: *Foundations of Software Science and Computation Structures - 19th International Conference, FOSACS 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings, Lecture Notes in Computer Science 9634*, Springer, pp. 145–161, doi:10.1007/978-3-662-49630-5_9.
- [18] Rob J. van Glabbeek (2004): *The meaning of negative premises in transition system specifications II*. *Journal of Logical and Algebraic Methods in Programming* 60-61, pp. 229–258, doi:10.1016/j.jlap.2004.03.007.
- [19] Jan F. Groote & Frits Vaandrager (1992): *Structured Operational Semantics and Bisimulation as a Congruence*. *Information and Computation* 100(2), pp. 202–260, doi:10.1016/0890-5401(92)90013-6.
- [20] Matthew Hennessy & Tim Regan (1995): *A Process Algebra for Timed Systems*. *Information and Computation* 117(2), pp. 221–239, doi:10.1006/inco.1995.1041.
- [21] Albert R. Meyer & Larry J. Stockmeyer (1972): *The Equivalence Problem for Regular Expressions with Squaring Requires Exponential Space*. In: *Proceedings of the 13th Annual Symposium on Switching and Automata Theory, SWAT'72*, IEEE Computer Society, USA, pp. 125–129, doi:10.1109/SWAT.1972.29.
- [22] Dale Miller & Gopalan Nadathur (2012): *Programming with Higher-Order Logic*, 1st edition. Cambridge University Press, New York, NY, USA, doi:10.1017/CB09781139021326.
- [23] Robin Milner, Joachim Parrow & David Walker (1992): *A calculus of mobile processes, I*. *Information and Computation* 100(1), pp. 1–40, doi:10.1016/0890-5401(92)90008-4.
- [24] Joachim Parrow, Johannes Borgström, Lars-Henrik Eriksson, Ramunas Gutkovas & Tjark Weber (2015): *Modal Logics for Nominal Transition Systems*. In: *26th International Conference on Concurrency Theory, CONCUR 2015, Madrid, Spain, Sept. 1-4, 2015*, pp. 198–211, doi:10.4230/LIPIcs.CONCUR.2015.198.
- [25] Alceste Scalas & Massimo Bartoletti (2015): *The LTS WorkBench*. In: *Proceedings 8th Interaction and Concurrency Experience, ICE 2015, Grenoble, France, 4-5th June 2015*, pp. 86–98, doi:10.4204/EPTCS.189.8.
- [26] Thomas Wilke (1999): *Classifying Discrete Temporal Properties*. In: *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science, STACS'99*, Springer-Verlag, Berlin, Heidelberg, pp. 32–46, doi:10.1007/3-540-49116-3_3.