

Behavioural Equivalences for Dynamic Web Data

Sergio Maffeis Philippa Gardner
Department of Computing, Imperial College London
`{maffeis,pg}@doc.ic.ac.uk`

Abstract

Peer-to-peer systems, exchanging dynamic documents through Web services, are a simple and effective platform for data integration on the internet. Dynamic documents can contain both data and references to external sources in the form of links, calls to web services, or coordination scripts. XML standards, and industrial platforms for web services, provide the technological basis for building such systems, and process algebras are a promising tool for studying and understanding their formal properties.

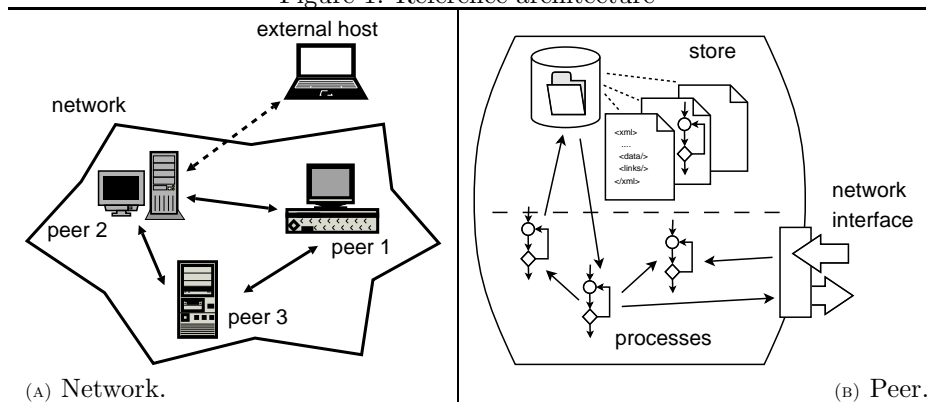
Core $Xd\pi$ is the explicitly located version of $Xd\pi$, a process calculus designed for reasoning about dynamic Web data, based on explicit repositories of higher-order semistructured data and π -calculus-like processes which can communicate with each other, query and update the local repository, or migrate to other peers to continue execution.

We study behavioural equivalences for Core $Xd\pi$ processes. To help with the proofs, which require a costly property of closure under contexts, we define a coinductive relation (called *domain bisimilarity*) which does not quantify over contexts and which entails process equivalence. Its definition is non-standard, because scripts are part of the values, and process equivalences are sensitive to the set of locations constituting the network. We apply our process equivalence to study some communication patterns used by servers in distributed query systems, and we propose a new pattern involving mobile code.

1 Introduction

The World Wide Web is a global network, used in daily activities to find information, communicate ideas, conduct business and carry out distributed computations. In order to fully exploit the potential of this massive network, there is a need for scalable mechanisms to organize and manipulate the available information. Peer-to-peer architectures help to deal with the issue of scalability, and technologies such as XML and Web services facilitate the development of distributed applications. XML [38] is a standardized data model, used to represent uniformly documents containing tagged information not adhering to a fixed structure. Web services [40] are Web sites which are designed to be used by applications rather than humans. Web service inter-operability is facilitated

Figure 1: Reference architecture



by the use of XML for data representation and of related standards for service invocation, description and discovery (SOAP, WSDL, UDDI [41, 39, 36]).

Data integration on the Web constitutes a challenging application for these technologies, because of the extreme heterogeneity of data sources involved, and the complexity of communication patterns which can arise. For example, translating a declarative request for networked data into a low-level execution plan may involve recursively invoking other declarative requests on different Web sites. Inspired by this problem, the $Xd\pi$ calculus [11] studies peer-to-peer architectures for exchanging Web data, schematically represented in Figure 1. Each *network* is composed by a variable number of interconnected peers, all sharing a similar internal structure, and each one identified by a unique name (Figure 1(A)). Peers share a common messaging protocol where the name of a peer is assumed to coincide with its network address: at this level of abstraction there are no restrictions to connectivity due to network domains or firewalls. Networks are *open* in the sense that it is always possible to add new peers or learn dynamically about their existence, and external hosts may participate in the data exchange too, typically playing a limited role. Each peer (schematized in Figure 1(B)) acts both as a provider and consumer of information. It contains a data repository, an internal working space where processes carry out local computations, and a network interface providing remote communication and services to other peers. Processes can communicate locally with each other, query and update the local repository and, when the architecture supports mobility, can migrate to other peers to continue execution. Repositories present to the processes a semi-structured view of their data. Data contains enough meta-information about its own structure to make it possible writing expressive queries.

Typically, data is not completely static. It may contain references to other data and services, in the form of URLs and queries, or *scripts*. A script is some code describing a process which can be interpreted by the working space to add dynamic content to documents. We refer to such data as *dynamic Web*

data. The World Wide Web itself is a very general example of architecture for dynamic Web data. Servers use the HTTP protocol to interact with each other, either requesting or providing information. HTML pages can contain hyperlinks, forms and client-side scripts, which provide dynamic behaviour. Web clients running a browser can be considered as the “external hosts” which participate to a smaller degree in the exchange of information, by mostly consuming rather than providing data.

A more specific example comes from the database world. The Active XML [33, 3] system for data integration (AXML for short) is based on networks of peers each containing a repository of documents and a set of service definitions. AXML service definitions typically consist of queries and updates on the local repository, but in general can consist of arbitrary Web services, providing an interface to hosts external to the AXML system. AXML documents are XML documents which can include special tags representing calls to services on other peers. The parameters to these service calls can be local queries (path expressions) or AXML data, hence service calls can be nested. Documents containing service calls are called *intensional documents*, and *materialization* is the process of invoking a service call and pasting its results in the original document. One interesting source of flexibility in AXML is the choice of when to materialize service calls. It can be done periodically, or when the data containing the call is fetched from the repository, or when it is returned to the client. Similarly, if a service call appears as a parameter to another service call, it can be materialized before calling the service or it can be passed on to it as an intensional parameter.

Besides Web browsers and AXML, a large class of other Web applications (such as file-sharing programs, personal Web portals, online bibliographic databases, etc.) can be seen as instances of the reference architecture given above, each with its own particular features and restrictions. The problems that these architectures have to address, in order to be practically useful, are varied. Firstly, it is well-known that interaction between concurrent processes is difficult to regulate. In the case of Web services, this problem is complicated by the difficulty in maintaining state across different Web service invocations, and requires the study of orchestration techniques.¹ Secondly, a major concern for systems dealing with dynamic Web data is security. Depending on the application domain, it may be crucial to have control for example over data integrity, confidentiality, or access control. The formal study of security properties needs to be grounded on a rigorous model of these architectures, and process algebraic techniques are particularly suited for the task, as they have already been successfully used to study concurrent, distributed and mobile systems, and analyze their formal properties.

¹By Web service orchestration, we mean a coordination infrastructure which allows modular applications to invoke different Web services and combine their results.

1.1 The $Xd\pi$ calculus

The $Xd\pi$ -calculus was defined with the aim of reasoning about dynamic Web data. $Xd\pi$ terms represent networks of peers where each peer consists of an XML data repository and a working space where π -like processes are allowed to run. We regard processes as agents with a simple set of functionalities: they communicate with each other, query and update the local repository, and migrate to other peers to continue execution. Process descriptions, in the form of scripts, can be included in documents and can be executed by other processes. The definition of $Xd\pi$ is parametric with respect to the choice of a specific language of query and update expressions.

Consider again the diagram of our reference architecture for dynamic Web data given in Figure 1. $Xd\pi$ models each peer as a *location* with a unique name corresponding to the peer identity (for example its IP address). A whole peer-to-peer system is modelled by the parallel composition of the locations corresponding to its peers, which we call a *network*. For example, the network in (A) could be represented by the term

$$peer_1 [tree_1 \parallel processes_1] \mid \dots \mid peer_n [tree_n \parallel processes_n]$$

The XML data stored at each peer is represented by an ordered, edge-labelled tree.² The choice of using edge-labelled rather than node-labelled trees is merely a matter of style. Following a common practice, we do not represent attributes explicitly, but we model them as edges labelled with the attribute name followed by a leaf containing the attribute value. We also embed pointers and scripts as leaves. In a concrete document, we would also expect them to be represented as attributes. Our results do not depend on these particular representation choices. To keep the model simple, we do not represent data values and XML-specific details such as name-spaces, ids and idrefs. The tree structure, along with scripts and pointers, provides a sufficiently accurate model for our purposes. Figure 2(A) shows a fragment of an XML document containing both a hyperlink and a service call, and Figure 2(B) shows its representation in $Xd\pi$ (the translation of the hyperlink and the service call are explained below).

Hyperlinks have been one of the main features responsible for the success of the Web. We abstract the concept of hyperlink into that of *pointer*, a pair consisting of a location name and a query to identify some data in the tree of the named location. For example, in Figure 2 we have translated the destination of the hyperlink “http://xdpi.net/papers/xdpi.pdf” into a pointer of the form *query@location* using the host name “xdpi.net” as the location name, and the path relative to the host “papers/xdpi.pdf” as the query. Pointers are declarative references which can be interpreted uniformly across locations. A pointer does not specify what to do with the data denoted by the query, but typically a process will read the location name and the query from a pointer in order to retrieve some data necessary to continue its execution. Clicking on

²Semi-structured data models are often unordered [1], in contrast with the ordered trees of XML documents. In previous work [11], we considered unordered trees, but here we prefer the ordered model, which has a straightforward correspondence to the textual syntax.

Figure 2: Representing XML in Core $Xd\pi$

```
<data>
  <a href = "http://xdpi.net/papers/xdpi.pdf"> Download </a>
  <call> xdpi.net/getRefs(bibtex,l) </call>
</data>
```

(A) A hyperlink and a service call in XHTML.

```
data[
  a[href[ papers/xdpi.pdf@xdpi.net ]|Download[]]
  call[ <go xdpi.net . getRefs(bibtex[,l)> ]
]
```

(B) The translation to $Xd\pi$.

an HTML hyperlink is a simple example, where the browser process reads the contents of the `href` attribute, retrieves the referenced data, and displays it in the browser window. We assume that the same query makes sense on different locations because we are assuming that all the peers export their data in the same semi-structured format.

Current Web technology is familiar with the use of scripts to provide Web pages with dynamic behaviour. Similarly, we propose to use scripts as a generalization of embedded service calls in the context of Web data integration. Since our scripts are used also for coordination, they are written in the same process language used to describe processes in the working space. A script is a static piece of code with some parameters. Scripts do not reference the global state except for names of locations and services, which are constants with a uniform meaning across the network. For example, the service call “`xdpi.net/getRefs(bibtex,l)`” embedded in (A) in Figure 2 could be (naively) translated to the script shown in (B), which specifies that a process should go from the local host to the host “`xdpi.net`” and there invoke the service “`getRefs`” with a data parameter “`bibtex`” and a return parameter l . We shall see a more realistic representation of service calls in Section 2.6. Scripts are atomic and cannot be combined together to form other scripts (for example, the parallel composition of two scripts is not defined).³

The working space of each peer is modelled by a parallel composition of processes inside the corresponding location. The interface between the working space and the data store is modelled by a single operation for updating or

³Some languages such as MetaOCaml [35] and TemplateHaskell [34] provide constructs for multi-stage programming, where pieces of code (possibly containing free variables) can be combined together to form bigger programs, which can then be executed. If desired, it is possible to support multi-stage programming in $Xd\pi$, defining an XML-like meta-syntax for scripts and interpreting it explicitly using parsing processes in the working space.

querying the local tree. Communication between locations is modelled through process migration, providing a flexible abstraction to model complex coordination protocols, and communication between processes is modelled by π -calculus communication. For example,

$$l [D_l \parallel \text{go } xdpi . \overline{\text{getRefs}}(\text{bibtex}[], l)] \mid xdpi [D_x \parallel !\text{getRefs}(x, y).[. . .]]$$

represents a network where a script similar to the one described above is run on location l . In parallel, on location $xdpi$, there is a service (represented by the replicated input $!\text{getRefs}(x, y).[. . .]$) ready to accept the request with parameters x and y . After migration, the output $\overline{\text{getRefs}}(\text{bibtex}[], l)$ will be at location $xdpi$, ready to be received by the service input

$$l [D_l \parallel \mathbf{0}] \mid xdpi [D_x \parallel \overline{\text{getRefs}}(\text{bibtex}[], l) \mid !\text{getRefs}(x, y).[. . .]]$$

The service may perform some computation and return some data, in the form of another output process which migrates to the location determined by the second parameter of getRefs (in this case the original location l).

The definition of $Xd\pi$ is minimal, including only the basic operations for asynchronous local communication based on pattern matching, execution of a query-update expression on the local repository, migration, spawning of scripted code and creation of fresh channels. From these, one can derive conditional statements, nondeterministic choices, constructs for parsing and iterating on list-like structures and remote communication in the style of Web services. Using these derived constructs in [18], we used $Xd\pi$ to give a precise semantics to AXML-like behaviour, and propose possible extensions.

1.2 Equivalences

The combination of Web services and scripted processes provides the data engineer with many alternative patterns for exchanging information. A theory of semantic equivalence for processes is therefore useful to show, for example, that some complex data-exchange protocol corresponds to some intuitive behaviour. Motivated by this consideration, we have defined network equivalences for $Xd\pi$ [11], which dictate when two networks can be considered indistinguishable with respect to the properties represented by a specific set of observations, in our case the attempts to interact with the local store.⁴ Network equivalences are parametric with respect to the language used for querying and updating documents (so the generic results are not tied to a particular choice), and can be instantiated to specific cases. Our objective is to define equivalence relations on processes such that, when we place equivalent processes in the same context, we obtain equivalent networks. Since we want to use the equivalences for example

⁴In [18] we considered as observations the shape of the data tree of a location, the presence of output actions in a process, and the attempts to interact with the local store, and we studied the formal relationship between the corresponding equivalences. In this paper, we focus on the latter kind of observation because the resulting equivalence coincides or implies the ones resulting from the other observations.

to optimize the interaction between different locations, we must compare several located processes at the same time, which possibly share some private channel names. Moreover, we want to make sure that the behaviour of the processes is robust with respect to changes in the data stored in each location and to the behaviour of other processes running in parallel. It is not straightforward to carry on the kind of reasoning mentioned above directly on $Xd\pi$ terms, because locations, processes and data are closely inter-twined. Instead, based on the ideas presented in [19], we propose a calculus called Core $Xd\pi$ which serves as an alternative representation of $Xd\pi$, where we locate processes explicitly and separate data from processes.

From $Xd\pi$ to Core $Xd\pi$. Core $Xd\pi$ is tailored to be exactly as expressive as $Xd\pi$ (a proof can be found in [18]), and is suitable for expressing a partial specification of a network by means of located processes running in parallel, possibly sharing private names. A $Xd\pi$ network consists of locations containing trees and processes. In contrast, a Core $Xd\pi$ network consists of a store, containing all the trees indexed by their location information, plus all the processes, augmented with explicit location information associated to every action they perform. For example, the $Xd\pi$ network described earlier is described in Core $Xd\pi$ as

$$\left((\{l \mapsto D_l\}, \{xdpi \mapsto D_x\}), \right. \\ \left. l \cdot \text{go } xdpi . \overline{xdpi} \cdot \text{getRefs}(\text{bibtex}[], l) \mid !xdpi \cdot \text{getRefs}(x, y) . [\dots] \right)$$

The store $(\{l \mapsto D_l\}, \{xdpi \mapsto D_x\})$ maps each location to its data. Each process action is prefixed by the location where the action takes place: for example, $l \cdot \text{go } xdpi . [\dots]$ shows that the migration step towards location $xdpi$ originates from location l ; after migration the process is located at $xdpi$. Although Core $Xd\pi$ is explicitly located, and does not necessarily require a migration operation, we left that as part of the syntax to have a closer correspondence with $Xd\pi$, and to serve as a hook for future security-sensitive checks.

In Section 3.1, we define the correspondent of $Xd\pi$ network equivalence for Core $Xd\pi$, and we define process equivalence as the closure of network equivalence under composition with different stores. This relation is hard to use directly because it requires a costly property of closure under all contexts. Instead, we define a coinductive equivalence relation (called *domain bisimilarity*) which does not quantify over contexts and which entails process equivalence.

The definition of domain bisimilarity is non-standard, due to the fact that scripts (which can appear in data) are part of the values, and process equivalences are sensitive to the set of locations constituting the network. We address these two problems by adapting existing techniques for translating messages containing scripts into ones where each script is replaced by a first-order value [30, 17], and by generalizing the notion of bisimulation to families of relations indexed by sets of locations.

As an application of our techniques, we use bisimilarity to study some communication patterns used by servers in distributed query systems to answer queries from clients. In Core $Xd\pi$, distributed queries take the form of processes which retrieve and combine data from different locations by using remote

communication and local requests. We show that some existing patterns [29] can be combined together obtaining a flexible infrastructure which is provably equivalent to an intuitive specification of the intended behaviour. By exploiting process migration, we also propose a new communication pattern, and we show that it is behaviourally equivalent to a naive, less efficient one.

Our work is one of the first attempts to integrate the study of mobile processes and semi-structured data, and is characterized by its emphasis on dynamic data. It is the first investigation of equivalence properties for (higher-order) data-centric applications based on the Web.

1.3 Related work

$Xd\pi$ was developed independently from the AXML system of Abiteboul *et al.* [33], which we have described above. The **ubQL** distributed query language of Sahuguet and Tannen [27] instead, constituted a source of inspiration for the design of $Xd\pi$. **ubQL** is built by adding process manipulation primitives to any “host” query language. These primitives, inspired by the π -calculus⁵, are used in a *deployment phase* to set up a network of processes which, in a successive *execution phase*, will query local repositories and forward their results to other sites, thus implementing a global query execution plan. **ubQL** processes can deal with streaming data, but there is no support for concurrent execution of query-processes on the same site (so in principle the system may not be able to execute more than one global query at a time). The main influences of **ubQL** on the design of $Xd\pi$ were on the choices of separating the queries from the process primitives, and maintaining independence from a specific query language. Also our examples on distributed query patterns of Section 4 are inspired by **ubQL**. Overall, $Xd\pi$ and **ubQL** have a significantly different focus and are studied using different methodologies. For example, an important part of the work on **ubQL** is the study of algorithms for query installation based on cost estimates, which we do not address, whereas behavioural equivalences are not studied in **ubQL**. Both AXML and **ubQL** are studied from a data-management viewpoint, which our process algebraic techniques could complement nicely. There are many specific issues which are important in databases, such as the use of meta-data to guide the optimization of queries, which we do not study. Instead we give a formal semantics to the distributed interaction between query-processes, arguing about their equivalence and providing a framework on which to base formal studies of security properties.

We now consider work related to the process algebraic approach. To the best of our knowledge, the only work relating the π -calculus with XML which pre-dates ours is the Iota concurrent XML scripting language of Bierman and Sewell [4], used to program Home Area Networks. Iota is a strongly typed functional language with concurrency primitives inspired by the π -calculus. Although the language has a formal semantics, its behavioural theory has not been

⁵The influence born by the π -calculus on **ubQL** can be better appreciated considering the preliminary joint work with Pierce [28].

studied. The programs for Home Area devices written in Iota are designed to run on the same Home Area server, and the communication with physical devices is modelled through input and output on special channels: distribution is not represented explicitly. Moreover, as opposed to $Xd\pi$, the application domain of Home Area Network programming is more control-oriented than data-oriented: there is no explicit representation of stores, which are central to our approach. Brown, Laneve and Meredith [6] have recently defined an (untyped) extension of the π -calculus with native XML datatypes called π Duce. They compare its expressivity to that of the functional language XDuce [16], and also consider a higher-order extension which enables dynamic content in documents. An interesting idea underlying the design of π Duce is that processes and data share a similar tree-like structure, and can inhabit the same semantic universe. The authors show a very simple encoding of an evaluator for the subset of the language without new name generation, into the language itself: the execution of processes represented as nested document elements can be simulated in the language. A similar approach could be taken in $Xd\pi$ to represent scripts as semi-structured data. In the case of dynamic Web data though, it is better to hide the internal structure of processes from the queries, so that one can replace a process by an equivalent one whilst preserving the observable behaviour of the system as a whole. Castagna, De Nicola and Varacca [8] propose $\mathbb{C}\pi$, a π -calculus extended with pattern matching and tuples of values (XML values can be represented through an encoding). The language comes with a very expressive type system featuring intersection and input-output types. The language itself is not distributed and does not include a concept of store. Acciai and Boreale [2] have recently proposed XPi, an extension of the asynchronous π -calculus with code mobility and ML-like pattern matching of structured values. A combination of static and dynamic typing ensures that each channel always exchanges values of the same types, which describe the partial structure of documents. Pattern matching plays a lesser role in $Xd\pi$, although it could be easily extended to the more expressive form adopted in XPi. Query expressions instead, which are separate entities from processes, are the primary means to extract information from XML trees.

2 Core $Xd\pi$

In this Section we introduce the formal definition of the syntax and semantics of Core $Xd\pi$. We begin with trees, data and queries, and then we pass to networks and processes.

Figure 3: Syntax: trees and data

$\mathbf{T}, \mathbf{S} ::=$	tree terms
$\mathbf{E} \mathbf{T}$	branch \mathbf{E} composed with tree \mathbf{T}
\emptyset	empty tree
x	tree variable
$\mathbf{E}, \mathbf{F} ::= \mathbf{a}[V] \mid x$	branch with edge label \mathbf{a} and data \mathbf{V} , or variable
$\mathbf{U}, \mathbf{V} ::=$	data terms
\mathbf{T}	tree \mathbf{T}
$\mathbf{p}@l$	pointer to location l with query \mathbf{p}
$\langle \mathbf{A} \rangle$	script \mathbf{A}
$l ::= l \mid x$	location name or variable
$\mathbf{p} ::= p \mid x$	query (see Definition 2.1) or variable
$\mathbf{A} ::= A \mid x$	script (see Figure 5) or variable
$\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathcal{E}$	(EDGE LABELS)
$l, m \in \mathcal{L}$ (countably infinite)	(LOCATIONS NAMES)
$p, q \in \mathcal{Q}$	(QUERIES)
$x, y, z \in \mathcal{V}$	(VARIABLES)
$\mathbf{E}, \mathbf{F} \in \mathcal{B} \stackrel{\text{def}}{=} \{\mathbf{E} : fv(\mathbf{E}) = \emptyset\}$	(BRANCHES)
$\mathbf{U}, \mathbf{V} \in \mathcal{D} \stackrel{\text{def}}{=} \{\mathbf{V} : fv(\mathbf{V}) = \emptyset\}$	(DATA)
$\mathbf{T}, \mathbf{S} \in \mathcal{T} \stackrel{\text{def}}{=} \{\mathbf{T} : fv(\mathbf{T}) = \emptyset\}$	(TREES)

Function fv is defined in Figure 23.

Notation: $\mathbf{a}[\] \stackrel{\text{def}}{=} \mathbf{a}[\emptyset] \quad E_1 \mid \dots \mid E_n \stackrel{\text{def}}{=} E_1 \mid \dots \mid E_n \mid \emptyset$.

2.1 Trees, data and queries

We represent semi-structured data using ordered labelled trees.⁶ The formal definition is given in Figure 3 (we use italic bold letters for arbitrary terms, which can contain variables (such as \mathbf{T}), and plain italic letters for closed terms (such as T)).

Trees and data. We represent a tree as a \emptyset -terminated list of branches $E_1 \mid \dots \mid E_n \mid \emptyset$ (abbreviated with $E_1 \mid \dots \mid E_n$) which start from the root. Each branch E_i has the form $\mathbf{a}[V]$ and denotes an edge labelled \mathbf{a} leading to a node containing the data V . A data item can be a subtree \mathbf{T} , a pointer $\mathbf{p}@l$ referencing the data selected at location l by query \mathbf{p} (described below), or a script

⁶A hypothetical encoding of trees into processes, which could be interesting in itself, would introduce unnecessary complexity, making it hard to reason directly on trees like we can do, thanks to our direct representation, in the equivalences of Chapter 3.

Figure 4: Syntax: Core $Xd\pi$ networks and contexts

$$\begin{aligned}
 D, B \in \mathcal{S} &\stackrel{\text{def}}{=} \mathcal{L} \rightarrow \mathcal{T} && \text{(STORES)} \\
 N, M \in \mathcal{N} &\stackrel{\text{def}}{=} \{(D, P) : D \in \mathcal{S}, P \in \mathcal{P}, \text{dom}(P) \subseteq \text{dom}(D)\} && \text{(NETWORKS)}
 \end{aligned}$$

Function dom is defined in Figure 22, processes are defined in Figure 5.

$\langle A \rangle$ (described in Section 2.3) which can be executed to collect data or perform coordination tasks. We show an example of a tree containing a script and a pointer:

$$\mathbf{a}[\mathbf{b}[\mathbf{c}[\langle A \rangle] \mathbf{d}[p@l]] \mathbf{e}[]]$$

We use the same identifiers x, y, z, \dots to range over all variables. When necessary, the kind of each variable can be understood by the place where the variable occurs.

Queries. First of all, it is important to clarify that, in this paper, we use the word “queries” to mean expressions used to query *or update* a tree. $Xd\pi$ is parametric on the choice of a particular query-update language chosen, as long as it is a language of expressions which can be evaluated against a tree to obtain a new tree (the result of updating the tree) and some data (the list of trees resulting from querying the tree). The only conditions that we need to impose on such a language are that the application of a substitution to a query must be well-defined and yield a query. The reasons why we need to define substitutions on queries will be clear after describing the semantics of processes in Section 2.3. In Section 3, we will impose additional conditions required to ensure that a query language is also compatible with our definitions of semantic equivalences.

Definition 2.1 (Query Language) *A query language consists of a triple $(\mathcal{Q}, fv, \mathfrak{E})$ where \mathcal{Q} is a set of queries ranged over by p, q, \dots , together with a function $fv : \mathcal{Q} \rightarrow \wp(\mathcal{V})$ giving the free variables of each query, and an evaluation function $\mathfrak{E} : (\mathcal{Q} \times \mathcal{T}) \rightarrow \mathcal{T} \times \text{lists}(\mathcal{D})$, which, given a query and a tree, returns an updated tree and a finite list of results. Additionally, \mathcal{Q} must be closed under substitution.*

Note that in the definition above the evaluation of queries is a partial function. This generality accounts for both the cases of ill-formed queries, which may not have a precise semantics, and Turing-equivalent query languages, which may not terminate. In Section 2.5, we give a concrete query language which will be used in the examples.

2.2 Networks

A Core $Xd\pi$ network represents a peer-to-peer system, where each location corresponds to a peer. Each peer can communicate with any other peer, and has

a unique name. A network is represented by a pair (D, P) where the first component (the *store*) is a finite partial function from location names to trees, and the second component is a process. The formal definition is given in Figure 4.⁷ For example, in the network $(\{l \mapsto T\}, P)$, the term $\{l \mapsto T\}$ says that the store of the peer at location l is the tree T , and the term P represents the processes running on the peer, which contain explicit location information. Interaction between processes and data is always local, as we shall see later from rule (CRED REQUEST) in Figure 7. In Figure 22, we define the function *dom* giving the domain of both networks, stores and processes. By definition, the domain of a network is the domain of the store, and a network is well-formed if the domain of the process is contained in the domain of the store.

2.3 Processes

The formal definition for Core $Xd\pi$ processes is given in Figure 5. We now describe the technical features of each construct.

Communication. The *output* process $\overline{l \cdot c}(\tilde{v})$ denotes a vector of values \tilde{v} waiting to be sent via channel c at location l , the *input* process $l \cdot c(\tilde{\pi}).P$ waits to receive values matching the patterns $\tilde{\pi}$ from an output process via channel c at l , and the *replicated input* is standard.⁸ The well-formedness condition *wf*(P) requires that the continuation on an input (or replicated input) process must be located at the same location where the input is defined. Channel names are partitioned into *private* and *service* channel names. The private channels denote “usual” π -calculus channels, which are typically used for coordination, and which can be kept secret in order to protect a protocol from external interferences. The service channels denote those channels which are used to implement the services which a peer offers to other peers, and which therefore are not meant to be restricted and can be referenced inside scripts.

Pattern matching. Both trees and pointers are data terms which processes need to parse. For this reason, we have added to π -calculus communication a very simple form of pattern matching. Patterns π_1, \dots, π_n are terms containing distinct variables which are instantiated, if pattern matching succeeds, with the values found in the corresponding position in the term to be matched. Our patterns do not include regular or recursive expressions, and we will avoid algorithmic issues by simply requiring the guessing of an appropriate substitution in order for pattern matching to take place. Pattern matching for XML-like data is an active research topic, which is orthogonal to our concerns. We believe that the specialized techniques studied elsewhere can be adapted to our setting. Our processes use patterns to parse data, and queries to query trees. This conceptual separation does not exclude the possibility for the query language to be

⁷The creation of new peers is not an operation which can be performed from within a system, and therefore we do not provide an operation to create new locations. Nevertheless, we will be able to carry on compositional reasoning, hence analyze networks with respect to arbitrary additions of peers.

⁸The communication constructs, which use polyadic synchronization, were inspired by the ϵ - π -calculus [7].

Figure 5: Syntax: Core $\lambda d\pi$ processes

$P, Q, R ::=$	process terms
$\mathbf{0}$	nil process
$P \mid P$	composition of processes
$(\nu c)P$	private channel c with scope P
$\overline{l \cdot c}(\tilde{v})$	at l , output values \tilde{v} on c
$l \cdot c(\tilde{\pi}).P$	at l , input on c of $\tilde{\pi}$, continue with P ($distinct(\tilde{\pi}), fv(\tilde{\pi}) \cap fv(l) = \emptyset$)
$!l \cdot c(\tilde{\pi}).P$	lazy replication of an input process ($distinct(\tilde{\pi}), fv(\tilde{\pi}) \cap fv(l) = \emptyset$)
$l \cdot go \ m.P$	at l , go to m , continue with P
$A \circ \langle l, \tilde{v} \rangle$	at l , run script A with parameters \tilde{v}
$l \cdot req_p \langle c \rangle$	at l , request query p with return channel c
$a, b, c ::= c \mid c \mid x$ private/service channel, or variable $v ::= c \mid l \mid p \mid A \mid E \mid T$ value terms	
$a, b, c \in \mathcal{C}_p$ (countably infinite)	(PRIVATE CHANNELS)
$a, b, c \in \mathcal{C}_s$	(SERVICE CHANNELS)
$v, u \in \mathcal{U} \stackrel{\text{def}}{=} \{v : fv(v) = \emptyset\}$	(VALUES)
$\pi \in \mathcal{K} \stackrel{\text{def}}{=} \mathcal{V} \cup \{V : cval(V) = \emptyset \text{ and } distinct(V)\}$	(PATTERNS)
$P, Q, R \in \mathcal{P} \stackrel{\text{def}}{=} \{P : fv(P) = \emptyset \text{ and } wf(P)\}$	(PROCESSES)
$A \in \mathcal{A} \stackrel{\text{def}}{=} \left\{ (x, \tilde{\pi})P : \begin{array}{l} fn(P) = \emptyset, fv(P) \subseteq fv(x, \tilde{\pi}), \\ distinct(x, \tilde{\pi}), dom(P) = \{x\} \end{array} \right\}$	(SCRIPTS)

We define wf , $distinct$, dom , $cval$ in Figure 22 and fv , fn in Figure 23.

Well formedness ensures that the continuation of a process is properly located (for example $wf(l \cdot c(\tilde{\pi}).P)$ and $wf(m \cdot go \ l.P)$ require that P is located at l).

Notation: $l \cdot P \stackrel{\text{def}}{=} P$ if $dom(P) = \{l\}$; $l \cdot \overline{m \cdot c}(\tilde{v}) \stackrel{\text{def}}{=} l \cdot go \ m \cdot \overline{m \cdot c}(\tilde{v})$.

based on pattern matching itself.

Scripts. A script $(x, \tilde{\pi})P$ represents the code of P parameterized on x , a placeholder for the location where the script is going to be run, and $\tilde{\pi}$, other optional parameters of P . By the side condition on the free variables of scripts, we impose that scripts remain statically defined until they are deployed dynamically by instantiation of their parameters. The application construct $A \circ \langle l, \tilde{v} \rangle$ passes the parameters \tilde{v} to the script A and runs it in the working space of l . Note that application is defined only when the first parameter passed to the

script is a location. Communication in Core $Xd\pi$ is higher-order, in the sense that processes may send scripts over channels, possibly as leaves inside trees.

Migration. Process migration, which we represent explicitly, models communication across locations: the process $m \cdot \mathbf{go} \ l.P$ represents a (higher-order) message from m addressed to l containing a request to run the (closed) code P . When P is an output process, we use the abbreviation $l \cdot \overline{m \cdot c}(\tilde{v})$ for $l \cdot \mathbf{go} \ m \cdot \overline{m \cdot c}(\tilde{v})$. The well-formedness condition $wf(P)$ requires that the continuation process must be correctly located at the destination location. Due to the peer-to-peer nature of our domain, each location is ready to receive and run any incoming code, so we do not need to provide an explicit operation to run a received process. In some cases, it may also be desirable to give control to each location regarding which code to accept and which to refuse. We leave that task to an eventual superimposed security infrastructure. Using an asynchronous form of communication offers a simple way to model failures within the system. The success of a migration step just depends on the existence of location l . In contrast, the migration rules for other mobile calculi (for example $d\pi$ [14]) assume that migration is always possible. Our choice has an important effect on the behavioural equivalences studied in Section 3.1.

Interaction with the store. Core $Xd\pi$ processes access the local tree by using a request operation $l \cdot \mathbf{req}_p(c)$ parametric in a query-update expression p and a channel c . The effect of evaluating expression p is to modify the local tree and to return a list of query results on the specified channel. Research on query and update languages for XML is still very active [25], and an in-depth study goes beyond the scope of this paper. Therefore, rather than committing to any particular choice, we parameterize our definitions with respect to an arbitrary query language. Our request operation is defined for any query which given a tree returns an updated tree and a list of results.

2.4 Reduction semantics

Network contexts are pairs of process and store contexts (see Figure 6). For example, if $C_{\mathcal{N}} = (- \uplus B, (\nu c) -)$ then $C_{\mathcal{N}}[(D, P)] = (D \uplus B, (\nu c)P)$. We omit the subscripts from contexts when no ambiguity can arise. The reduction relation \longrightarrow for Core $Xd\pi$ describes process interaction, the interaction between processes and data, and the movement of processes across locations. The formal definition is given in Table 7. It relies on a standard notion of structural congruence for processes and networks defined in the Appendix, in Figure 21.

Rules (CRED CONTEXT) and (CRED STRUCT) are standard contextual rules which allow reduction under parallel composition, restriction and structural congruence. There are two rules for process movement between locations: rule (CRED STAY) describes the case where the process is already at the target location, and rule (CRED GO) allows a process $l \cdot \mathbf{go} \ m.P$ to move from l to m . Rule (CRED COM) states that if an output $l \cdot \overline{a}(\tilde{v})$ and an input $l \cdot a(\tilde{\pi}).P$ on the same channel a are in the same location l (part of the store), and the values \tilde{v} match the input patterns $\tilde{\pi}$ (there is a substitution σ such that $\tilde{v} = \tilde{\pi}\sigma$), then communication

Figure 6: Syntax: contexts

$$\begin{aligned}
\mathcal{K}_{\mathcal{P}} &\stackrel{\text{def}}{=} C_{\mathcal{P}}[-] ::= - \mid P \mid C_{\mathcal{P}}[-] \mid C_{\mathcal{P}}[-] \mid P \mid (\nu c)C_{\mathcal{P}}[-] && \text{(PROCESS CONTEXTS)} \\
\mathcal{K}_{\mathcal{S}} &\stackrel{\text{def}}{=} C_{\mathcal{S}}[-] ::= - \mid C_{\mathcal{S}}[-] \uplus D && \text{(STORE CONTEXTS)} \\
\mathcal{K}_{\mathcal{N}} &\stackrel{\text{def}}{=} C_{\mathcal{N}}[-, -] ::= (C_{\mathcal{S}}[-], C_{\mathcal{P}}[-]), \text{ dom}(C_{\mathcal{P}}) \subseteq \text{ dom}(C_{\mathcal{S}}) && \text{(NETWORKS CONTEXTS)} \\
(C_{\mathcal{S}}[-], C_{\mathcal{P}}[-])[D, P] &\stackrel{\text{def}}{=} (C_{\mathcal{S}}[D], C_{\mathcal{P}}[P]) && \text{(CONTEXT APPLICATION)}
\end{aligned}$$

$$\text{Notation: } \{l \mapsto T\}(l) \stackrel{\text{def}}{=} T; \quad (D \uplus B)(l) \stackrel{\text{def}}{=} \begin{cases} D(l) & \text{if } l \in \text{dom}(D) \\ B(l) & \text{if } l \in \text{dom}(B) \end{cases}.$$

Convention: $D \uplus B$ is defined if and only if $\text{dom}(D) \cap \text{dom}(B) = \emptyset$.
Function dom is defined in Figure 22.

Figure 7: Semantics: reduction relation for Core $\text{Xd}\pi$

$$\begin{aligned}
(\{l \mapsto T\}, l \cdot \text{go } l.P \mid Q) &\longrightarrow (\{l \mapsto T\}, P \mid Q) && \text{(CRED STAY)} \\
(\{l \mapsto T\} \uplus \{m \mapsto S\}, l \cdot \text{go } m.P \mid Q) &\longrightarrow (\{l \mapsto T\} \uplus \{m \mapsto S\}, P \mid Q) && \text{(CRED GO)} \\
(\{l \mapsto T\}, \overline{l \cdot c}(\tilde{\pi}\sigma) \mid l \cdot c(\tilde{\pi}).\mathbf{P} \mid Q) &\longrightarrow (\{l \mapsto T\}, \mathbf{P}\sigma \mid Q) && \text{(CRED COM)} \\
(\{l \mapsto T\}, \overline{l \cdot c}(\tilde{\pi}\sigma) \mid !l \cdot c(\tilde{\pi}).\mathbf{P} \mid Q) &\longrightarrow (\{l \mapsto T\}, !l \cdot c(\tilde{\pi}).\mathbf{P} \mid \mathbf{P}\sigma \mid Q) && \text{(CRED COM!)} \\
(\{l \mapsto T\}, (x, \tilde{\pi})\mathbf{P} \circ \langle l, \tilde{\pi}\sigma \rangle \mid Q) &\longrightarrow (\{l \mapsto T\}, \mathbf{P}\{l/x\}\sigma \mid Q) && \text{(CRED RUN)} \\
\frac{\mathfrak{E}(p, T) = (T', U_1 \mid \dots \mid U_n \mid \emptyset)}{(\{l \mapsto T\}, l \cdot \text{req}_p \langle c \rangle \mid Q) \longrightarrow (\{l \mapsto T'\}, \overline{l \cdot c}(\mathbf{r}[U_1] \mid \dots \mid \mathbf{r}[U_n] \mid \emptyset) \mid Q)} &&& \text{(CRED REQUEST)} \\
\frac{\text{(CRED CONTEXT)} \quad N \longrightarrow N'}{C_{\mathcal{N}}[N] \longrightarrow C_{\mathcal{N}}[N']} && \frac{\text{(CRED STRUCT)} \quad N \equiv M \longrightarrow M' \equiv N'}{N \longrightarrow N'} &&
\end{aligned}$$

Convention: in this table c ranges over $\mathcal{C}_{\mathcal{P}} \cup \mathcal{C}_{\mathcal{S}}$.

takes place and execution proceeds with $\mathbf{P}\sigma$. Rule (CRED COM!) is similar, but leaves the replicated input process $!l \cdot a(\tilde{\pi}).\mathbf{P}$ in place for further use. We show an example of the communication of a private channel over a service channel below. The reduction step involves the use of structural congruence to extend the scope of the restricted name before communication (*scope extrusion*):

$$\begin{aligned}
&(\{l \mapsto T\}, (\nu c)(\overline{l \cdot a} \langle c, \mathbf{b}[\] \rangle) \mid l \cdot a(x, \mathbf{b}[y]).(\overline{l \cdot x} \langle y \rangle \mid \mathbf{P})) \\
&\longrightarrow (\{l \mapsto T\}, (\nu c)(\overline{l \cdot c} \langle \emptyset \rangle \mid \mathbf{P}\{c/x, \emptyset/y\}))
\end{aligned}$$

Rule (CRED RUN) runs a script, passing as the first parameter the name of the location where it is going to run. Rule (CRED REQUEST) applies the query denoted

Figure 8: Syntax: Sam queries

$\widehat{p}, \widehat{q} ::=$	path expressions
ε	empty path
\mathbf{A}/\widehat{p}	follow an edge with label in set \mathbf{A} , then \widehat{p}
$\square\widehat{p}$	follow occurrences of \widehat{p} anywhere
$p, q ::= \widehat{p}(\pi)\mathbf{V}$	queries: follow \widehat{p} , match π and insert \mathbf{V}
$\mathbf{A}, \mathbf{B} \in \wp(\mathcal{E})$	(LABEL SETS)
$p, q \in \mathcal{Q}$	(QUERIES)

Notation: $\text{copy}_{\widehat{p}}(\pi) \stackrel{\text{def}}{=} \widehat{p}(\pi)\pi$; $\text{cut}_{\widehat{p}}(\pi) \stackrel{\text{def}}{=} \widehat{p}(\pi)\emptyset$; $\text{paste}_{\widehat{p}}\langle E \rangle \stackrel{\text{def}}{=} \widehat{p}(x)E|x$;
 $*/\widehat{p} \stackrel{\text{def}}{=} \mathcal{E}/\widehat{p}$; $\mathbf{a}/\widehat{p} \stackrel{\text{def}}{=} \{\mathbf{a}\}/\widehat{p}$.

Convention: we omit a trailing ε from a path, for example we write $\mathbf{A}/$ for \mathbf{A}/ε .

by p on the local tree T , obtaining an updated tree T' which replaces T , and a list of results $U_1 | \dots | U_n | \emptyset$ which is turned into a tree of results labelled with r (a reserved label used to denote results) sent on channel c at l . We show a simple example of update, supposing that p is a query which extracts from a tree the data found by following the path \mathbf{a}/\mathbf{b} :

$$\begin{aligned} & (\{l \mapsto \mathbf{a}[\mathbf{b}[V_1] | \mathbf{b}[V_2] | \mathbf{a}[U]]\}, l \cdot \text{req}_p \langle c \rangle | P) \\ \longrightarrow & \longrightarrow (\{l \mapsto \mathbf{a}[\mathbf{b}[] | \mathbf{b}[] | \mathbf{a}[U]]\}, \overline{l \cdot c} \langle r[V_1] | r[V_2] \rangle | P) \end{aligned}$$

Note that the subtrees V_i removed from the store are returned as results by the output on c .

2.5 A sample query and update language

In this section, we define a particular query and update language inspired by XPath [37] which will be used in the examples later on. The result of evaluating a query against a piece of data (when defined) is a pair consisting of a new piece of data, intended to replace the original one, and a list of results, intended to be used by the continuation of the process that executed the query.

The syntax for queries is given in Figure 8. A query $\widehat{p}(\pi)\mathbf{V}$ is formed by a path expression \widehat{p} followed by an update expression $(\pi)\mathbf{V}$. A path expression \mathbf{A}/\widehat{p} applied to a tree $\mathbf{a}[V] | T$ evaluates p on V if \mathbf{a} is in the set \mathbf{A} , and evaluates itself on the rest of the tree T . A recursive expression $\square p$ applied to a tree $\mathbf{a}[V] | T$ evaluates p on any node in the tree in a bottom up fashion.⁹ First it evaluates $\square p$ on V and T , obtaining the updated items V' and T' , then it

⁹Suppose we chose a top-down strategy instead. A simple query like “add a subtree $\mathbf{a}[]$ inside any branch labelled \mathbf{a} ” on the tree $\mathbf{a}[]$ should be ruled out, because its evaluation diverges: each time a new subtree is added there is a new branch to update. Inconsistencies

Figure 9: Semantics: query evaluation for Sam

$\mathfrak{E}((\pi)\mathbf{V}, \pi\sigma) = (\mathbf{V}\sigma, \pi\sigma \upharpoonright \emptyset)$	(EVAL MATCH)
$\mathfrak{E}((\pi)\mathbf{V}, U) = (U, \emptyset) \quad \text{if } U \neq \pi\sigma$	(EVAL MISMATCH)
$\frac{\mathbf{a} \in \mathbf{A} \quad \mathfrak{E}(p, V) = (V', L) \quad \mathfrak{E}(\mathbf{A}/p, T) = (T', L')}{\mathfrak{E}(\mathbf{A}/p, \mathbf{a}[V] \upharpoonright T) = (\mathbf{a}[V'] \upharpoonright T', L \upharpoonright L')}$	(EVAL EDGE FOLLOW)
$\frac{\mathbf{a} \notin \mathbf{A} \quad \mathfrak{E}(\mathbf{A}/p, T) = (T', L')}{\mathfrak{E}(\mathbf{A}/p, \mathbf{a}[V] \upharpoonright T) = (\mathbf{a}[V] \upharpoonright T', L')}$	(EVAL EDGE DISCARD)
$\mathfrak{E}(\mathbf{A}/p, U) = (U, \emptyset) \quad \text{if } U \neq \mathbf{a}[V] \upharpoonright T$	(EVAL NOT EDGE)
$\frac{\begin{array}{l} \mathfrak{E}(\square p, V) = (V', L) \\ \mathfrak{E}(\square p, T) = (T', L') \\ \mathfrak{E}(p, \mathbf{a}[V'] \upharpoonright T') = (T'', L'') \end{array}}{\mathfrak{E}(\square p, \mathbf{a}[V] \upharpoonright T) = (T'', L \upharpoonright L' \upharpoonright L'')}$	(EVAL ANYWHERE TREE)
$\mathfrak{E}(\square p, U) = \mathfrak{E}(p, U) \quad \text{where } U \neq \mathbf{a}[V] \upharpoonright T$	(EVAL ANYWHERE ELSE)

The infix operator on lists \upharpoonright appends its second argument to its first argument. Query evaluation \mathfrak{E} is a partial function from $\mathcal{Q} \times \mathcal{D}$ to $\mathcal{D} \times \text{lists}(\mathcal{D})$.

evaluates p on $\mathbf{a}[V'] \upharpoonright T'$, combining the results together. The update expression $(\pi)\mathbf{V}$ is a binding pattern π followed by a data term \mathbf{V} . When $(\pi)\mathbf{V}$ is applied to each data item U selected by \hat{p} , such that $U = \pi\sigma$ for a closing substitution σ , the expression returns the new data item $\mathbf{V}\sigma$ and the result U . If there is no such substitution, the expression returns the original data U and the empty result \emptyset . The formal definition of query evaluation is given in Figure 9.

Definition 2.2 (Sample Query Language) *The sample query language Sam is the triple $(\mathcal{Q}, fv, \mathfrak{E})$ where \mathcal{Q} is defined in Figure 8, \mathfrak{E} is defined in Figure 9 and fv is defined as $fv(\hat{p}(\pi)\mathbf{V}) = fv(\mathbf{V}) \setminus fv(\pi)$.*

We show now how Sam is capable of expressing some intuitive tree manipulations, using the macros defined in Figure 8. The query $q = \text{copy}_{\mathbf{b}/}(y@x)$ reads the query and the location of any pointer contained in branches labelled \mathbf{b} at the top level. For example,

$$\mathfrak{E}(q, \mathbf{b}[T] \upharpoonright \mathbf{b}[p@l] \upharpoonright \mathbf{b}[p'@l']) = (\mathbf{b}[T] \upharpoonright \mathbf{b}[p@l] \upharpoonright \mathbf{b}[p'@l'], p@l \upharpoonright p'@l')$$

The query $q = \text{cut}_{\mathbf{a}/\square\mathbf{b}}(x)$ removes the contents of any branch labelled \mathbf{b} found after an initial branch \mathbf{a} , and returns the removed data as the results. For

of this kind are well-known in languages for updating trees, and there is no general agreement on which strategy should be preferred. Our results do not depend on the strategy chosen for our query language.

example,

$$\mathfrak{E}(q, \mathbf{b}[V] \mid \mathbf{a}[c[\mathbf{b}[U]]]) = (\mathbf{b}[V] \mid \mathbf{a}[c[\mathbf{b}[U]]], U)$$

The query $q = \text{paste}_{\mathbf{a}/\ast}/\langle \mathbf{e}[] \rangle$ adds a branch $\mathbf{e}[]$ to any child of \mathbf{a} , by reading the contents of each child and pasting them back with prefixed the new branch $\mathbf{e}[]$. For example,

$$\mathfrak{E}(q, \mathbf{a}[\mathbf{b}[] \mid c[\mathbf{d}[]]]) = (\mathbf{a}[\mathbf{b}[\mathbf{e}[]] \mid c[\mathbf{e}[] \mid \mathbf{d}[]]], \emptyset \mid (\mathbf{d}[]))$$

where the results are the list $\emptyset \mid (\mathbf{d}[])$ where the first element is the empty tree (the contents of \mathbf{b}) and the second element is tree $\mathbf{d}[]$ (the contents of \mathbf{c}). Note that the query for pasting data is defined only if each selected node (each child of \mathbf{a}) contains a tree¹⁰, since otherwise the resulting tree would be ill-formed. The query $\square(\pi)\mathbf{V}$ where $\pi = \mathbf{b}[\langle x \rangle] \mid y$ and $\mathbf{V} = c[\langle x \rangle] \mid y$ relabels each branch \mathbf{b} containing a script to \mathbf{c} :

$$\mathfrak{E}(\square(\pi)\mathbf{V}, \mathbf{a}[\mathbf{b}[\langle A \rangle] \mid \mathbf{b}[T] \mid \mathbf{b}[\langle A' \rangle]]) = (\mathbf{a}[c[\langle A \rangle] \mid \mathbf{b}[T] \mid c[\langle A' \rangle]], L)$$

where the results $L = (\mathbf{b}[\langle A \rangle] \mid \mathbf{b}[T] \mid c[\langle A' \rangle]) \mid (\mathbf{b}[\langle A' \rangle])$ correspond to the two trees to which the pattern was applied successfully. Note that in the first result, which is the last computed, the last branch has already been relabelled.

On the other hand, our sample query language is not sophisticated enough to express (atomically) a query like “delete each branch labelled \mathbf{a} which contains a branch labelled \mathbf{b} ”, because if we do not know in advance where a branch labelled \mathbf{b} occurs in the contents of \mathbf{a} , we cannot write a pattern to select only the nodes containing \mathbf{b} . This limitation can be easily overcome by adopting a richer pattern language or by enriching the path expressions with conditions on the contents of nodes.

2.6 Example: Web services

We now describe a simple implementation of macros for defining and calling services in *Core Xdπ*.

Service definition and service call. A service definition is characterized by the invoking location \mathbf{l} , the name of the service \mathbf{a} , its input pattern $\tilde{\pi}$, its body $(x)\mathbf{P}$ and its output pattern $\tilde{\omega}$. The service at location \mathbf{l} receives on channel \mathbf{a} the input parameters $\tilde{\pi}$, a location name y and a channel name z (the latter parameters are used to return the result). The body $(x)\mathbf{P}$ takes a fresh channel name c (bound to x) in input, performs some arbitrary computation, and outputs the result on channel c . Note that the variables in $\tilde{\pi}$ may bind in \mathbf{P} . A forwarding process inputs the result from channel c according to the output pattern $\tilde{\omega}$, and forwards it to location y on channel z .

$$\begin{array}{l} \text{(SERVICE DEFINITION)} \quad \mathbf{l} \cdot \text{Define } \mathbf{a}(\tilde{\pi}) \text{ as } (x)\mathbf{P} \text{ output } \langle \tilde{\omega} \rangle \\ \quad \stackrel{\text{def}}{=} \mathbf{l} \cdot \mathbf{a}(\tilde{\pi}, y, z).(\nu c)((x)\mathbf{P} \circ \langle c \rangle \mid \mathbf{l} \cdot c(\tilde{\omega}).\mathbf{l} \cdot \overline{y} \cdot z(\tilde{\omega})) \end{array}$$

¹⁰A type system regulating the contents of trees could prevent processes from getting stuck because of undefined queries.

The service call is dual. It specifies the location l from which the service is invoked (and to which it is returned), the location m and the name a of the service, its parameters \tilde{v} , and a continuation process Q with patterns $\tilde{\pi}$ for parsing the results.

$$\begin{aligned} \text{(SERVICE CALL)} \quad & l \cdot \text{Call } m \cdot a(\tilde{v}) \text{ return } (\tilde{\pi})Q \\ & \stackrel{\text{def}}{=} (\nu c)(l \cdot \overline{m \cdot a}(\tilde{v}, l, c) \mid l \cdot c(\tilde{\pi}).Q) \end{aligned}$$

The parameters l and c sent on a are used by the forwarding process in the service definition to return the result to the continuation process Q . For example, a service providing querying capabilities on its local store, and the corresponding service call, could be defined respectively as

$$\begin{aligned} m \cdot \text{Define } \text{query}(x_1) \text{ as } (x) \text{req}_{x_1} \langle x \rangle \text{ output } \langle x_2 \rangle \\ l \cdot \text{Call } m \cdot \text{query}(p) \text{ return } (x)Q \end{aligned}$$

The service takes as input a query x_1 and executes the corresponding request on the local store. The forwarding part of the service definition will intercept the request result and send it on c at l , where it is passed on to Q on variable x .

Subscriptions. We can easily generalize service definitions to cover the case of *push services*, which send a stream of results to a client in reply to a single service call. The only difference between the code below and (SERVICE DEFINITION) is the presence of a replicated input in the forwarding process

$$\begin{aligned} \text{(PUSH SERVICE)} \quad & l \cdot \text{Push } a(\tilde{\pi}) \text{ as } (x)P \text{ output } \langle \tilde{\omega} \rangle \\ & \stackrel{\text{def}}{=} !l \cdot a(\tilde{\pi}, y, z).(\nu c)((x)P \circ \langle c \rangle \mid !l \cdot c(\tilde{\omega}).l \cdot \overline{y \cdot z}(\tilde{\omega})) \end{aligned}$$

The corresponding service subscription waits for multiple results on channel c :

$$\begin{aligned} \text{(SUBSCRIPTION)} \quad & l \cdot \text{Subscribe } m \cdot a(\tilde{v}) \text{ return } (\tilde{\pi})Q \\ & \stackrel{\text{def}}{=} (\nu c)(l \cdot \overline{m \cdot a}(\tilde{v}, l, c) \mid !l \cdot c(\tilde{\pi}).Q) \end{aligned}$$

If desired, the streamed results received by the client can be combined together using a loop.

Result forwarding. In order to have complete control on the return parameters, in certain cases we will bypass the service call code, and use only a migration step followed by a service invocation. For example, let

$$\text{Service} = m \cdot \text{Define } \text{query}(x_1) \text{ as } (x) \text{req}_{x_1} \langle x \rangle \text{ output } \langle x_2 \rangle$$

and consider the network

$$N = (D, (\nu c) \left(\begin{array}{l} l \cdot \overline{m \cdot \text{query}}((w)w, n, c) \\ | \text{Service} \\ | n \cdot c(x). \text{req}_{(w)x} \langle c \rangle \end{array} \right))$$

where $D = \{l \mapsto T_0, m \mapsto S, n \mapsto \emptyset\}$. The service invocation migrates to m and triggers the service, passing as return parameters n and c . The local request at m copies the whole tree S and forwards it to c on n :

$$\begin{aligned} N &\longrightarrow^* (D, (\nu c)(Service \mid m \cdot \overline{n \cdot c} \langle S \rangle \mid n \cdot c(x).req_{(w)x}(c))) \\ &\longrightarrow^* (\{l \mapsto T_0, m \mapsto S, n \mapsto S\}, Service \mid (\nu c)(\overline{n \cdot c} \langle \emptyset \rangle)) \end{aligned}$$

At n the code listening on c receives the result and replaces the local tree. We will follow this strategy in several examples, redirecting the results of a service to a location different from the one that issued the service call.

3 Behavioural Equivalences

We investigate behavioural equivalences for Core $Xd\pi$. First we define network equivalence, which dictates when two networks can be considered indistinguishable with respect to an externally defined comparison. Our network equivalence is parametric with respect to the language used for querying and updating documents.

We also define a process equivalence, which establishes when two processes can replace each other in a network without affecting network equivalence. We would like to reason about the equivalence of groups of processes, possibly interacting across several locations, and obtain results which are robust with respect to changes in the data stored in the local repositories and the behaviour of other parallel processes. Even if we decided to restrict our optimizations to the processes running in a single local peer, we must be ready to reason about partial network specifications in the case where parts of the local process migrate to other locations in order to interact with remote data or services, reinforcing the case for global reasoning. The structure of Core $Xd\pi$ networks, where processes are located explicitly and are separated from the data store, facilitates this process. For example, we can express a partial specification of a network by means of located processes running in parallel, possibly sharing private names. Located processes are equivalent if the networks obtained by composing them with arbitrary stores are equivalent. However, process equivalence is hard to use directly because it requires a costly property of closure under contexts. Instead, we use a labelled transition system to define a coinductive equivalence relation (called *domain bisimilarity*), which does not quantify over contexts and which entails process equivalence.

The definition of domain bisimilarity is non-standard, due to the fact that scripts (which can appear in data) are part of the values, and process equivalences are sensitive to the set of locations constituting the network. We address these two problems by adapting existing techniques for translating messages containing scripts into ones where each script is replaced by a first-order value [30, 17], and by generalizing the notion of bisimulation to families of relations indexed by sets of locations.

3.1 Reduction and Domain Congruence

Reduction Congruence. Network equivalence for Core $Xd\pi$ is a standard reduction-closed, contextual equivalence which preserves some observation predicates. In [18] we consider different choices of observation predicates (in particular the shape of the data tree of a location and the presence of output actions in a process), and we study the formal relationship between the corresponding reduction congruences. Here, we focus on request observations because it seems natural to observe the effect processes have on data. In addition, the resulting reduction congruence coincides with the one resulting from output observation and implies the one resulting from tree-shape observations.

Definition 3.1 (Request observation predicate) We define the request observation predicate $\downarrow_{l.p}$ as $(D, P) \downarrow_{l.p} \iff \exists C, c, Q. P \equiv C[l \cdot \text{req}_p \langle c \rangle \mid Q]$ and the weak observation predicate $\Downarrow_{l.p}$ as $N \Downarrow_{l.p} \iff \exists N'. N \xrightarrow{*} N' \text{ and } N' \downarrow_{l.p}$.

Definition 3.2 (Reduction Congruence) Reduction congruence \simeq on Core $Xd\pi$ networks is the largest symmetric relation \simeq which is

- *observation preserving:* $N \simeq M \implies \forall l, p. N \Downarrow_{l.p} \implies M \Downarrow_{l.p}$
- *reduction closed:* $N \simeq M \implies \forall N'. N \xrightarrow{*} N' \implies \exists M'. M \xrightarrow{*} M' \text{ and } N' \simeq M'$
- *contextual:* $N \simeq M \implies \forall C. C[N] \simeq C[M]$.

For example, we have

$$(\{l \mapsto T\}, \overline{l \cdot a} \langle c \rangle) \not\approx (\{l \mapsto T\}, \overline{l \cdot a} \langle b \rangle)$$

because the context $K = (-, - \mid l \cdot a(x). \overline{l \cdot x} \mid l \cdot c.l \cdot \text{req}_p \langle a \rangle)$ can tell a difference between the two processes.

In order to use equational reasoning, it is important to remark that reduction congruence is an equivalence relation (the proof is completely standard).

Observation 3.3 (Equivalence) Reduction congruence \simeq is an equivalence relation.

Since reduction congruence is based on contexts which do not inhibit reduction, a simple test for equivalence consists in checking if two terms can reduce to each other.

Lemma 3.4 (Mutual Reduction) If $N \xrightarrow{*} M$ and $M \xrightarrow{*} N$ then $N \simeq M$.

Proof. We show that the relation

$$\simeq = \left\{ (C[N], C[M]) : N \xrightarrow{*} M, M \xrightarrow{*} N \right\}$$

is contained in \simeq . The proof is symmetric. Consider an arbitrary pair $(N, M) \in \simeq$. Suppose $N \Downarrow_{l.p}$. It must be the case that $N \xrightarrow{*} N' \downarrow_{l.p}$. By hypothesis

$M \xrightarrow{*} N$, hence $M \Downarrow_{l.p}$. Suppose $N \xrightarrow{*} N'$. By hypothesis $M \xrightarrow{*} N$, hence $M \xrightarrow{*} N'$. By definition of $\xrightarrow{*}$, $N' \simeq N'$. We need to show that for an arbitrary $C[-]$, $C[N] \simeq C[M]$. Since $C[-]$ is a reduction context, $N \xrightarrow{*} M \implies C[N] \xrightarrow{*} C[M]$ and $M \xrightarrow{*} N \implies C[M] \xrightarrow{*} C[N]$. By definition, $C[N] \simeq C[M]$. \square

Domain Congruence. We now define process equivalence for Core $Xd\pi$. This equivalence depends on the locations present in the network, but not on the actual contents of the stores. Consider replacing the definition of a service at location l , which uses only local data, with one located at m (where there is a cached copy of the same data) and providing an equivalent functionality. If location m is connected, then the behaviour of the services is the same. On the other hand, if location m is not connected, the behaviour of the services is different. With network equivalence, the connected locations are those in the domain of the store. With process equivalence, we must state explicitly the locations which we assume to be part of the network. As a consequence, process equivalence is indexed by a *domain* (a set of locations) Λ .

A Core $Xd\pi$ process can be seen as a partial specification of a network, describing only some of the processes running in some of the locations. This point of view is useful for reasoning about replacing components which are part of some distributed data-exchange protocol. Accordingly, we say that two processes P and Q are equivalent with respect to a domain Λ if all the networks containing *at least* the locations in Λ and either P or Q , are equivalent.

Besides comparing partial network specifications, process equivalences can be useful for example to replace optimized pieces of code inside a specific process. For that purpose, we need a more general class of process contexts which include prefixes.

Definition 3.5 (Extended Contexts) Extended process contexts \mathcal{K}_f are the terms generated by

$$C ::= - \mid C \mid \mathbf{P} \mid \mathbf{P} \mid C \mid (\nu c)C \mid \mathbf{l}\cdot\mathbf{a}(\tilde{\pi}).C \mid !\mathbf{l}\cdot\mathbf{a}(\tilde{\pi}).C \mid \mathbf{l}\cdot\mathbf{go}\ m.C$$

Unless we specify otherwise, from now on we use $C[-]$ to denote extended contexts.

Definition 3.6 (Domain Congruence) Given a set of location names Λ , we define the induced domain congruence \sim^Λ on processes by

$$\sim^\Lambda = \{(\mathbf{P}, \mathbf{Q}) : \forall D, C[-]. \Lambda \subseteq \text{dom}(D) \implies (D, C[\mathbf{P}]) \simeq (D, C[\mathbf{Q}])\}$$

where each $C[-]$ is closing for both \mathbf{P} and \mathbf{Q} .

Domain congruence is monotonic: the larger the set of locations which we assume to be part of the network, the larger the number of processes which we can equate.

Observation 3.7 (Monotonicity) If $\Lambda \subseteq \Lambda'$ then $\sim^\Lambda \subseteq \sim^{\Lambda'}$.

Figure 10: Notation for asynchronous processes

(FORWARDER)	$\mathbf{l}\cdot\text{FW}(\mathbf{a}, \mathbf{b}, \tilde{\pi}) \stackrel{\text{def}}{=} \mathbf{l}\cdot\mathbf{a}(\tilde{\pi})\cdot\overline{\mathbf{l}\cdot\mathbf{b}}(\tilde{\pi})$
(EQUATOR)	$\mathbf{l}\cdot\text{EQ}(\mathbf{a}, \mathbf{b}, \tilde{\pi}) \stackrel{\text{def}}{=} !\mathbf{l}\cdot\text{FW}(\mathbf{a}, \mathbf{b}, \tilde{\pi}) \mid !\mathbf{l}\cdot\text{FW}(\mathbf{b}, \mathbf{a}, \tilde{\pi})$
(DISTRIBUTED FORWARDER)	$\mathbf{l}\cdot\text{dFW}(\mathbf{a}, \mathbf{m}, \mathbf{b}, \tilde{\pi}) \stackrel{\text{def}}{=} \mathbf{l}\cdot\mathbf{a}(\tilde{\pi})\cdot\overline{\mathbf{l}\cdot\mathbf{m}\cdot\mathbf{b}}(\tilde{\pi})$
(DISTRIBUTED EQUATOR)	$\text{dEQ}(\mathbf{l}, \mathbf{a}, \mathbf{m}, \mathbf{b}, \tilde{\pi}) \stackrel{\text{def}}{=} !\mathbf{l}\cdot\text{dFW}(\mathbf{a}, \mathbf{m}, \mathbf{b}, \tilde{\pi}) \mid !\mathbf{m}\cdot\text{dFW}(\mathbf{b}, \mathbf{l}, \mathbf{a}, \tilde{\pi})$

Proof. Follows easily by Definition 3.6. □

Due to the several explicit (and implicit) universal quantifications involved in Definition 3.6, it is very difficult to show directly that two processes are domain congruent. For this reason, in Section 3.2 we will introduce a proof method which does not require closure under contexts and which entails domain congruence.

Asynchronous Laws. Core $Xd\pi$ is an extension of the asynchronous π -calculus, so we consider some equational laws inspired by the latter. Consider the process definitions given in Figure 10. The *asynchrony law*, stating that the presence of a communication buffer cannot be observed, holds also in Core $Xd\pi$ (see Section 3.2.2 for a proof):

$$!\mathbf{l}\cdot\text{FW}(\mathbf{a}, \mathbf{a}, \tilde{\pi}) \sim_r^\Lambda \mathbf{0}$$

The law stating that two channels \mathbf{a} and \mathbf{b} cannot be distinguished if they are part of the same equator does not hold. For example,

$$\mathbf{l}\cdot\text{EQ}(\mathbf{a}, \mathbf{b}, \tilde{\pi}) \mid \overline{\mathbf{l}\cdot\mathbf{c}}(\mathbf{a}) \not\sim_r^\Lambda \mathbf{l}\cdot\text{EQ}(\mathbf{a}, \mathbf{b}, \tilde{\pi}) \mid \overline{\mathbf{l}\cdot\mathbf{c}}(\mathbf{b})$$

because a context could intercept the channel name \mathbf{a} and use it in some fresh location m where \mathbf{a} and \mathbf{b} are not equated. We have instead a new law about equating located channels across different locations:

$$\text{dEQ}(\mathbf{l}, \mathbf{a}, \mathbf{m}, \mathbf{b}, \tilde{\pi}) \mid \overline{\mathbf{l}\cdot\mathbf{a}}(\tilde{\pi}\sigma) \sim_r^{\{l,m\}} \text{dEQ}(\mathbf{l}, \mathbf{a}, \mathbf{m}, \mathbf{b}, \tilde{\pi}) \mid \overline{\mathbf{m}\cdot\mathbf{b}}(\tilde{\pi}\sigma)$$

This law could be useful to show that we can replicate Web services (improving efficiency) without the clients needing to be aware of the change.

3.2 Bisimilarity

In this Section, we define a coinductive equivalence relation (*bisimilarity*), which does not quantify over contexts and which entails domain congruence.

3.2.1 Labelled transition system

A typical proof that processes are bisimilar involves a universal quantification over labelled transitions. Since Core $Xd\pi$ values include scripts, and labels

Figure 11: Syntax: configurations

$\mathbf{K} ::=$	configuration terms	
\mathbf{P}	process terms (built as for Core $Xd\pi$)	
$\mathbf{K} \mid \mathbf{K}$	parallel composition	
$(\nu c)\mathbf{K}$	restriction	
$\langle k \Leftarrow A \rangle$	definition for trigger name k with script A	
$\mathbf{A} ::= A \mid x \mid k$ script or variable or trigger name		
$h, i, j, k \in \mathcal{Y}$ ($\mathcal{Y} \cap \mathcal{C}_p = \emptyset$, \mathcal{Y} countably infinite)		(TRIGGER NAMES)
$K, L \in \mathcal{W} \stackrel{\text{def}}{=} \{\mathbf{K} : fv(\mathbf{K}) = \emptyset, \text{unique}(\mathbf{K})\}$		(CONFIGURATIONS)
$A \in \mathcal{A}_W \stackrel{\text{def}}{=} \{A : A \in \mathcal{A}, \text{triggers}(A) = \emptyset\}$		(SCRIPTS)
$C_W[-] ::= - \mid C_W[-] \mid \mathbf{K} \mid \mathbf{K} \mid C_W[-] \mid (\nu c)C_W[-]$		(CONFIGURATION CONTEXTS)

The function *triggers* and the predicate *unique* are defined in Figure 12. Apart from the redefinition of \mathbf{A} , A and p (see below), the grammars for trees, processes and values are the same as for Core $Xd\pi$.

For trigger definitions, we adopt the following notation:

$$\begin{aligned} \Theta^{\tilde{k}} &\stackrel{\text{def}}{=} \langle k_1 \Leftarrow A_1 \rangle \mid \dots \mid \langle k_n \Leftarrow A_n \rangle, \text{ where all } k_i \text{ are distinct, } n \geq 0; \\ \Theta^{\tilde{k} \curvearrowright \tilde{j}} &\stackrel{\text{def}}{=} \Theta^{\tilde{k}} \{ \tilde{j} / \tilde{k} \}, \text{ when } \{ \tilde{j} / \tilde{k} \} \text{ is defined;} \\ \Theta &\stackrel{\text{def}}{=} \Theta^{\tilde{k}}, \text{ when } \tilde{k} \text{ is not important.} \\ t^{\langle k \Leftarrow A \rangle} &\stackrel{\text{def}}{=} t\{A/k\}, \text{ for any term } t, \text{ and similarly for } t^{\Theta^{\tilde{k}}}. \end{aligned}$$

typically include values, we risk falling back to quantifying over processes. Following the approach of [30, 17], we avoid this problem by translating messages containing scripts into ones where each script is replaced by a *trigger name* (a first-order value), and by placing in parallel to the process being analyzed some *definitions* associating to each trigger name the code of the corresponding script. By including these definitions in the code, we are able to analyze also the interaction between scripts and their contexts.

Configurations. We introduce *configurations*, which are processes extended with the trigger names and definitions mentioned above. The formal syntax is given in Figure 11. Note that \mathbf{A} denotes now a script, a variable or a trigger name, hence processes can syntactically contain triggers. Nonetheless, scripts and queries are not allowed to contain triggers. In fact, trigger names and definitions are merely intermediate terms arising during the analysis of the transition of a process, and are not meant to be part of the user syntax. For a configuration K to be well-formed, no two definitions in K can have the same trigger name (predicate *unique*(K), defined in Figure 12). As a convention, we

Figure 12: Function *triggers* and predicate *unique*.

$$\text{triggers}(t) = \text{fn}(t) \cap \mathcal{Y}$$

$$\frac{\text{unique}(\mathbf{K}) \quad \text{unique}(\mathbf{K}')}{\text{dfs}(\mathbf{K}) \cap \text{dfs}(\mathbf{K}') = \emptyset} \quad \frac{\text{unique}(\mathbf{K})}{\text{unique}((\nu c)\mathbf{K})} \quad \text{unique}(\langle k \leftarrow A \rangle)$$

The function *dfs*, returning the triggers defined by a configuration, is given by

$$\text{dfs}(\mathbf{K} \mid \mathbf{K}') = \text{dfs}(\mathbf{K}) \cup \text{dfs}(\mathbf{K}') \quad \text{dfs}((\nu c)\mathbf{K}) = \text{dfs}(\mathbf{K}) \quad \text{dfs}(\langle k \leftarrow A \rangle) = \{k\}$$

let Θ , Ω and Φ range on groups of definitions. Note also that two groups of definitions $\Theta^{\tilde{k}}$ and $\Theta^{\tilde{j}}$ identified by the same name but by different vectors of triggers can in principle be arbitrarily different: it is an important syntactic convention which helps to simplify the notation, and is often used in the rest of the section. In the Appendix, in Figure 21 and Figure 23, we extend \equiv , *fv* and *fn* to configurations. In Figure 24 we extend the function *dom* of Figure 22 to configurations, and we define a function *scripts* returning the scripts present in a piece of data.

Queries. Queries used for updating can mention constant data, which may contain scripts. We assume two functions, *scripts* and *triggers*, which given a query return respectively the set of scripts and triggers it contains. The only condition that we need to impose on query evaluation consists of it not being dependent on the particular structure of scripts. In other words, if we replace a script in a query with a trigger name, then the result of the query should be equivalent up to substitution of the script for the trigger. Moreover, any script returned by the query must occur in the input tree or in the query itself. The condition is formalized below.

Definition 3.8 (Script Independence) *Let $\mathcal{L} = (\mathcal{Q}, \text{fv}, \mathfrak{E})$ be an arbitrary query language, let p, T be such that $\mathfrak{E}(p, T) = (S, L)$, and let p_0, T_0 be their first-order versions, such that $\text{scripts}(p_0) = \text{scripts}(T_0) = \emptyset$ and $p = p_0^{\Theta^{\tilde{j}}}, T = T_0^{\Omega^{\tilde{k}}}$ for some $\Theta^{\tilde{j}}, \Omega^{\tilde{k}}$.*

The query language \mathcal{L} is script independent if for all $\Theta^{\tilde{j}}, \Omega^{\tilde{k}}$ there exist $\Theta^{\tilde{h}}$ and $\Omega^{\tilde{i}}$ such that

- *query evaluation does not depend on the structure of scripts: there are S_0, L_0 with $\text{scripts}(S_0) = \text{scripts}(L_0) = \emptyset$ such that $\mathfrak{E}(p_0^{\Theta^{\tilde{j}}}, T_0^{\Omega^{\tilde{k}}}) = (S_0^{\Theta^{\tilde{h}}}, L_0^{\Omega^{\tilde{i}}})$*
- *no new scripts are introduced: for any definition $\langle k \leftarrow A \rangle$ occurring in $\Theta^{\tilde{h}}$ or $\Omega^{\tilde{i}}$ there must be a definition $\langle k' \leftarrow A \rangle$ occurring in $\Theta^{\tilde{j}}$ or $\Omega^{\tilde{k}}$.*

Figure 13: Extraction relation

$$\begin{array}{l}
\mathfrak{X}(l) = (l; \mathbf{0}) \\
\mathfrak{X}(p) = \mathfrak{X}_{\mathcal{Q}}(p) \\
\mathfrak{X}(A) = (k; \langle k \leftarrow A \rangle) \\
\mathfrak{X}(c) = (c; \mathbf{0}) \\
\frac{\mathfrak{X}(\mathbf{E}) = (\mathbf{E}'; \Theta^{\tilde{k}}) \quad \mathfrak{X}(\mathbf{T}) = (\mathbf{T}'; \Omega^{\tilde{h}}) \quad (\{\tilde{k}, \tilde{h}\} \cap \text{fn}(\mathbf{E} \upharpoonright \mathbf{T})) \cup (\{\tilde{k}\} \cap \{\tilde{h}\}) = \emptyset}{\mathfrak{X}(\mathbf{E} \upharpoonright \mathbf{T}) = (\mathbf{E}' \upharpoonright \mathbf{T}'; \Theta^{\tilde{k}} \mid \Omega^{\tilde{h}})} \\
\frac{\mathfrak{X}(\mathbf{v}) = (\mathbf{v}'; \Theta^{\tilde{k}}) \quad \mathfrak{X}(\tilde{\mathbf{v}}) = (\tilde{\mathbf{v}}'; \Omega^{\tilde{h}}) \quad (\{\tilde{k}, \tilde{h}\} \cap \text{fn}(\mathbf{v}, \tilde{\mathbf{v}})) \cup (\{\tilde{k}\} \cap \{\tilde{h}\}) = \emptyset}{\mathfrak{X}(\mathbf{v}, \tilde{\mathbf{v}}) = (\mathbf{v}', \tilde{\mathbf{v}}'; \Theta^{\tilde{k}} \mid \Omega^{\tilde{h}})} \\
\frac{\mathfrak{X}(T) = (T'; \Theta^{\tilde{k}}) \quad \mathfrak{X}(D) = (D'; \Omega^{\tilde{h}}) \quad (\{\tilde{k}, \tilde{h}\} \cap \text{fn}(T, D)) \cup (\{\tilde{k}\} \cap \{\tilde{h}\}) = \emptyset}{\mathfrak{X}(\{l \mapsto T\} \uplus D) = (\{l \mapsto T'\} \uplus D'; \Theta^{\tilde{k}} \mid \Omega^{\tilde{h}})} \\
\mathfrak{X}(\mathbf{p}) = (\mathbf{p}'; \Theta) \\
\mathfrak{X}(\mathbf{p} \uplus l) = (\mathbf{p}' \uplus l; \Theta) \\
\mathfrak{X}(\emptyset) = (\emptyset; \mathbf{0}) \\
\mathfrak{X}(x) = (x; \mathbf{0}) \\
\mathfrak{X}(k) = (k; \mathbf{0}) \\
\frac{\mathfrak{X}(\mathbf{A}) = (\mathbf{v}; \Theta)}{\mathfrak{X}(\langle \mathbf{A} \rangle) = (\langle \mathbf{v} \rangle; \Theta)} \\
\frac{\mathfrak{X}(\mathbf{V}) = (\mathbf{V}'; \Theta)}{\mathfrak{X}(\mathbf{a}[\mathbf{V}]) = (\mathbf{a}[\mathbf{V}']; \Theta)} \\
\mathfrak{X}(\emptyset) = (\emptyset; \mathbf{0})
\end{array}$$

Notation: in this table c ranges over $\mathcal{C}_p \cup \mathcal{C}_s$.
The relation $\mathfrak{X}_{\mathcal{Q}}$ is specified in Definition 3.9.

Extracting scripts from values. Our strategy consists of translating values containing scripts into values containing trigger names only, extracting at the same time the corresponding definitions. For that purpose, we define in Figure 13 an extraction relation \mathfrak{X} which applies to Core Xdπ data and stores, and returns the corresponding first-order terms and the definitions extracted.

The definition of \mathfrak{X} is straightforward. The only points worth noting are that the premises of the rules for tuples, tree and store composition make sure that the trigger names remain disjoint, and that the rule for scripts replaces a script with a trigger and records the corresponding definition. The rule for queries invokes a specialized extraction relation $\mathfrak{X}_{\mathcal{Q}}$ which depends on the query language. $\mathfrak{X}_{\mathcal{Q}}$ can behave similarly to \mathfrak{X} , relating each query with its first-order version and the corresponding definitions, or can behave differently (for example being the identity function on queries and the constant 0 on configurations), as long as it satisfies the basic properties requested by the definition given below.

Definition 3.9 (Query Extraction) *The relation $\mathfrak{X}_{\mathcal{Q}}$ can be any subset of $\mathcal{Q}_C \times \mathcal{Q}_C \times \mathcal{W}$ satisfying the condition that if $\mathfrak{X}_{\mathcal{Q}}(p) = (p'; K)$ then*

1. K are well-formed definitions: $K = \Theta^{\tilde{k}}$;

Figure 14: Labels for the transition system

$\alpha_l ::=$	transition labels
$(\tilde{a}, \tilde{k})\overline{l \cdot c}(\tilde{v})$	output of \tilde{v} on c at l , extruding \tilde{a}, \tilde{k}
$l \cdot c(\tilde{v})$	input of \tilde{v} on c at l
$l \cdot \tau$	internal reduction at l
$(\tilde{k})l \cdot \text{req}(p)(T)$	request p at l extruding \tilde{k} , obtaining result T
$l \cdot k(\tilde{v})$	run the script defined by k with parameters l, \tilde{v}
$(\tilde{a}, \tilde{k})\overline{l \cdot j}(\tilde{v})$	assume running the script defined by trigger j with parameters l, \tilde{v} , extruding \tilde{a}, \tilde{k}

Convention: in the syntax above, c ranges over $\mathcal{C}_p \cup \mathcal{C}_s$, and t over $\mathcal{C}_p \cup \mathcal{C}_s \cup \mathcal{Y}$.

labels α_l	names n	bound names bn
$(\tilde{a}, \tilde{k})\overline{l \cdot t}(\tilde{v})$	$\{\tilde{a}\} \cup \{\tilde{k}\} \cup \{t\} \cup fn(\tilde{v})$	$\{\tilde{a}\} \cup \{\tilde{k}\}$
$l \cdot t(\tilde{v})$	$\{t\} \cup fn(\tilde{v})$	\emptyset
$l \cdot \tau$	\emptyset	\emptyset
$(\tilde{k})l \cdot \text{req}(p)(T)$	$\{\tilde{k}\} \cup \text{triggers}(p) \cup fn(T)$	$\{\tilde{k}\}$

$fn(\alpha_l) = n(\alpha_l) \setminus bn(\alpha_l)$

2. *trigger names can be extended as long as there are no clashes: $\text{triggers}(p') = \text{triggers}(p) \cup \{\tilde{k}\}$ and $\text{triggers}(p) \cap \{\tilde{k}\} = \emptyset$;*
3. *the new trigger names are defined up-to renaming: for all \tilde{j} distinct from $\text{triggers}(p)$, $\mathfrak{X}_{\mathcal{Q}}(p) = (p' \{ \tilde{j} / \tilde{k} \}; \Theta^{\tilde{k} \rightsquigarrow \tilde{j}})$;*
4. *substitution is the inverse of extraction: if $\mathfrak{X}_{\mathcal{Q}}(p) = (p'; \Theta)$ then $p = p'^{\Theta}$.*

Under the assumption (that we adopt henceforth) that $\mathfrak{X}_{\mathcal{Q}}$ respects Definition 3.9, the effects of relation \mathfrak{X} can be reversed by replacing, in the extracted first-order term, the new trigger names by the corresponding definitions.

Observation 3.10 (Extraction) *For any given term t , if $\mathfrak{X}(t) = (t'; \Theta)$ then $t = t'^{\Theta}$.*

Proof. By induction on the derivation of $\mathfrak{X}(t) = (t'; \Theta)$. □

Labelled transition system. Our labelled transition system incorporates ideas on asynchronous transitions from [15], and on translating higher-order actions into first order actions from [30, 17].

Figure 15: Labelled transition system

<p>(LTS COM)</p> $\frac{}{\overline{l \cdot c(\tilde{\pi}\sigma) \mid l \cdot c(\tilde{\pi}) \cdot \mathbf{P}} \xrightarrow{l \cdot \tau} \mathbf{P}\sigma}$	<p>(LTS !COM)</p> $\frac{}{\overline{l \cdot c(\tilde{\pi}\sigma) \mid !l \cdot c(\tilde{\pi}) \cdot \mathbf{P}} \xrightarrow{l \cdot \tau} \mathbf{P}\sigma \mid !l \cdot c(\tilde{\pi}) \cdot \mathbf{P}}$	
<p>(LTS RUN)</p> $\frac{}{(x, \tilde{\pi})\mathbf{P} \circ \langle l, \tilde{\pi}\sigma \rangle \xrightarrow{l \cdot \tau} \mathbf{P}\{l/x\}\sigma}$	<p>(LTS IN)</p> $\frac{\mathit{scripts}(\tilde{v}) = \emptyset}{0 \xrightarrow{l \cdot c(\tilde{v})} \overline{l \cdot c(\tilde{v})}}$	
<p>(LTS OUT)</p> $\frac{\mathfrak{X}(\tilde{v}) = (\tilde{v}'; \Theta^{\tilde{k}})}{\overline{l \cdot c(\tilde{v})} \xrightarrow{(\tilde{k})\overline{l \cdot c(\tilde{v}')}} \Theta^{\tilde{k}}}$	<p>(LTS TRIGGER)</p> $\frac{\mathfrak{X}(\tilde{v}) = (\tilde{v}'; \Theta^{\tilde{k}}) \quad j \notin \{\tilde{k}\}}{j \circ \langle l, \tilde{v} \rangle \xrightarrow{(\tilde{k})\overline{l \cdot j(\tilde{v}')}} \Theta^{\tilde{k}}}$	
<p>(LTS OPEN)</p> $\frac{a \in \mathit{fn}(\tilde{v}) \setminus \{\tilde{a}, c\} \quad K \xrightarrow{(\tilde{a}, \tilde{k})\overline{l \cdot c(\tilde{v})}} K'}{(\nu a)K \xrightarrow{(a, \tilde{a}, \tilde{k})\overline{l \cdot c(\tilde{v})}} K'}$	<p>(LTS REQ)</p> $\frac{\mathfrak{X}_{\mathcal{Q}}(p) = (p', \Theta^{\tilde{k}}) \quad \mathit{scripts}(T) = \emptyset \quad T = \mathbf{r}[U_1] \mid \dots \mid \mathbf{r}[U_n] \mid \emptyset}{l \cdot \mathit{req}_p \langle c \rangle \xrightarrow{(\tilde{k})\overline{l \cdot \mathit{req}(p')(T)}} \overline{l \cdot c(T)} \mid \Theta^{\tilde{k}}}$	
<p>(LTS DEF)</p> $\frac{\mathit{scripts}(\sigma) = \emptyset}{\langle k \leftarrow (x, \tilde{\pi})\mathbf{P} \rangle \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} \langle k \leftarrow (x, \tilde{\pi})\mathbf{P} \rangle \mid \mathbf{P}\{l/x\}\sigma}$	<p>(LTS GO)</p> $\frac{}{l \cdot \mathit{go} m \cdot \mathbf{P} \xrightarrow{m \cdot \tau} \mathbf{P}}$	
<p>(LTS RES)</p> $\frac{K \xrightarrow{\alpha_l} K' \quad a \notin n(\alpha_l)}{(\nu a)K \xrightarrow{\alpha_l} (\nu a)K'}$	<p>(LTS PAR)</p> $\frac{K \xrightarrow{\alpha_l} K' \quad \mathit{rel}(\alpha_l, L)}{K \mid L \xrightarrow{\alpha_l} K' \mid L}$	<p>(LTS STRUCT)</p> $\frac{K \equiv L \xrightarrow{\alpha_l} L' \equiv K'}{K \xrightarrow{\alpha_l} K'}$

Notation: $\mathit{rel}(\alpha_l, K) \stackrel{\text{def}}{=} \mathit{bn}(\alpha_l) \cap \mathit{fn}(K) = \emptyset$.

Convention: in this table c ranges over $\mathcal{C}_p \cup \mathcal{C}_s$, and t over $\mathcal{C}_p \cup \mathcal{C}_s \cup \mathcal{Y}$.

The labels of the transition system record what kind of interaction with the external environment is necessary for a configuration to evolve into another. Labels, along with the notions of their names, free names and bound names, are defined in Figure 14. Each label, including the one for internal reduction, shows explicitly the location where interaction takes place. By using appropriate conditions on the function $\mathit{scripts}$ in the rules of the labelled transition system (lts for short), we will guarantee that labels are first-order, as planned. The formal definition of the lts is given in Figure 15. We discuss the more interesting transition rules. Rules (LTS COM), (LTS !COM) and (LTS RUN) closely mimic the corresponding reduction rules. These transitions do not require interaction with the external environment, so the label $l \cdot \tau$ requires only the existence of loca-

tion l . Rule (LTS IN) provides a first-order output message from the environment which can be used to analyze the continuation of an input process by deriving a further transition using the communication rules. Rule (LTS OUT) states that a potentially higher-order output $\overline{l-c}\langle\tilde{v}\rangle$ evolves to the definitions that are extracted from \tilde{v} to obtain \tilde{v}' , and carries in the label the first-order version of the process. The intuition is that a bisimilar process will be required to perform the same first-order transition, and a potential incompatibility between the original higher-order messages will be detected by analyzing the resulting configurations (Θ^k for the first process). Rule (LTS TRIGGER) states that the application of a trigger name to the potentially higher-order parameters \tilde{v} evolves to the definitions that are extracted from \tilde{v} to obtain \tilde{v}' , and carries in the label the first-order version of the process, similarly to the case for output. Rule (LTS OPEN) is standard. Not that it applies to transitions originated using (LTS OUT) or (LTS TRIGGER). Rule (LTS REQ) can be interpreted as the combination of an output of p and an input of T on a special name `req`. Rule (LTS DEF) analyzes the script of a definition for all its possible (first-order) input parameters.¹¹ It is akin to performing an asynchronous input transition to receive the parameters for the script from the context, and a communication step to instantiate the script.

The sample query and update language. We conclude this section by extending `Sam` (Definition 2.2) to deal with trigger names, and showing that it respects our assumptions (Observation 3.12).

Definition 3.11 (Sam[#]) *The query language `Sam#` is defined as `Sam`, with the exception that trees can contain also trigger names in the same position as scripts (i.e. within $\langle-\rangle$). The extraction relation $\mathfrak{X}_{\mathcal{Q}}$, and the functions `scripts` and `triggers` are defined on `Sam#` by*

$$\frac{\mathfrak{X}(\mathbf{V}) = (\mathbf{U}; \Theta)}{\mathfrak{X}_{\mathcal{Q}}(\widehat{p}(\pi)\mathbf{V}) = (\widehat{p}(\pi)\mathbf{U}; \Theta)}$$

$$\text{scripts}(\widehat{p}(\pi)\mathbf{V}) = \text{scripts}(\mathbf{V}) \quad \text{triggers}(\widehat{p}(\pi)\mathbf{V}) = \text{triggers}(\mathbf{V})$$

Observation 3.12 (Properties of `Sam#`) (i) `Sam#` is script independent, and (ii) $\mathfrak{X}_{\mathcal{Q}}$ for `Sam#` respects Definition 3.9.

Proof. Point (i) follows by induction on the derivation of \mathfrak{E} . The idea is that query evaluation does not depend on the structure of scripts, and by rule (EVAL MATCH) only scripts coming from the query and the input tree can occur in the result and the output tree. Point (ii) follows by induction on the derivation of $\mathfrak{X}_{\mathcal{Q}}$. \square

¹¹Both in (LTS DEF) and (LTS IN) we do not need to consider higher-order values. This is due to the fact that the bisimilarity relation that we will consider turns out to be closed with respect to parallel composition with definitions (Theorem B.15), and hence already takes into account the effects of scripts received from the environment.

3.2.2 Domain bisimilarity

We introduce our bisimulation equivalence. The intuition is that, when two bisimilar processes are running in a location domain Λ , if a process makes an action α_l with $l \in \Lambda$ then the other one must be able to mimic it, possibly relying on the existence of other locations in Λ . Since the location domain can be extended to $\Lambda \cup \Lambda'$ by composing networks, we need to make sure that also the actions mentioning locations in Λ' are matched, this time within a larger relation parameterized by $\Lambda \cup \Lambda'$.

The definition of bisimilarity relies on the following derived transition relations.

Definition 3.13 (Derived Transition Relations) *Consider the lts defined in Figure 15. Given $l \in \Lambda$, we use the notation*

$$\xrightarrow{\tau}_{\Lambda} \stackrel{\text{def}}{=} l \cdot \tau \rightarrow; \quad l \cdot \tau \rightarrow_{\Lambda} \stackrel{\text{def}}{=} \tau^* \rightarrow_{\Lambda}; \quad \alpha_l \rightarrow_{\Lambda} \stackrel{\text{def}}{=} \tau^* \rightarrow_{\Lambda} \circ \alpha_l \rightarrow \circ \tau^* \rightarrow_{\Lambda} \quad \text{when } \alpha_l \neq l \cdot \tau.$$

Definition 3.14 (Domain Bisimilarity) *A family of symmetric relations on configurations (indexed with sets of locations) $\approx = \{\approx_{\Lambda} : \Lambda \subseteq \mathcal{L}\}$ is a domain bisimulation if $K \approx_{\Lambda} L$ and $K \xrightarrow{\alpha_l} K'$ implies:*

1. if $l \in \Lambda$ with $\text{rel}(\alpha_l, L)$ then $L \xrightarrow{\alpha_l} L'$ and $K' \approx_{\Lambda} L'$;
2. if $l \notin \Lambda$ then $K \approx_{\Lambda \cup \{l\}} L$.

Domain bisimilarity $\approx = \{\approx_{\Lambda} : \Lambda \subseteq \mathcal{L}\}$ is the (point-wise) largest domain bisimulation: if \approx is a domain bisimulation, then $\approx_{\Lambda} \subseteq \approx_{\Lambda}$ for all Λ . Two open processes \mathbf{P}, \mathbf{Q} are Λ -bisimilar if and only if for all closing substitutions σ , $\mathbf{P}\sigma \approx_{\Lambda} \mathbf{Q}\sigma$.

Remark 3.15 (Initial Elements) *To show $K \approx_{\Lambda} L$ for a specific Λ , we can exhibit a domain bisimulation $\approx = \{\approx_{\Delta} : \Delta \subseteq \mathcal{L}\}$ such that $K \approx_{\Lambda} L$ and \approx_{Δ} is the empty set for all Δ smaller than Λ .*

A proof that the largest domain bisimulation indeed exists in our non-standard setting can be found in [18]. Domain bisimilarity is a coinductive relation, preserved by structural congruence and monotonic in the domain Λ . Under the mild assumption that the query language does not depend on scripts (Definition 3.8), domain bisimilarity enjoys two important properties which make it a useful proof method for domain congruence:

- domain bisimilarity is a *congruence*, that is embedding open processes in extended contexts preserves bisimilarity;
- domain bisimilarity is a *sound approximation* of the domain congruence induced by request observables, that is if two processes are bisimilar then they are request-congruent.

These important properties are summarized by the theorem below.

Theorem 3.16 (Properties of Domain Bisimilarity) *Assuming a script-independent query language:*

1. If $K \approx_{\Lambda} L$, $K \equiv K'$ and $L \equiv L'$, then $K' \approx_{\Lambda} L'$.
2. For all sets of locations Λ, Λ' , if $\Lambda \subsetneq \Lambda'$ then $\approx_{\Lambda} \subsetneq \approx_{\Lambda'}$.
3. For all full process contexts $C \in \mathcal{K}_f$ (Definition 3.5), if $P \approx_{\Lambda} Q$ then $C[P] \approx_{\Lambda} C[Q]$.
4. For all Λ, P, Q (where P and Q have no free trigger names), if $P \approx_{\Lambda} Q$ then $P \sim^{\Lambda} Q$.

Proof. The proofs for (1) and (2) are in Appendix B.1. The proofs for (3) and (4) are respectively in Appendix B.2 and Appendix B.3. \square

We now give a first example of the proof method. Larger examples are given in Chapter 4.

Example 3.17 (Proof Method) *Recall the asynchrony law of Section 3.1. It states that a communication buffer cannot be distinguished from the empty process. By definition, $!\mathbf{l}\cdot\text{FW}(\mathbf{a}, \mathbf{a}, \tilde{\pi}) \approx_{\Lambda} \mathbf{0}$ if for any closing substitution σ , $(!\mathbf{l}\cdot\text{FW}(\mathbf{a}, \mathbf{a}, \tilde{\pi}))\sigma \approx_{\Lambda} \mathbf{0}\sigma$. Given an arbitrary σ we have that $(!\mathbf{l}\cdot\text{FW}(\mathbf{a}, \mathbf{a}, \tilde{\pi}))\sigma = !\mathbf{l}\cdot\text{FW}(a, a, \tilde{\pi})$ for some a, l . To show that $!\mathbf{l}\cdot\text{FW}(a, a, \tilde{\pi}) \approx_{\Lambda} \mathbf{0}$, we need to give a domain bisimulation $\approx = \{\approx_{\Delta}\}$ such that \approx_{Λ} contains the two processes. Since structural congruence preserves bisimilarity (Point 1 of Theorem 3.16), we reason up-to \equiv .*

For each Δ , we begin with a relation $\mathcal{R}_{\Delta}^0 = \{(!\mathbf{l}\cdot\text{FW}(a, a, \tilde{\pi}), \mathbf{0})\}$ containing the pair that we want to prove bisimilar. By definition of bisimilarity, we must close the relation under transitions. Due to (LTS IN) we must close the relation under parallel compositions with arbitrary output processes: $\mathcal{R}_{\Delta}^1 = \{(M \mid !\mathbf{l}\cdot\text{FW}(a, a, \tilde{\pi}), M)\}$ where $M = \prod_{0 \leq i \leq n} \overline{l_i} \cdot c_i \langle \tilde{v}_i \rangle$, $\text{dom}(M) \subseteq \Delta$ and $\text{scripts}(\tilde{v}_i) = \emptyset$ (note that $\mathcal{R}_{\Delta}^1 = \mathcal{R}_{\Delta}^0$ if $n = 0$). The possible tau transitions arising from the interaction of $!\mathbf{l}\cdot\text{FW}(a, a, \tilde{\pi})$ and an output $\overline{l} \cdot a \langle \tilde{v} \rangle$ where $\tilde{v} = \tilde{\pi}\sigma$ are already covered because by (LTS !COM) and (LTS STRUCT), $\overline{l} \cdot a \langle \tilde{v} \rangle \mid !\mathbf{l}\cdot\text{FW}(a, a, \tilde{\pi}) \xrightarrow{l \cdot \tau} \overline{l} \cdot a \langle \tilde{v} \rangle \mid !\mathbf{l}\cdot\text{FW}(a, a, \tilde{\pi})$. Again by definition of bisimilarity, we must make the relation symmetric, hence we conclude with $\approx_{\Delta} = \mathcal{R}_{\Delta}^1 \cup (\mathcal{R}_{\Delta}^1)^{-1}$.

Incompleteness. In general, domain bisimilarity is a more restrictive equivalence than domain congruence. The property is intrinsic to our choice of giving a proof method parametric in the chosen query and update language. In fact, without specializing the labelled transition system to a particular language, we are forced to distinguish request transitions as soon as queries are syntactically different. On the other hand, equivalences dependent on specific knowledge of

the semantics of queries would lead to optimizations which are no longer correct when the query language changes.¹²

Example 3.18 (Incompleteness) *Consider the query language $\text{Sam}^\#$ and the process definition*

$$X(\mathbf{a}, \mathbf{b}) \stackrel{\text{def}}{=} (\nu c)(\overline{l \cdot c} \mid !l \cdot c.(\nu e) \left(\begin{array}{l} l \cdot \text{req}(x)\mathbf{a}[]\langle e \rangle \mid \\ l \cdot e(x).(\nu e')(l \cdot \text{req}(x)\mathbf{b}[]\langle e' \rangle \mid l \cdot e'(x).\overline{l \cdot c}) \end{array} \right))$$

The process loops, replacing at each iteration whatever tree is at l first with $\mathbf{a}[]$ and then with $\mathbf{b}[]$. We have $X(\mathbf{a}, \mathbf{b}) \sim_{\{l\}} X(\mathbf{b}, \mathbf{a})$, because once the two processes are inserted in the same store, they can always reduce to each other. On the other hand, we have that $X(\mathbf{a}, \mathbf{b}) \not\approx_{\{l\}} X(\mathbf{b}, \mathbf{a})$ because the request transitions cannot be matched.

4 Distributed Query Patterns

In this Section, we use bisimilarity to study some communication patterns used by servers in distributed query systems to answer queries from clients. In *Core Xd π* , distributed queries take the form of processes which retrieve and combine data from different locations by using remote communication and local requests. We show that some existing patterns [29] can be combined together obtaining a flexible infrastructure which is provably equivalent to an intuitive specification of the intended behaviour. By exploiting process migration, we also propose a new communication pattern, and we show that it is behaviourally equivalent to a naive, less efficient one.

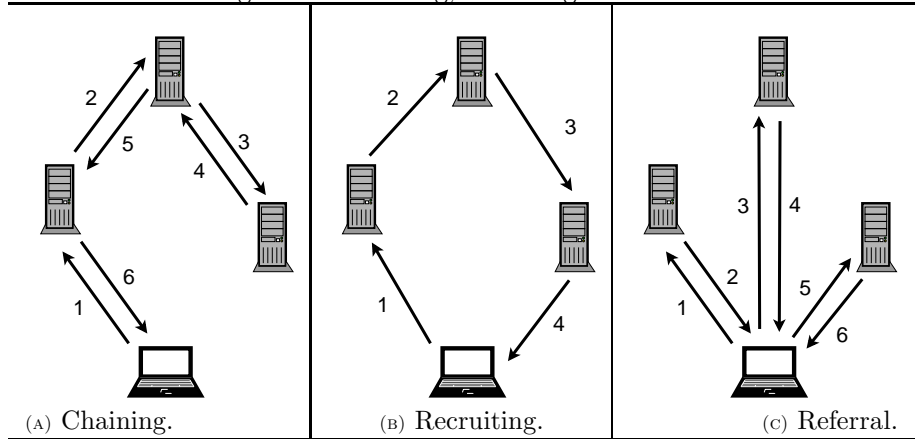
4.1 Chaining, recruiting and referral

We now consider *chaining*, *recruiting* and *referral*, three distributed query patterns studied by Sahuguet *et al.* in [29, 27] and described below. These patterns are interesting because, as will soon be apparent, they are simple yet can express ways of answering requests which are non-trivial, and display different levels of cooperation between the parties involved.

The usage of these patterns presupposes an architecture of servers sharing a common communication protocol for answering cooperatively the queries issued by clients. The protocol consists of alternative actions which depend on the contents of a query and on the local data, and is implemented by dedicated services running on each peer. The distributed querying infrastructure obtained

¹²Even if, for the sake of argument, we fixed a concrete query-update language and knew everything about its semantic equivalences, it would still be unclear how to deal with the case of Example 3.18. There, the initial updates that two processes can perform are by no means equivalent, yet by an “idempotence” argument the overall behaviours turn out to be equivalent. A complete bisimilarity would need to be able to consider in some way the cumulative effect of request transitions, also when interleaved with communication steps, in order to equate sequences of updates with the same global effect on the data-store.

Figure 16: Chaining, recruiting and referral



by combining the three query patterns is very flexible and can provide location independence to the clients. In fact, a client simply needs to invoke a service on a peer acting as the “entry point” to the network in order to get access to data which may reside on some other server unknown to the client itself.

We now describe the three patterns. In each case, a server receiving a query will try to execute it locally, and if that is not possible, will take alternative action.

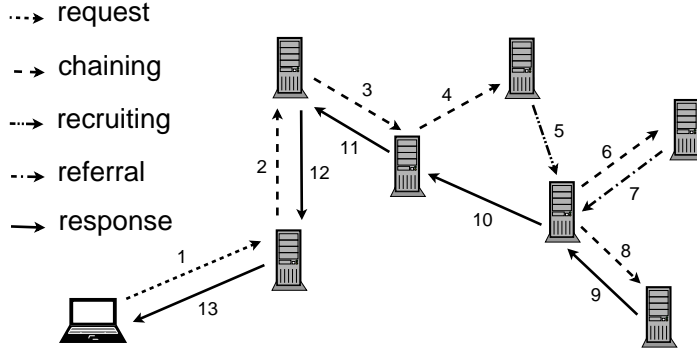
Chaining (Figure 16.(A)): if a server cannot deal directly with the call, it re-issues it to an alternative server, waits for the answer, and then returns the answer to the client.

Recruiting (Figure 16.(B)): if a server cannot deal directly with the call, it forwards it to another server (without noticing the client), so that the result will eventually return to the client without further intervention of the first server. To implement this pattern the address for returning the result must be a parameter of the call, and the client must be willing to accept asynchronous connections.

Referral (Figure 16.(C)): if a server cannot deal directly with the call, it suggests to the client an alternative server which might be able to. This strategy requires active collaboration from the client, which must be ready to contact the alternative server.

When each server involved in answering a request is able to use any of the patterns above, the flow of the data from the initial service call to the final answer can become complex and involve arbitrary combinations of the patterns, as in the example shown in Figure 17.

Figure 17: Combining the query patterns



4.1.1 Implementing the patterns

We now describe, step by step, some Core $Xd\pi$ code which implements a system where a client request can be answered by servers using an arbitrary combination of chaining, recruiting and referral. The code is based on services which retrieve and combine data from different locations by exploiting remote communication and local requests. In Section 2.6, we have seen how to represent service calls in Core $Xd\pi$:

$$(\nu c)(l.\overline{m}.a(\tilde{v}, l, c) \mid l.c(\tilde{\pi}).\mathbf{P})$$

where a is the name of the service to be invoked at location m with parameters \tilde{v} yielding a result to be passed on the channel c local to l , and \mathbf{P} is the code for handling the results, which are expected to match pattern $\tilde{\pi}$. In this section, a service call will carry four parameters: a tree T used to represent a condition, checked using pattern matching, that a server must satisfy in order to provide the right service (for example specifying the kind of result expected), a query p which is meant to be run on the store of the service matching tag T , and the return parameters m and c stating the location and the channel where the result should be returned. This approach can be easily applied also to service call having more parameters.

A client must be able to deal with the referral query pattern, therefore its code consists essentially of a loop. The loop consists of calling a first server (which could in principle provide the final result, terminating the loop), and then repeating the same call at the alternative addresses received in unsuccessful replies, until a reply containing the final result is received. The context defined below implements the loop at location m :

$$m.\text{Ref}_{(n,l,s,T,p,z)}[-] \stackrel{\text{def}}{=} (\nu c) \left(\frac{m.c(\text{OK}[], z). - \mid}{\overline{m.c}(\text{REF}[], l) \mid !m.c(\text{REF}[], x).m.\overline{x.s}(T, p, n, c)} \right)$$

It is parametric in the location n where the result must be returned, the location l of the first server to be interrogated, and the parameters of the call: the service

Figure 18: Syntax: abbreviations

$$\text{(INTERNAL CHOICE)} \quad \bigoplus_{i \in 1..n} {}^l P_i \stackrel{\text{def}}{=} (\nu c)(\overline{l \cdot c} \mid \prod_{i \in 1..n} l \cdot c.P_i) \quad c \notin \text{fn}(P_i)$$

$$\text{Notation: } \prod_{i \in 1..n} P_i \stackrel{\text{def}}{=} P_1 \mid \dots \mid P_n; \quad P_1 \oplus^l \dots \oplus^l P_n \stackrel{\text{def}}{=} \bigoplus_{i \in 1..n} {}^l P_i.$$

name s and condition T , the actual query p and the variable z for binding the result in the continuation. The context uses a private channel c to implement the referral loop and uses the tags $\text{OK}[]$ and $\text{REF}[]$ as guards to exit or continue the loop. Any process built using this context always starts the referral loop by invoking s at l and then waiting for two possible answers: either a referral message with the tag $\text{REF}[]$ and the name of an alternative location (bound to x), which starts another iteration of the loop against the corresponding server, or a result message with the tag $\text{OK}[]$ and the result of the service call (bound to z), which terminates the loop and passes the result on to the process which replaces the context hole “—”.

The server filters calls based on the parameter T in order to decide whether they can be served locally or not. Its code, which uses the conventions of Figure 18, consists of the following two processes, run in parallel:

$$l \cdot \text{Local}_{(s,T)} \stackrel{\text{def}}{=} !l \cdot s(T, x, y, z).(\nu c)(l \cdot \text{req}_x \langle c \rangle \mid l \cdot c(w).l \cdot \overline{y \cdot z} \langle \text{OK}[], w \rangle)$$

$$l \cdot \text{Remote}_{(s,\Delta)} \stackrel{\text{def}}{=} \prod_{(m, S_m) \in \Delta} !l \cdot s(S_m, x, y, z). \left(\begin{array}{l} l \cdot \overline{m \cdot s} \langle S_m, x, y, z \rangle \\ \oplus^l l \cdot \overline{y \cdot z} \langle \text{REF}[], m \rangle \\ \oplus^l l \cdot \text{Ref}_{(l, m, s, S_m, x, w)} [l \cdot \overline{y \cdot z} \langle \text{OK}[], w \rangle]. \end{array} \right)$$

If the first parameter matches T , the server runs the query (bound to x) on the local data and sends the result back to the client on channel z at y . If the first parameter does not match T , the server selects another server more appropriate for that request out of the set Δ relating servers to tags specifying their services (the outermost parallel composition of the remote process). It then invokes s on the chosen server using either chaining (third branch of the choice), or recruiting (first branch), or referral (second branch). In the case of chaining, the server runs the same code as the client with different parameters. Notice that the code handling the result forwards the result to the client instead of using it locally.

Installation. In order to use these patterns, the code implementing the services must be installed somehow on each participating server. We can assume that it is pre-installed on each peer, or we can install it “on demand” using either process migration or a specialized service which runs scripts. For example, consider the code of a service P parametric in the location x where the service is run and some other initialization pattern π . If we assume that an arbitrary location l exists then, given an arbitrary initialization parameter $v = \pi\sigma$, it is easy to see that running the code at m is equivalent to installing the service

code from l

$$\mathbf{P}\{m/x\}\sigma \sim^{\{l\}} l.\mathbf{go}\ m.(x, \pi)\mathbf{P} \circ \langle m, v \rangle.$$

Alternatively, one could use a dedicated installation service $Inst$ at location m which receives an abstraction and some parameters, and runs the abstraction locally

$$\mathbf{P}\{m/x\}\sigma \sim^{\{l\}} (\nu Inst)(l.\overline{m.Inst}\langle(x, \pi)\mathbf{P}, v\rangle \mid !m.Inst(y, z).y \circ \langle m, z \rangle).$$

4.1.2 Relating the patterns to a specification

We use a simple system with a client and two servers as an example of how to reason using our equivalences. The reasoning is analogous in the case of multiple servers. The client is on peer m , and runs the code

$$m.\mathbf{Client}_{(l,s,\mathbf{a}[]) } \stackrel{\text{def}}{=} m.\mathbf{Ref}_{(m,l,s,\mathbf{a}[],p,z)}[m.\mathbf{P}]$$

where the service is requested to match the tag $\mathbf{a}[]$, and the continuation process \mathbf{P} is an arbitrary process located at m which does not contain free occurrences of channel c mentioned in the definition of $\mathbf{Ref}_{(-)}$. A server is composed by the parallel composition of the branches dealing with local and remote processing, as described in Section 4.1.1:

$$l_1.\mathbf{Server}_{(s,T,l_2,S)} \stackrel{\text{def}}{=} l_1.\mathbf{Remote}_{(s,\{l_2,S\})} \mid l_1.\mathbf{Local}_{(s,T)}.$$

We consider two processes P_1 and P_2 , where a client requests from the server at l_1 the data specified by $\mathbf{a}[]$ (served locally) or $\mathbf{b}[]$ (served remotely at l_2).

$$\mathbf{Servers} \stackrel{\text{def}}{=} l_1.\mathbf{Server}_{(s,\mathbf{a}[],l_2,\mathbf{b}[]) } \mid l_2.\mathbf{Server}_{(s,\mathbf{b}[],l_1,\mathbf{a}[]) }$$

$$P_1 \stackrel{\text{def}}{=} m.\mathbf{Client}_{(l_1,s,\mathbf{a}[]) } \mid \mathbf{Servers}$$

$$P_2 \stackrel{\text{def}}{=} m.\mathbf{Client}_{(l_1,s,\mathbf{b}[]) } \mid \mathbf{Servers}$$

We compare P_1 and P_2 with Q_1 and Q_2 defined below, which provide a specification of the expected behaviour respectively of P_1 and P_2 . Each process goes directly from m to the relevant location, fetches the data returned by query p , and goes back to paste it as the new data tree of m :

$$m.\mathbf{Spec}_{(l)} \stackrel{\text{def}}{=} m.\mathbf{go}\ l.(\nu c)(l.\mathbf{req}_p\langle c\rangle \mid l.c(z).l.\mathbf{go}\ m.m.\mathbf{P})$$

$$Q_1 \stackrel{\text{def}}{=} m.\mathbf{Spec}_{(l_1)} \mid \mathbf{Servers}$$

$$Q_2 \stackrel{\text{def}}{=} m.\mathbf{Spec}_{(l_2)} \mid \mathbf{Servers}.$$

An important difference between the client and the specification is that the client sends an output message to a service which can in principle be intercepted by some process external to the protocol which performs an input on the same service channel. To rule out this undesired interference, we restrict the name

of the service s both in each P_i and Q_i , with the side effect of preventing also the unarmful case in which several clients use the services at the same time.¹³

We can show, in a domain containing both l_1 and l_2 , the following equivalences:

$$(\nu s)P_1 \sim^{\{l_1, l_2\}} (\nu s)Q_1 \quad (\nu s)P_2 \sim^{\{l_1, l_2\}} (\nu s)Q_2$$

Hence, by definition, we can replace $(\nu s)P_i$ by $(\nu s)Q_i$ in any network, and preserve network equivalence.

Proof of equivalence. By virtue of Theorem B.21, a formal proof of each equivalence above would involve showing the existence of an appropriate domain bisimulation containing the relevant pair, along the lines of the examples of Section 3.2.2.

We show here a sketch for the case for P_2 and Q_2 . In order to make the proof more manageable, we adopt the simplifying assumption that the query p does not contain scripts. The proof of the general case follows a similar structure. Moreover, we use implicitly the closure of bisimilarity under structural congruence.

We start by analyzing the non-input transitions of the two processes, and then we indicate how to build a domain bisimulation by pairing compatible states and dealing with input transitions. Consider $S_0 = (\nu s)Q_2$. By structural congruence, $S_0 \equiv m \cdot \text{Spec}_{(l_2)} \mid (\nu s)\text{Servers}$, which makes it easy to see that Servers does not have transitions, because of the restriction on s . Hence, we concentrate on $m \cdot \text{Spec}_{(l_2)}$. All it can do is a tau transition at l_2 corresponding to a migration step to reach a state S_1 , followed by a request transition at l_2 to become

$$S_2 = (\nu s)((\nu c)(\overline{l_2 \cdot c}(V) \mid l_2 \cdot c(z) \cdot l_2 \cdot \text{go } m \cdot m \cdot \mathbf{P}) \mid \text{Servers})$$

where V stands for a generic result obtained by the request. In turn, this process can only do a local communication followed by a migration (both tau transitions, respectively at l_2 and m) to become

$$S_3^V = (\nu s)((\nu c)(m \cdot \mathbf{P}\{V/z\}) \mid \text{Servers}).$$

Using the hypothesis $c \notin \text{fn}(\mathbf{P})$, we obtain

$$(\nu s)((\nu c)(m \cdot \mathbf{P}\{V/z\}) \mid \text{Servers}) \equiv (\nu s)(m \cdot \mathbf{P}\{V/z\} \mid \text{Servers}).$$

Hence, we have shown that the transition for S_0 that we need to match is

$$S_0 \xrightarrow{l_2 \cdot \text{req}(p)(V)} S_3^V = S_0 \xrightarrow{\tau} S_1 \xrightarrow{l_2 \cdot \text{req}(p)(V)} S_2^V \xrightarrow{\tau} S_3^V.$$

All we need to show now is that also starting from $(\nu s)P_2$, for each possible execution path, we can only do an analogous weak request transitions, and reach a state equivalent either to S_3 above. We now analyze the transitions of $S'_0 = (\nu s)P_2$. First the client $m \cdot \text{Client}_{(l_1, s, \mathbf{b}[\])}$ performs a tau transition at m

¹³In future, we plan to consider less restrictive ways to rule out this kind of interference using the type based techniques for linearity of Yoshida *et al.* [42].

corresponding to the initialization of the loop and then one at l_1 corresponding to the migration of the service call. The whole process becomes

$$S'_1 = (\nu s, c)(C_0 \mid \overline{l_1 \cdot s}(\mathbf{b}[], p, m, c) \mid \text{Servers})$$

where

$$C_0 = m \cdot s(\text{OK}[], z) \cdot m \cdot \mathbf{P} \mid !m \cdot s(\text{REF}[], x) \cdot m \cdot \overline{x \cdot s}(\mathbf{b}[], p, n, c).$$

Because of $\mathbf{b}[]$, $l_1 \cdot \text{Server}_{(s, \mathbf{a}[], l_2, \mathbf{b}[])}$ receives the call in the remote branch (the one for l_2 with tag $\mathbf{b}[]$). Nondeterministically, the process at l_1 evolves to either

$$\text{(recruiting)} \quad S'_2 = l_1 \cdot \text{Server}_{(s, \mathbf{a}[], l_2, \mathbf{b}[])} \mid l_1 \cdot \overline{l_2 \cdot s}(\mathbf{b}[], p, m, c)$$

$$\text{(referral)} \quad S'_3 = l_1 \cdot \text{Server}_{(s, \mathbf{a}[], l_2, \mathbf{b}[])} \mid l_1 \cdot \overline{m \cdot c}(\text{REF}[], l_2)$$

$$\text{(chaining)} \quad S'_4 = l_1 \cdot \text{Server}_{(s, \mathbf{a}[], l_2, \mathbf{b}[])} \mid l_1 \cdot \text{Ref}_{(l_1, l_2, s, \mathbf{b}[], p, w)}[l_1 \cdot \overline{m \cdot c}(\text{OK}[], w)].$$

Both the communication for receiving the call and the choice of the branch to execute are two tau transitions at l_1 , so $S'_1 \xrightarrow{\tau} S'_i$ for $i \in 2..4$. We now consider the transitions of each choice branch.

Recruiting. If recruiting is chosen, the server at l_1 performs a tau transition at l_2 corresponding to the forwarding of the client request, becoming

$$S'_{2,1} = (\nu s, c)(C_0 \mid R_1 \mid l_1 \cdot \text{Server}_{(s, \mathbf{a}[], l_2, \mathbf{b}[])}) \mid \overline{l_2 \cdot s}(\mathbf{b}[], p, m, c) \mid l_2 \cdot \text{Server}_{(s, \mathbf{b}[], l_1, \mathbf{a}[])}$$

where R_1 is a deadlocked process containing the code for the two discarded choice branches. Due to the parameter $\mathbf{b}[]$, server l_2 receives the call in the local branch, performing another tau transition at l_2 :

$$S'_{2,2}(\nu s, c)(\dots \mid l_2 \cdot \text{Server}_{(s, \mathbf{b}[], l_1, \mathbf{a}[])}) \mid (\nu c')(l_2 \cdot \text{req}_p(c') \mid l \cdot c'(w) \cdot l_2 \cdot \overline{m \cdot c}(\text{OK}[], w)).$$

This process can only do a request transition (say obtaining data V) to $S'_{2,3}$, followed by a tau transition (corresponding to local communication) at l_2 , becoming

$$S'_{2,4} = (\nu s, c)(C_0 \mid R_1 \mid \text{Servers} \mid l_2 \cdot \overline{m \cdot c}(\text{OK}[], V)),$$

where V stands for a generic result obtained by the request. After two tau transitions at m , corresponding to migration and local communication between C_0 and $\overline{m \cdot c}(\text{OK}[], V)$, we obtain

$$S'_{2,5} = (\nu s)(m \cdot \mathbf{P}\{V/x\} \mid (\nu c)(!m \cdot s(\text{REF}[], x) \cdot m \cdot \overline{x \cdot s}(\mathbf{b}[], p, n, c)) \mid R_1 \mid \text{Servers}),$$

where we stop, with

$$S'_2 \xrightarrow{\tau} S'_{2,2} \xrightarrow{l_2 \cdot \text{req}(p)(V)} S'_{2,3} \xrightarrow{\tau} S'_{2,5} V.$$

Referral. In the case of referral, the server performs two tau transitions at m which correspond to the forwarding of the message referring location l_2 , and to a second iteration of the referral loop of the client. The client then sends a new call to l_2 (a tau transition at l_2) and becomes

$$S'_{3,1} = (\nu s, c)(C_0 \mid R_2 \mid l_1 \cdot \text{Server}_{(s, a[], l_2, b[])}) \mid \overline{l_2} \cdot s \langle b[], p, m, c \rangle \mid l_2 \cdot \text{Server}_{(s, b[], l_1, a[])}$$

where R_2 is a deadlocked process containing the code for the two discarded choice branches. From now on, the transitions are the same as in the case for recruiting until we obtain the process

$$S_{3,4}^V = (\nu s)(m \cdot \mathbf{P}\{V/x\} \mid (\nu c)(!m \cdot s(\text{REF}[], x) \cdot m \cdot \overline{x} \cdot s \langle b[], p, n, c \rangle) \mid R_2 \mid \text{Servers}),$$

where we stop, with

$$S'_3 \xrightarrow{\tau} S'_{3,2} \xrightarrow{l_2 \cdot \text{req}(p)(V)} S_{3,3}^V \xrightarrow{\tau} S_{3,4}^V.$$

Chaining. In the case of chaining, the reasoning is similar. The first transitions correspond to a local communication at l_1 starting the referral loop of the server and a migration followed by communication at l_2 to start the local branch of that service. At this point, the process performs a request transition analogous to the one in the previous cases, and tau transitions corresponding to a local communication to get the result at l_2 , a migration to l_1 , and a local communication to terminate the referral loop of the server. The continuation process performs the transitions corresponding to the migration of the final result to m and to the communication to terminate the referral loop of the client, reaching

$$S_{4,3}^V = (\nu s)(m \cdot \mathbf{P}\{V/x\} \mid (\nu c)(!m \cdot s(\text{REF}[], x) \cdot m \cdot \overline{x} \cdot s \langle b[], p, n, c \rangle) \mid R_3 \mid \text{Servers}),$$

where R_3 is a deadlocked process containing the code for the two discarded choice branches and the residual of the referral loop at l_1 , and we stop, with

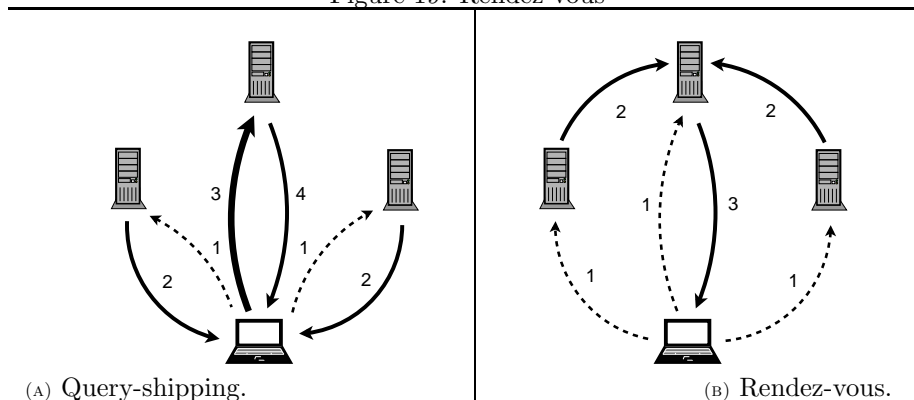
$$S'_4 \xrightarrow{\tau} S'_{4,1} \xrightarrow{l_2 \cdot \text{req}(p)(V)} S_{4,2}^V \xrightarrow{\tau} S_{4,3}^V.$$

Intuitively, S_0, S'_2, S'_3 and S'_4 are all equivalent states, because they are structurally equivalent to a process of the form

$$(\nu s)(m \cdot \mathbf{P}\{V/z\} \mid \text{Servers} \mid \delta).$$

The bisimulation relation we are looking for is obtained in three steps. First, we pair each of the states S_0, S_1 preceding the request transition in the lts of $(\nu s)Q_2$ with each $S'_0, \dots, S'_{2,1}, S'_{2,2}, \dots, S'_{4,1}$ preceding a request transition in the lts of $(\nu s)P_2$ (and vice versa). Second, we pair each of the states S_2^V, S_3^V following the request transition giving a particular result V in the lts of $(\nu s)Q_2$ with each $S_{i,j}^V$ following a request transition giving the same result in the lts of $(\nu s)P_2$ (and vice versa). Third, we close the relation obtained so far under parallel composition with the output messages derived by input transitions (as shown explicitly in Example 3.17). The relation defined above can be shown to be a domain bisimulation by formally checking the definition.

Figure 19: Rendez-vous



4.2 Rendez-vous and shipping

In the previous example, the infrastructure of servers implementing the distributed query patterns was fixed in advance, while the actual interactions between them were determined at run-time. The messages exchanged between different locations were always service calls or their results. Now, we consider a more flexible scenario which exploits code mobility.

Data-shipping and *query-shipping* are two traditional database techniques for distributed query evaluation: the first consists of evaluating locally a query on remote data by asking for the relevant data to be sent from the remote sources; the second consists of delegating the evaluation of a query to one of the remote sources in order to reduce the bandwidth used by data transfers. In the next section, we propose a distributed query pattern, called *rendez-vous*, which combines data and query shipping by using code mobility and private channels. The idea is to give a client the ability to ship result-handling code to another location, and to redirect the results of arbitrary service calls towards the location containing the result-handler. Within an infrastructure of services such as the one used above for chaining, recruiting and referral, this pattern can help to save bandwidth by eliminating unnecessary data transfers.

4.2.1 The rendez-vous query pattern

We now compare the query-shipping and rendez-vous patterns by giving a concrete example where a client calls a remote service using as parameters two large sets of data obtained by other remote service calls.

Suppose that on location l there is a specialized service $l\text{-Join}(x_1, x_2, y, z)$ which returns on channel z at location y the result of joining the data bound to x_1 with the data bound to x_2 . Suppose moreover that a client running on location m wants to join some data obtained by query p at location l_1 with other data obtained by query q at location l_2 . We assume that l_1 and l_2 run

the services described in Section 4.1.2, that l_1 (respectively l_2) serves locally the requests tagged by $\mathbf{a}[]$ (respectively $\mathbf{b}[]$).

Query shipping. The client can use query shipping: it first invokes the query services at locations l_1 and l_2 , then passes on the results as inputs to the join service on location l (see Figure 19(A)). Below we give the code of a client implementing this approach:

$$m\text{-ClientQ} \stackrel{\text{def}}{=} (\nu c, c_1, c_2) \left(\begin{array}{l} m\overline{l_1}\cdot\mathbf{s}\langle\mathbf{a}[], p, m, c_1\rangle \\ | m\overline{l_2}\cdot\mathbf{s}\langle\mathbf{b}[], q, m, c_2\rangle \\ | m\cdot c_1(\mathbf{OK}[], x_1).m\cdot c_2(\mathbf{OK}[], x_2).m\overline{l}\cdot\mathbf{Join}\langle x_1, x_2, m, c\rangle \\ | m\cdot c(z).m\cdot\mathbf{P} \end{array} \right)$$

It starts sending off the two service calls to l_1 and l_2 and then waits for the results respectively on c_1 and c_2 to bind them to x_1 and x_2 . The remaining code is a standard service call for the join service at l with parameters x_1 and x_2 , binding the final result to z in the continuation $m\cdot\mathbf{P}$, which can be an arbitrary process.

Rendez-vous. In order to save bandwidth, a better strategy is to request the query services at l_1 and l_2 to forward their results to location l , and to install at l a process which collects the two results and invokes the join service locally, asking for the final result to be returned at location m (see Figure 19(B)). Below we give a context implementing the general pattern, with two holes for inserting the code to handle the intermediate results at l and the final result at m . The code is parametric in the tags T_i and the queries p_i used to determine the partial results, the variables x_i for binding them in the intermediate code at “ $-_1$ ”, and the variable z for binding the final result in the continuation code at “ $-_2$ ”:

$$m\text{-RzV}_{(T_1, p_1, x_1, T_2, p_2, x_2, z)}[-]_1[-]_2 \stackrel{\text{def}}{=} (\nu c, c_1, c_2) \left(\begin{array}{l} m\overline{l_1}\cdot\mathbf{s}\langle T_1, p_1, l, c_1\rangle \\ | m\overline{l_2}\cdot\mathbf{s}\langle T_2, p_2, l, c_2\rangle \\ | m\cdot\mathbf{go}\ l.l\cdot c_1(\mathbf{OK}[], x_1).l\cdot c_2(\mathbf{OK}[], x_2).\text{-}_1 \\ | m\cdot c(z).\text{-}_2 \end{array} \right)$$

The code given above can be easily parameterized also on the number, the names and the locations of the services involved, and can be adapted to return the final results at an arbitrary location on an arbitrary channel.

We give below the code for a client, equivalent to ClientQ, which uses the rendez-vous strategy:

$$m\text{-ClientR} \stackrel{\text{def}}{=} m\text{-RzV}_{(\mathbf{a}[], p, x_1, \mathbf{b}[], q, x_2, z)}[\overline{l}\cdot\mathbf{Join}\langle x_1, x_2, m, c\rangle][m\cdot\mathbf{P}]$$

The code for handling the intermediate results consists in a local call to the join service, whereas the continuation is the same generic process used for ClientQ.

4.2.2 Equivalence of the patterns

Consider the process `Servers` defined in Section 4.1 consisting in the parallel compositions of the servers for implementing chaining, recruiting and referral at locations l_1 and l_2 . The clients given above, each in parallel with `Servers`, are equivalent in any network regardless of what locations are present.

$$(\nu s)(m \cdot \text{ClientQ} \mid \text{Servers}) \sim^\emptyset (\nu s)(m \cdot \text{ClientR} \mid \text{Servers})$$

In order to make the proof more manageable, we adopt once again the simplifying assumption that p and q do not contain scripts, and we use implicitly the closure of bisimilarity under structural congruence.

First of all, we simplify the problem further by studying an equation relating only the parts of the client processes above which are different from each other and which play a significant role in the proof. The full result follows by exploiting the closure of bisimilarity under parallel composition, restriction and structural congruence (Theorem B.15 and Proposition B.1) to recover the processes of the original statement.

Consider the definitions

$$m \cdot \text{ClientQ}' \stackrel{\text{def}}{=} (\nu c_1, c_2) \left(\begin{array}{l} m \cdot \overline{l_1} \cdot \mathbf{s} \langle \mathbf{a}[], p, m, c_1 \rangle \\ | m \cdot \overline{l_2} \cdot \mathbf{s} \langle \mathbf{b}[], q, m, c_2 \rangle \\ | m \cdot c_1(\text{OK}[], x_1) \cdot m \cdot c_2(\text{OK}[], x_2) \cdot m \cdot \overline{l} \cdot \mathbf{Join} \langle x_1, x_2, m, c \rangle \end{array} \right)$$

$$m \cdot \text{ClientR}' \stackrel{\text{def}}{=} (\nu c_1, c_2) \left(\begin{array}{l} m \cdot \overline{l_1} \cdot \mathbf{s} \langle \mathbf{a}[], p, l, c_1 \rangle \\ | m \cdot \overline{l_2} \cdot \mathbf{s} \langle \mathbf{b}[], q, l, c_2 \rangle \\ | m \cdot \text{go} \cdot l \cdot c_1(\text{OK}[], x_1) \cdot l \cdot c_2(\text{OK}[], x_2) \cdot \overline{l} \cdot \mathbf{Join} \langle x_1, x_2, m, c \rangle \end{array} \right)$$

Our goal is to build a domain bisimulation containing the pair

$$((\nu s)(m \cdot \text{ClientQ}' \mid \text{Servers}), (\nu s)(m \cdot \text{ClientR}' \mid \text{Servers})).$$

The construction of the proof is summarized by the diagrams in Figure 20. We represent the transitions of the two processes above in the form of lattices of states related by the *lts* (to be read in the direction of the arrows). Like in Section 4.1.2, we do not consider input transitions at this stage. The dotted arcs indicate the states from the two diagrams which will be paired in the bisimulation.


Building the transition diagrams. We describe the steps leading to the transitions. Later, we will explain how to build the bisimulation relation.

We begin with $(\nu s)(m \cdot \text{ClientQ}' \mid \text{Servers})$, corresponding to the top diagram of Figure 20. The starting state is the one pointed to by an arrow on the left of the diagram. We follow the top-left border of the diagram. Consider the sub-processes $m \cdot \overline{l_1} \cdot \mathbf{s} \langle \mathbf{a}[], p, m, c_1 \rangle$ and

$$l_1 \cdot \text{Local}_{(s, \mathbf{a}[]) } \stackrel{\text{def}}{=} !l_1 \cdot s \langle \mathbf{a}[], x, y, z \rangle \cdot (\nu c') (l_1 \cdot \text{req}_x \langle c' \rangle \mid l_1 \cdot c'(w) \cdot l_1 \cdot \overline{y} \cdot z \langle \text{OK}[], w \rangle).$$

Together they can perform, in order:

Figure 20: Bisimulation Diagrams



1. a migration step from m to l_1 ;
2. an internal communication on s at l_1 ;
3. a request transition generating an output $\overline{l_1 \cdot c'} \langle V_1 \rangle$, where V_1 is the data obtained by query p ;
4. an internal communication on c' ;

5. a migration to m .

We are left with the processes $l_1 \cdot \text{Local}_{(s, \mathbf{a}[\cdot])}$ and $\overline{m \cdot c_1} \langle \text{OK}[\cdot], V_1 \rangle$. The second process can communicate with

$$m \cdot c_1 \langle \text{OK}[\cdot], x_1 \rangle \cdot m \cdot c_2 \langle \text{OK}[\cdot], x_2 \rangle \cdot m \cdot \overline{l \cdot \text{Join}} \langle x_1, x_2, m, c \rangle,$$

and we are left with

$$m \cdot c_2 \langle \text{OK}[\cdot], x_2 \rangle \cdot m \cdot \overline{l \cdot \text{Join}} \langle V_1, x_2, m, c \rangle.$$

Independently, $m \cdot \overline{l_2 \cdot \mathbf{s}} \langle \mathbf{b}[\cdot], q, m, c_2 \rangle$ and $l_2 \cdot \text{Local}_{(s, \mathbf{b}[\cdot])}$ can mimic the 5 steps above (represented in the diagram by the bottom-left border), becoming $l_2 \cdot \text{Local}_{(s, \mathbf{b}[\cdot])}$ and $\overline{m \cdot c_2} \langle \text{OK}[\cdot], V_2 \rangle$. These two independent groups of respectively 6 and 5 ordered transitions give a lattice of 42 processes related by the lts, where the bottom element is the initial process, and the top element (at the intersection between the top and bottom-right borders) is the process

$$(\nu s) \left(\begin{array}{l} \overline{m \cdot c_2} \langle \text{OK}[\cdot], V_2 \rangle \\ | m \cdot c_2 \langle \text{OK}[\cdot], x_2 \rangle \cdot m \cdot \overline{l \cdot \text{Join}} \langle V_1, x_2, m, c \rangle \\ | \text{Servers} \end{array} \right)$$

This process can only perform a communication on channel c_2 and a migration from m to l , becoming the point on the right with the outgoing arrow

$$(\nu s) (\overline{l \cdot \text{Join}} \langle V_1, V_2, m, c \rangle \mid \text{Servers})$$

By a similar reasoning, we can derive for the process $(\nu s) (m \cdot \text{ClientQ}' \mid \text{Servers})$ the lattice of $36 + 42 + 1$ processes reported in the bottom diagram of Figure 20. The vertical transition possible from each of the 36 states in the lower layer of the diagram is the initial migration step from m to l of the result handling code, where

$$m \cdot \text{go } l \cdot l \cdot c_1 \langle \text{OK}[\cdot], x_1 \rangle \cdot l \cdot c_2 \langle \text{OK}[\cdot], x_2 \rangle \cdot \overline{l \cdot \text{Join}} \langle x_1, x_2, m, c \rangle$$

becomes process

$$l \cdot c_1 \langle \text{OK}[\cdot], x_1 \rangle \cdot l \cdot c_2 \langle \text{OK}[\cdot], x_2 \rangle \cdot \overline{l \cdot \text{Join}} \langle x_1, x_2, m, c \rangle.$$

Only when this transition has occurred, can communication on c_1 at l happen (hence the additional 7 states appearing only in the upper layer). The final state, after communication on c_2 at l has happened, is once again

$$(\nu s) (\overline{l \cdot \text{Join}} \langle V_1, V_2, m, c \rangle \mid \text{Servers}).$$

Building the bisimulation relation. We now describe how to pair-up the processes (states) of Figure 20 to build a suitable bisimulation. First, we relate each of the 25 processes in the top diagram reachable from the initial one after at most 4 transitions along each axis, with the two corresponding processes in the bottom diagram (as shown by the left-most dotted arc in Figure 20). Then,

we relate each of the 10 processes obtained in the top diagram after 5 transitions along one axis and at most 4 on the other axis with the corresponding 4 processes in the bottom diagram (as shown by the second dotted arc in Figure 20). Next, we relate each of the 5 processes obtained in the top diagram after 6 transitions along the first axis and at most 4 along the second axis, with the corresponding 5 processes in the bottom diagram (as shown by the third dotted arc in Figure 20). Then, we relate the two processes reachable after 11 and 12 transitions in the top diagram with the 10 processes of the bottom diagram to which they are joined by the fourth arc in Figure 20. We also relate the process obtained in the top diagram after 5 transitions along each axis with the 8 processes at the vertices of the corresponding cube in the bottom diagram. Finally, we associate the final process in the top diagram with the final process in the bottom diagram (the fifth dotted arc in Figure 20).

We take the symmetric closure of this relation, and we close it under parallel composition with output messages like in Example 3.17. Note that in the diagram we have shown the transitions for one possible choice of the data items V_1 and V_2 obtained as results of the request transition. To be completely formal, the reasoning above must be quantified on all possible result values, by pairing the corresponding states as in the example of Section 4.1.2.

Following a simulation step. We consider now an example to explain the rationale behind the pairing of states in the relation. We focus on the simulation of $(\nu s)(m \cdot \text{ClientQ}' \mid \text{Servers})$ by $(\nu s)(m \cdot \text{ClientR}' \mid \text{Servers})$ which is subtle. The other direction is straightforward.

In the top diagram, the process connected to the leftmost dotted arc can perform a weak transition $\xrightarrow{l_1 \cdot \text{req}(p)(V_1)} \gg_{\{l_1, m\}}$ to become the process connected to the third dotted arc. In the bottom diagram, the top process connected to the arc can simulate the step by performing a weak transition $\xrightarrow{l_1 \cdot \text{req}(p)(V_1)} \gg_{\{l_1\}}$ to become the top-left process of those connected to the bottom end of the third arc. The bottom process connected to the first arc cannot simulate the transition by reaching the same process, because that would involve using location l . Instead, it performs a transition $\xrightarrow{l_1 \cdot \text{req}(p)(V_1)} \gg_{\{l_1\}}$ to become the bottom-left process of those connected to the bottom end of the third arc, which is also in the relation. Also the states reached by the process on the top diagram by performing one or two transitions the less are related to the two states of the bottom diagram mentioned above. The idea is that the last two tau transitions at m cannot be matched by tau transitions at l , in fact they do not need to be matched at all, so the processes in the bottom diagram stay the same.

The association between the states of the two diagrams above is not completely straightforward because we are showing the most general result in which domain congruence holds for the empty domain (\sim^\emptyset). In this case, we had to make sure that for each transition between processes in the top diagram there was a corresponding transition in the second diagram (possibly null) which related two processes bisimilar to the original ones *involving only locations used also by the original transition*. The problems are due to the additional transi-

tions at l which are possible in the second diagram. If we tried to prove $\sim^{\{l\}}$ instead, the association between processes would have been straightforward.

5 Conclusions

We previously introduced $Xd\pi$ [11], a calculus for describing the interaction between processes and active data across distributed locations. In this paper, we have concentrated on Core $Xd\pi$, the explicitly-located version of $Xd\pi$, in order to study process equivalences. Both calculi, and their formal relationship, are studied in detail in [18]. In Section 3.1, we defined two contextual equivalences for Core $Xd\pi$ networks and processes. Network equivalence dictates when two networks can be considered indistinguishable by an observer looking at the interface between processes and local stores. Process equivalence is such that, when we place equivalent processes in equivalent network contexts, we obtain equivalent networks. Both equivalences are parametric with respect to the language used for querying and updating documents, and can be instantiated to specific cases. Contextual equivalences are difficult to use directly. In Section 3.2 we defined *domain bisimilarity*, a coinductive equivalence relation which entails process equivalence. Its non-standard definition, due to the fact that scripts (which can appear in data) are part of the values, and process equivalences are sensitive to the set of locations constituting the network. Our investigations required new reasoning techniques, as well as well-established ones. Domain bisimilarity is intrinsically incomplete, due to its being parametric on a query and update language.

An important design choice, enabling us to study how properties of data can be affected by process interaction, was to model data and processes at the same level of abstraction, rather than encoding data into processes, as customary in the π -calculus [20, 21, 22, 32]. Whilst such an encoding makes sense when using the π -calculus as a low level concurrency modelling language, it becomes a burden when reasoning about the coordination of higher-level processes. Our choice also gave us the opportunity to keep our language modular with respect to the choice of a query language, which can be easily adapted from the existing literature on XML [5].

We delegate migration control to external security checks that can be superimposed on the language. It would be interesting to include an explicit construct to constrain process migration at the location level, perhaps based on types, along the lines of [13]. We have indeed already considered such an extension, which is not included in the present work because it is of limited interest in the untyped setting. Migration has been included in Core $Xd\pi$ only to maintain a closer correspondence with $Xd\pi$, but is not necessary. In fact, each located action already contains information about where it is to be executed. For example, in the process below we can imagine that each input on a_i at l_i is followed by an implicit migration step $l_i \cdot \mathbf{go} \ l_{i+1}$ to the next location:

$$l_1 \cdot a_1(x) \cdot l_2 \cdot a_2(y) \cdot \overline{l_3} \cdot a_3(x, y),$$

intuitively corresponds to

$$l_1 \cdot a_1(x) \cdot l_1 \cdot \mathbf{go} \quad l_2 \cdot l_2 \cdot a_2(y) \cdot l_2 \cdot \mathbf{go} \quad l_3 \cdot \overline{l_3} \cdot a_3(x, y).$$

Overall, if explicit migration were to be discarded, the presentation of our model would be slightly simpler. Nevertheless, we have decided to remain with our original presentation because the encoding of $Xd\pi$ in $\text{Core } Xd\pi$ and its full abstraction (reported in [18]) are interesting results in their own right, confirming the informal thesis of [7] that locations can be encoded, without divergence, in the π -calculus with polyadic synchronization.

Using domain bisimilarity, in Section 4 we studied some communication patterns employed by servers in distributed query systems to answer queries from clients. Queries took the form of processes which retrieve and combine data from different locations by using remote communication and local requests. In particular, we considered chaining, recruiting and referral, three distributed query patterns studied in [29, 27] which are interesting because, despite their simplicity, they can express ways of answering requests which are non-trivial and display different levels of cooperation between the parties involved. By exploiting process migration, we have proposed the rendez-vous query pattern which can help to save bandwidth in certain applications. A challenging application for future work, is to extend the distributed query pattern examples to model a robust system where the servers return streams of results which can dry out or be restarted, and show its equivalence to a simpler, non-streaming specification. We believe that our techniques are suitable for this task, but we need automated tools and symbolic techniques (such as the *open bisimulation* of Sangiorgi [31], or up-to techniques [23] such as “up to confluent reduction”), to help producing manageable bisimilarity proofs. For example, already in the example given in Section 4.2.2, the bisimulation relation was difficult to represent succinctly because each state of one process could be related to several states of the other process.

Studying types for $\text{Core } Xd\pi$ is ongoing work. A type system would be useful to guarantee the absence of run-time errors, refine the behavioural equivalences, guarantee the conformance of data trees to schemas, and study security properties. Given the use of mobile code in our systems, in the absence of trust, we face the problem of protecting a host from a potentially malicious agent. This problem could be tackled by type-checking each agent dynamically entering a location [13] (possibly relying on the ability of a location to infer the type of the agent), or by using the Proof Carrying Code [24] approach (to send a migrating process along with its type), or by a combination of both techniques. A type system usually restricts the number of terms that are admissible in a calculus. Hence, the behavioural equivalences can become easier to verify (since there can be less-counter-examples), and some laws that do not hold in the untyped calculus become valid for the typed fragment. An obvious theoretical question arising from the definition of a type system for $\text{Core } Xd\pi$, is to understand how the behavioural equivalences are affected by typing, for example along the lines of [26, 12].

References

- [1] Serge Abiteboul, Peter Buneman, and Dan Suciu. *Data on the Web: from relations to semistructured data and XML*. Morgan Kaufmann, 2000.
- [2] Lucia Acciai and Michele Boreale. XPi: A typed process calculus for XML messaging. In *Proc. of FMOODS'05*, 2005.
- [3] Omar Benjelloun. Active XML: A data-centric perspective on Web services. Ph.D. Thesis, University of Paris XI, 2002.
- [4] Gavin Bierman and Peter Sewell. Iota: a concurrent XML scripting language with application to Home Area Networks. University of Cambridge Technical Report 557, 2003.
- [5] Angela Bonifati and Stefano Ceri. Comparative analysis of five XML query languages. *SIGMOD Record*, 29(1):68–91, 2000.
- [6] Allen Brown, Cosimo Laneve, and Greg Meredith. PiDuce: a process calculus with native XML datatypes. In *Proc. of WSFM'05*. LNCS, 2005.
- [7] Marco Carbone and Sergio Maffeis. On the expressive power of polyadic synchronisation in π -calculus. *Nordic Journal of Computing*, 10(2):70–98, 2003.
- [8] Giuseppe Castagna, Rocco De Nicola, and Daniele Varacca. Semantic subtyping for the pi-calculus. In *Proc. of LICS'05*, pages 92–101. IEEE Computer Society Press, 2005.
- [9] Murdoch J. Gabbay. The π -calculus in FM. In *Thirty-five years of Automath*, volume 28, pages 71–123. Kluwer Academic Press, 2003.
- [10] Murdoch J. Gabbay and Andrew M. Pitts. A new approach to abstract syntax with variable binding. *Formal Aspects of Computing*, 13:341–363, 2002. Special issue in honour of Rod Burstall.
- [11] Philippa Gardner and Sergio Maffeis. Modelling dynamic Web data. *Theoretical Computer Science*, 342:104–131, 2005.
- [12] Matthew Hennessy, Massimo Merro, and Julian Rathke. Towards a behavioural theory of access and mobility control in distributed systems. *Theoretical Computer Science*, 322(3):615–669, 2004.
- [13] Matthew Hennessy, Julian Rathke, and Nobuko Yoshida. safeDpi: A language for controlling mobile code. In *Proc. of FoSSaCS 2004*, volume 2987 of LNCS, pages 241–256. Springer Verlag, 2004.
- [14] Matthew Hennessy and James Riely. Resource access control in systems of mobile agents. In *HLCL '98*, volume 16.3 of *ENTCS*, pages 3–17. Elsevier Science Publishers, 1998.

- [15] Kohei Honda and Nobuko Yoshida. On reduction-based process semantics. *Theoretical Computer Science*, 151(2):437–486, 1995.
- [16] Haruo Hosoya and Benjamin C. Pierce. Xduce: A statically typed xml processing language. *ACM Transactions on Internet Technology*, 3(2):117–148, 2003.
- [17] Alan Jeffrey and Julian Rathke. Contextual equivalence for higher-order pi-calculus revisited. *Logical Methods in Computer Science*, 1(1:4), 2005.
- [18] Sergio Maffei. *Dynamic Web Data: A Process Algebraic Approach*. PhD thesis, Imperial College London, September 2005.
- [19] Sergio Maffei and Philippa Gardner. Behavioural equivalences for dynamic Web data. In *Proc. of Ifip WCC-TCS'04*, Kluwer Academic Press, pages 535–548, August 2004.
- [20] Robin Milner. *Communication and concurrency*. Prentice-Hall, Inc., 1989.
- [21] Robin Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, May 1999.
- [22] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I and II. *Information and Computation*, 100(1):1–40,41–77, September 1992.
- [23] Robin Milner and Davide Sangiorgi. Techniques of weak bisimulation up-to. In *Proc. of CONCUR '92*, volume 2987 of *LNCS*, pages 41–77. Springer-Verlag, 1992.
- [24] George Necula and Peter Lee. Safe, untrusted agents using proof-carrying code. In *Mobile Agents and Security*, pages 61–91. Springer Verlag, 1998.
- [25] David Neven, Thomas Schwentick, and Dan Suciu. Foundations of semistructured data. Dagstuhl Seminar Proceedings 05061. Available online from <http://drops.dagstuhl.de/portals/05061/>, 2005.
- [26] Benjamin C. Pierce and Davide Sangiorgi. Behavioral equivalence in the polymorphic pi-calculus. In *Proc. of POPL '97*, pages 242–255. ACM Press, 1997.
- [27] Arnaud Sahuguet. *ubQL: A Distributed Query Language to Program Distributed Query Systems*. PhD thesis, University of Pennsylvania, 2002.
- [28] Arnaud Sahuguet, Benjamin Pierce, and Val Tannen. Distributed Query Optimization: Can Mobile Agents Help? Unpublished draft, 2000.
- [29] Arnaud Sahuguet and Val Tannen. ubql, a language for programming distributed query systems. In *Proc. of webDB'01*, pages 37–42, 2001.

- [30] Davide Sangiorgi. Expressing mobility in process algebras: First-order and higher-order paradigms. PhD thesis, University of Edinburgh, 1992.
- [31] Davide Sangiorgi. A theory of bisimulation for the pi-calculus. In *Proc. of CONCUR '93*, pages 127–142, Springer Verlag, 1993.
- [32] Davide Sangiorgi and David Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
- [33] Serge Abiteboul, et al. Active XML primer. INRIA Futurs, GEMO Report number 275, 2003.
- [34] Tim Sheard and Simon Peyton Jones. Template meta-programming for haskell. In *Proc. of Haskell '02*, pages 1–16. ACM Press, 2002.
- [35] Walid Taha. Metaocaml: A compiled, type-safe multi-stage programming language. Available online from <http://www.cs.rice.edu/taha/MetaOCaml/>, 2001.
- [36] UDDI. Universal Description, Discovery, and Integration of Business for the Web (UDDI) 3.0. <http://www.uddi.org>, 2005.
- [37] W3C. XML Path Language (XPath) Version 1.0. <http://w3.org/TR/xpath>, 1999.
- [38] W3C. Extensible Markup Language (XML) 1.0 (2nd edition). <http://www.w3.org/TR/REC-xml.html>, 2000.
- [39] W3C. Web Services Description Language (WSDL) 1.1. <http://w3.org/TR/wsdl>, 2001.
- [40] W3C. Web Services Activity. <http://www.w3.org/2002/ws>, 2002.
- [41] W3C. Simple Object Access Protocol (SOAP) Version 1.2. <http://w3.org/TR/SOAP>, 2003.
- [42] Nobuko Yoshida, Martin Berger, and Kohei Honda. Strong normalisation in the π -calculus. *Information and Computation*, 191(2):145–202, 2004.

A Figures

Figure 21: Full structural congruence

$(\nu c)\mathbf{0} \equiv \mathbf{0}$	(CSTRUCT RES PNil)
$c \notin \text{fn}(P) \implies P \mid (\nu c)Q \equiv (\nu c)(P \mid Q)$	(CSTRUCT RES PPAR)
$(\nu c)(\nu d)P \equiv (\nu d)(\nu c)P$	(CSTRUCT RES PRES)
$P \mid (Q \mid Q') \equiv (P \mid Q) \mid Q'$	(CSTRUCT PAR ASSOC)
$P \equiv Q \implies Q \equiv P$	(CSTRUCT PAR COMM)
$P \mid \mathbf{0} \equiv P$	(CSTRUCT PAR ZERO)
$P \equiv Q \implies (\nu c)P \equiv (\nu c)Q$	(CSTRUCT CONG RES)
$P \equiv P' \implies P \mid Q \equiv P' \mid Q$	(CSTRUCT CONG PAR)
$P \equiv Q \implies l \cdot a(\tilde{\pi}).P \equiv l \cdot a(\tilde{\pi}).Q$	(CSTRUCT CONG IN)
$P \equiv Q \implies !l \cdot a(\tilde{\pi}).P \equiv !l \cdot a(\tilde{\pi}).Q$	(CSTRUCT CONG !IN)
$P \equiv Q \implies l \cdot \text{go } m.P \equiv l \cdot \text{go } m.Q$	(CSTRUCT CONG GO)
$P \equiv Q \implies (D, P) \equiv (D, Q)$	(CSTRUCT PROC)
$(\nu c)\mathbf{0} \equiv \mathbf{0}$	(STRUCT RES CNil)
$c \notin \text{fn}(K) \implies K \mid (\nu c)K' \equiv (\nu c)(K \mid K')$	(STRUCT RES CPAR)
$(\nu c)(\nu d)K \equiv (\nu d)(\nu c)K$	(STRUCT RES CRES)
$K \mid (K' \mid L) \equiv (K \mid K') \mid L$	(STRUCT PAR CASSOC)
$K \equiv K' \implies K' \equiv K$	(STRUCT PAR CCOMM)
$K \mid \mathbf{0} \equiv K$	(STRUCT PAR CZERO)
$K \equiv K' \implies (\nu c)K \equiv (\nu c)K'$	(STRUCT CONG CRES)
$K \equiv L \implies K \mid K' \equiv L \mid K'$	(STRUCT CONG CPAR)
$t \equiv t$	(CSTRUCT REFL)
$t \equiv t' \implies t' \equiv t$	(CSTRUCT SYMM)
$t \equiv t'' \text{ and } t'' \equiv t' \implies t \equiv t'$	(CSTRUCT TRANS)

Notation: t ranges over P or K .

Figure 22: Predicates wf , $distinct$ and functions dom , $cval$

$$\begin{array}{c}
 \frac{}{wf(\mathbf{0})} \quad \frac{}{wf(\overline{l \cdot c} \langle \tilde{v} \rangle)} \quad \frac{wf(\mathbf{P}) \quad wf(\mathbf{Q})}{wf(\mathbf{P} | \mathbf{Q})} \\
 \frac{wf(\mathbf{P})}{wf((\nu c)\mathbf{P})} \quad \frac{dom(\mathbf{P}) = \{l\} \quad wf(\mathbf{P})}{wf(l \cdot c(\tilde{\pi}) \cdot \mathbf{P})} \quad \frac{dom(\mathbf{P}) = \{l\} \quad wf(\mathbf{P})}{wf(!l \cdot c(\tilde{\pi}) \cdot \mathbf{P})} \\
 \frac{dom(\mathbf{P}) = \{m\} \quad wf(\mathbf{P})}{wf(l \cdot go \ m \cdot m \cdot \mathbf{P})} \quad wf(\mathbf{A} \circ \langle l, \tilde{v} \rangle) \quad wf(l \cdot req_p \langle c \rangle) \\
 \\
 \frac{distinct(\pi) \quad distinct(\tilde{\pi}) \quad fv(\pi) \cap fv(\tilde{\pi}) = \emptyset}{distinct(\pi, \tilde{\pi})} \quad \frac{distinct(\mathbf{A})}{distinct(\langle \mathbf{A} \rangle)} \\
 \frac{distinct(\mathbf{E}) \quad distinct(\mathbf{T}) \quad fv(\mathbf{E}) \cap fv(\mathbf{T}) = \emptyset}{distinct(\mathbf{E} | \mathbf{T})} \quad \frac{distinct(\mathbf{V})}{distinct(\mathbf{a}[\mathbf{V}])} \\
 distinct(\emptyset) \quad distinct(x) \quad \frac{fv(l) \cap fv(p) = \emptyset}{distinct(p @ l)} \\
 \\
 dom(D) = \text{domain of } D \quad dom(\mathbf{0}) = \emptyset \\
 dom(\mathbf{P} | \mathbf{Q}) = dom(\mathbf{P}) \cup dom(\mathbf{Q}) \quad dom((\nu c)\mathbf{P}) = dom(\mathbf{P}) \\
 dom(\overline{l \cdot c} \langle \tilde{v} \rangle) = \{l\} \quad dom(l \cdot c(\tilde{\pi}) \cdot \mathbf{P}) = \{l\} \\
 dom(!l \cdot c(\tilde{\pi}) \cdot \mathbf{P}) = \{l\} \quad dom(l \cdot go \ m \cdot m \cdot \mathbf{P}) = \{l\} \\
 dom(\mathbf{A} \circ \langle l, \tilde{v} \rangle) = \{l\} \quad dom(l \cdot req_p \langle c \rangle) = \{l\} \\
 \\
 dom(-) = \emptyset \quad dom(C_S[-] \uplus D) = dom(D) \cup dom(C_S[-]) \\
 \left. \begin{array}{l} dom(\mathbf{P} | C_{\mathcal{P}}[-]) \\ dom(C_{\mathcal{P}}[-] | \mathbf{P}) \end{array} \right\} = dom(\mathbf{P}) \cup dom(C_{\mathcal{P}}[-]) \quad dom((\nu c)C_{\mathcal{P}}[-]) = dom(C_{\mathcal{P}}[-]) \\
 \\
 \frac{t \in \mathcal{A} \cup \mathcal{Q}}{cval(t) = \{t\}} \quad cval(\mathbf{E} | \mathbf{T}) = cval(\mathbf{E}) \cup cval(\mathbf{T}) \\
 cval(\mathbf{a}[\mathbf{V}]) = cval(\mathbf{V}) \quad cval(\langle \mathbf{A} \rangle) = cval(\mathbf{A}) \\
 cval(p @ l) = cval(p) \quad cval(\emptyset) = cval(l) = cval(x) = \emptyset
 \end{array}$$

Figure 23: Free variables and free names for Core $\lambda d\pi$

$$\begin{array}{ll}
fv((x, \tilde{\pi})\mathbf{P}) = fv(\mathbf{P}) \setminus fv(x, \tilde{\pi}) & fv(\mathbf{0}) = \emptyset \\
fv(\mathbf{P} \mid \mathbf{Q}) = fv(\mathbf{P}) \cup fv(\mathbf{Q}) & fv((\nu c)\mathbf{P}) = fv(\mathbf{P}) \\
fv(\mathbf{l} \cdot \mathbf{c}(\tilde{\pi}).\mathbf{l} \cdot \mathbf{P}) = fv(\mathbf{l}) \cup fv(\mathbf{c}) \cup (fv(\mathbf{P}) \setminus fv(\tilde{\pi})) & fv(\overline{\mathbf{l}} \cdot \mathbf{c}(\tilde{\nu})) = fv(\mathbf{l}) \cup fv(\mathbf{c}) \cup fv(\tilde{\nu}) \\
fv(!\mathbf{l} \cdot \mathbf{c}(\tilde{\pi}).\mathbf{l} \cdot \mathbf{P}) = fv(\mathbf{l}) \cup fv(\mathbf{c}) \cup (fv(\mathbf{P}) \setminus fv(\tilde{\pi})) & fv(\mathbf{l} \cdot \mathbf{go} \ \mathbf{m} \cdot \mathbf{m} \cdot \mathbf{P}) = fv(\mathbf{l}) \cup fv(\mathbf{m}) \cup fv(\mathbf{P}) \\
fv(\mathbf{A} \circ \langle \mathbf{l}, \tilde{\nu} \rangle) = fv(\mathbf{A}) \cup fv(\mathbf{l}) \cup fv(\tilde{\nu}) & fv(\mathbf{l} \cdot \mathbf{req}_p \langle \mathbf{c} \rangle) = fv(\mathbf{l}) \cup fv(\mathbf{p}) \cup fv(\mathbf{c}) \\
\\
fn(\mathbf{0}) = \emptyset & fn(\mathbf{P} \mid \mathbf{Q}) = fn(\mathbf{P}) \cup fn(\mathbf{Q}) \\
fn((\nu c)\mathbf{P}) = fn(\mathbf{P}) \setminus \{c\} & fn(\overline{\mathbf{l}} \cdot \mathbf{c}(\tilde{\nu})) = fn(\mathbf{c}) \cup fn(\tilde{\nu}) \\
fn(\mathbf{l} \cdot \mathbf{c}(\tilde{\pi}).\mathbf{P}) = fn(\mathbf{c}) \cup fn(\mathbf{P}) & fn(!\mathbf{l} \cdot \mathbf{c}(\tilde{\pi}).\mathbf{P}) = fn(\mathbf{c}) \cup fn(\mathbf{P}) \\
fn(\mathbf{l} \cdot \mathbf{go} \ \mathbf{m} \cdot \mathbf{m} \cdot \mathbf{P}) = fn(\mathbf{P}) & fn(\mathbf{A} \circ \langle \mathbf{l}, \tilde{\nu} \rangle) = fn(\tilde{\nu}) \\
fn(\mathbf{l} \cdot \mathbf{req}_p \langle \mathbf{c} \rangle) = fn(\mathbf{c}) & fn((D, P)) = fn(P) \\
\\
fv(\mathbf{K} \mid \mathbf{K}') = fv(\mathbf{K}) \cup fv(\mathbf{K}') & fv((\nu c)\mathbf{K}) = fv(\mathbf{K}) \\
fv(\langle k \Leftarrow A \rangle) = \emptyset & fv(k) = \emptyset \\
\\
fn(\mathbf{K} \mid \mathbf{K}') = fn(\mathbf{K}) \cup fn(\mathbf{K}') & fn((\nu c)\mathbf{K}) = fn(\mathbf{K}) \\
fn(\langle k \Leftarrow A \rangle) = \{k\} & fn(k) = \{k\}
\end{array}$$

Figure 24: Functions *dom* and *scripts* for configurations

$$\begin{array}{lll}
dom(\mathbf{K} \mid \mathbf{K}') = dom(\mathbf{K}) \cup dom(\mathbf{K}') & dom((\nu c)\mathbf{K}) = dom(\mathbf{K}) & dom(\langle k \Leftarrow A \rangle) = \emptyset \\
\\
scripts(v, \tilde{\nu}) = scripts(v) \cup scripts(\tilde{\nu}) & scripts(\{v/x\}) = scripts(v) & \\
scripts(E \uparrow T) = scripts(E) \cup scripts(T) & scripts(\emptyset) = \emptyset & \\
scripts(\mathbf{a} \mid V) = scripts(V) & scripts(\langle A \rangle) = \{A\} & \\
scripts(A) = \{A\} & scripts(c) = scripts(c) = \emptyset & \\
scripts(p \textcircled{a} l) = scripts(p) & scripts(l) = \emptyset &
\end{array}$$

Assumption: *scripts* on queries is given as part of the query language definition.

B Results and proofs

This section gives the formal proofs of the properties of domain bisimilarity.

To follow more easily certain common steps in the proofs, it may be helpful to keep in mind that: private and service channel names are distinct; a script is well-formed only if it has no free private channel or trigger names; configurations are well-formed if for any trigger name there is at most one definition.

B.1 Basic properties

We study some basic properties of domain bisimilarity which will be useful to prove the main results of congruence and soundness. A first property is that structural congruence preserves bisimilarity. We will use this implicitly in the rest of the section.

Proposition B.1 (Bisimilarity Up-To Structural Congruence) *If $K \approx_\Lambda L$, $K \equiv K'$ and $L \equiv L'$, then $K' \approx_\Lambda L'$.*

Proof. The family of relations \approx with generic element

$$\approx_\Delta = \{(K', L') : K' \equiv K \approx_\Delta L \equiv L'\}$$

is a domain bisimulation. Follows by using rule (LTS STRUCT). \square

By definition of bisimilarity, the smaller the domain Λ , the less likely that two processes are bisimilar. In fact, we need to check for matching actions first in Λ , then in any Λ' containing Λ . The underlying intuition is that if we can rely on a larger set of locations to be connected to the network, then we can perform more optimizations.

Proposition B.2 (Monotonicity) *Domain bisimilarity is monotonic: for all sets of locations Λ, Λ' , if $\Lambda \subsetneq \Lambda'$ then $\approx_\Lambda \subsetneq \approx_{\Lambda'}$.*

Proof.

(\subseteq) Follows by Definition 3.14, noticing that using rule (LTS IN) it is always possible to make an input action at location l , for any l not in Λ .

(\subsetneq) If $\Lambda \subsetneq \Lambda'$ then there exists an m such that $m \in \Lambda' \setminus \Lambda$. Consider the two processes $P = \bar{l} \cdot a$ and $Q = l \cdot \mathbf{go} \ m \cdot \bar{l} \cdot a$, where $l \neq m$. Clearly, $P \not\approx_\Lambda Q$ because $P \xrightarrow{\bar{l} \cdot a} \mathbf{0}$ but there is no Q' such that $Q \xrightarrow{\bar{l} \cdot a} Q'$. To show that $P \approx_{\Lambda'} Q$, let \mathcal{R}_Δ be the set containing the pairs

$$(M | P, M | Q), \quad (M | P, M | m \cdot \bar{l} \cdot a), \quad (M, M)$$

for any M of the form

$$\prod_{0 \leq i \leq n} \bar{l}_i \cdot c_i \langle \tilde{v}_i \rangle, \quad \text{dom}(M) \subseteq \Delta$$

where $\text{scripts}(\tilde{v}_i) = \emptyset$ for all i . The family $\tilde{\approx}$, where $\tilde{\approx}_\Delta = \mathcal{R}_\Delta \cup (\mathcal{R}_\Delta)^{-1}$ for each Δ containing Λ' and $\tilde{\approx}_\Delta = \emptyset$ otherwise, is a domain bisimulation containing (P, Q) , hence $P \approx_{\Lambda'} Q$.

□

The lemma given below is a standard technical lemma relating the transitions in the lts with the syntactic structure of configurations, up-to structural congruence. It is used in many proofs, sometimes implicitly.

Lemma B.3 (Transition Correspondence) *The transitions of the lts are in close correspondence with the structure of configurations.*

1. $K \xrightarrow{(\tilde{a}, \tilde{k})\overline{l \cdot c}(\tilde{v})} K'$ if and only if $K \equiv (\nu \tilde{a})(L | \overline{l \cdot c}(\tilde{v}'))$ where $c \notin \{\tilde{a}\}$, $\{\tilde{a}\} \subseteq \text{fn}(\tilde{v}')$ and $K' \equiv L | \Theta^{\tilde{k}}$ where $\mathfrak{X}(\tilde{v}') = (\tilde{v}; \Theta^{\tilde{k}})$.
2. $K \xrightarrow{l \cdot c(\tilde{v})} K'$ if and only if $K' \equiv K | \overline{l \cdot c}(\tilde{v})$, $\text{scripts}(\tilde{v}) = \emptyset$ and $\text{rel}(l \cdot c(\tilde{v}), K)$.
3. $K \xrightarrow{l \cdot \tau} K'$ if and only if
 - $K \equiv (\nu \tilde{a})(L | l \cdot c(\tilde{\pi}) \cdot \mathbf{P} | \overline{l \cdot c}(\tilde{\pi}\sigma))$ and $K' \equiv (\nu \tilde{a})(L | \mathbf{P}\sigma)$, or
 - $K \equiv (\nu \tilde{a})(L | !l \cdot c(\tilde{\pi}) \cdot \mathbf{P} | \overline{l \cdot c}(\tilde{\pi}\sigma))$ and $K' \equiv (\nu \tilde{a})(L | !l \cdot c(\tilde{\pi}) \cdot \mathbf{P} | \mathbf{P}\sigma)$, or
 - $K \equiv (\nu \tilde{a})(L | (x, \tilde{\pi})\mathbf{P} \circ \langle l, \tilde{\pi}\sigma \rangle)$ and $K' \equiv (\nu \tilde{a})(L | \mathbf{P}\{l/x\}\sigma)$, or
 - $K \equiv (\nu \tilde{a})(L | m \cdot \text{go } l.P)$ and $K' \equiv (\nu \tilde{a})(L | P)$.
4. $K \xrightarrow{(\tilde{k})l \cdot \text{req}(p)(T)} K'$ if and only if $K \equiv (\nu a)(L | l \cdot \text{req}_{p'}(c))$ and $K' \equiv (\nu a)(L | \overline{l \cdot c}(T) | \Theta^{\tilde{k}})$ for some p' such that $\mathfrak{X}_{\mathcal{Q}}(p') = (p, \Theta^{\tilde{k}})$ and some T such that $\text{scripts}(T) = \emptyset$ and T has the form $\mathfrak{r}[U_1] \dots \mathfrak{r}[U_n]$.
5. $K \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} K'$ if and only if $K \equiv L | \langle k \Leftarrow (x, \tilde{\pi})\mathbf{P} \rangle$ and $K' \equiv K | \langle k \Leftarrow (x, \tilde{\pi})\mathbf{P} \rangle | \mathbf{P}\{l/x\}\sigma$ and $\text{scripts}(\sigma) = \emptyset$.
6. $K \xrightarrow{(\tilde{a}, \tilde{k})\overline{l \cdot j}(\tilde{v})} K'$ if and only if $K \equiv (\nu \tilde{a})(L | j \circ \langle l, \tilde{v}' \rangle)$ where $j \notin \{\tilde{k}\}$, $\{\tilde{a}\} \subseteq \text{fn}(\tilde{v}')$ and $K' \equiv L | \Theta^{\tilde{k}}$ where $\mathfrak{X}(\tilde{v}') = (\tilde{v}; \Theta^{\tilde{k}})$.

Proof.

(\Leftarrow) Follows easily by definition of lts.

(\Rightarrow) By induction on the depth n of the derivation tree in the premise for the labelled transition. We give the case for bound output as an example (point 1).

($n = 0$) Suppose $K \xrightarrow{(\tilde{a}, \tilde{k})\overline{l \cdot c}(\tilde{v})} K'$ is derived by directly applying (LTS OUT). It must be the case that $K = \overline{l \cdot c}(\tilde{v}')$, where $\mathfrak{X}(\tilde{v}') = (\tilde{v}; \Theta^{\tilde{k}})$ and $K' = \Theta^{\tilde{k}}$.

($n = m + 1$) A derivation of depth $m + 1$ for $K \xrightarrow{(\tilde{a}, \tilde{k})\overline{l \cdot c}(\tilde{v})} K'$ must be obtained by applying one of the rules (LTS RES), (LTS PAR), (LTS STRUCT) or (LTS OPEN) to a derivation of depth m . The case for (STRUCT) or (RES) is applied then we must have that $(\nu d)K_1 \xrightarrow{(\tilde{a}, \tilde{k})\overline{l \cdot c}(\tilde{v})} (\nu d)K'_1$, where $d \notin n((\tilde{a}, \tilde{k})\overline{l \cdot c}(\tilde{v}))$, follows from the premise $K_1 \xrightarrow{(\tilde{a}, \tilde{k})\overline{l \cdot c}(\tilde{v})} K'_1$. By inductive hypothesis, $K_1 \equiv (\nu \tilde{a})(L | \overline{l \cdot c}(\tilde{v}'))$ where $c \notin \{\tilde{a}\}$, $\{\tilde{a}\} \subseteq \text{fn}(\tilde{v}')$ and $K'_1 \equiv L | \Theta^{\tilde{k}}$ where $\mathfrak{X}(\tilde{v}') = (\tilde{v}; \Theta^{\tilde{k}})$. Since \mathfrak{X} does not affect session channels, $d \notin \text{fn}(\tilde{v}') \setminus \{\tilde{a}\}$. By structural congruence, $(\nu d)K_1 \equiv (\nu \tilde{a})(\nu d)(L | \overline{l \cdot c}(\tilde{v}'))$, and $(\nu d)K'_1 \equiv (\nu d)(L | \Theta^{\tilde{k}})$. The cases for (PAR) and (OPEN) are similar. □

The next step towards the main proofs consists of generalizing the variant lemma of [15] to bijective substitutions (here called *switchings*) on channel and trigger names.¹⁴ By using switchings rather than generic substitutions we obtain a purely coinductive proof. Below we let a, b, c range over channel or trigger names, and we consider only well-sorted substitutions (replacing channels for channels and triggers for triggers).

Definition B.4 (Switching) *Given a term t with a function fn returning its free names, a switching $a \frown b$ is a bijective substitution $\{c/a, a/b\} \{b/c\}$ such that $c \notin \text{fn}(t) \cup \{a, b\}$. We denote by $\tilde{a} \frown \tilde{b}$ the switching $a_1 \frown b_1 \dots a_n \frown b_n$ where both \tilde{a} and \tilde{b} are vectors of distinct names.*

Observation B.5 (Switching Properties) *Switching is self-dual $K = (K^{a \frown b})^{a \frown b}$ and symmetric $K^{a \frown b} = K^{b \frown a}$.*

Proof. Follows from the definition of switching and substitution. □

We note below that both the extraction and the transition relations do not depend on specific names, hence they are fully compatible with switching, and α -conversion.

Lemma B.6 (Switching Extraction) *Extraction preserves switching: if $\mathfrak{X}(\tilde{v}) = (\tilde{v}', \Theta^{\tilde{k}})$ then $\mathfrak{X}(\tilde{v}^{a \frown b}) = (\tilde{v}'^{a \frown b}; (\Theta^{\tilde{k}})^{a \frown b})$.*

¹⁴Our approach is reminiscent of the permutation-based approach to abstract syntax developed by Gabbay and Pitts [10]. In particular, it may be interesting in future work to compare our use of switchings with the work of Gabbay [9] on the π -calculus.

Proof. By a simple induction on the derivation of $\mathfrak{X}(\tilde{v}) = (\tilde{v}', \Theta^{\tilde{k}})$, where the base case for queries uses Definition 3.9. \square

Lemma B.7 (Switching Transitions) *If $K \xrightarrow{\alpha_l} K'$ then $K^{a\smallfrown b} \xrightarrow{\alpha_l^{a\smallfrown b}} K'^{a\smallfrown b}$, provided that $\text{bn}(\alpha_l) \cap \text{fn}(K^{a\smallfrown b}) \cap \{a, b\} = \emptyset$.*

Proof. By case analysis on α_l , using Lemma B.3. We show the case for the bound output when the name of the channel used for output occurs in the switching; the other cases are simpler. Let $\rho =^{a\smallfrown b}$. Without loss of generality, suppose that $K \xrightarrow{\alpha_l} K'$ where $\alpha_l = (b, \tilde{b})\overline{l.a}(\tilde{v}')$ and b does not occur in \tilde{b} . By Lemma B.3, $K \equiv (\nu b, \tilde{b})(L | \overline{l.a}(\tilde{v}))$ where $a \notin \{b, \tilde{b}\}$, $\{b, \tilde{b}\} \subseteq \text{fn}(\tilde{v})$ and $K' \equiv L | \Theta^{\tilde{k}}$ where $\mathfrak{X}(\tilde{v}) = (\tilde{v}', \Theta^{\tilde{k}})$. Let c be a fresh name. By α -conversion and Observation B.5, we have $K \equiv (\nu c, \tilde{b})(L\{c/b\} | \overline{l.a}(\tilde{v}\{c/b\}))\rho$. Applying the inner switching, we obtain $K \equiv (\nu c, \tilde{b})(L\{c/b\}\rho | \overline{l.b}(\tilde{v}\{c/b\}\rho))\rho$. Since $\{c/b\}$ has replaced b with c ,

$$K \equiv (\nu c, \tilde{b})(L\{c/b\}\{b/a\} | \overline{l.b}(\tilde{v}\{c/b\}\{b/a\}))\rho$$

Since a does not appear free anymore, we can alpha-convert c with a in the term above, obtaining

$$K \equiv (\nu a, \tilde{b})(L\{c/b\}\{b/a\}\{a/c\} | \overline{l.b}(\tilde{v}\{c/b\}\{b/a\}\{a/c\}))\rho$$

By definition of switching, $K \equiv (\nu a, \tilde{b})(L\rho | \overline{l.b}(\tilde{v}\rho))\rho$. By Observation B.5, $K\rho \equiv (\nu a, \tilde{b})(L\rho | \overline{l.b}(\tilde{v}\rho))$. By Lemma B.6, $\mathfrak{X}(\tilde{v}\rho) = (\tilde{v}'\rho; \Theta^{\tilde{k}}\rho)$. By (LTS OUT), $K\rho \xrightarrow{\alpha_l\rho} K'\rho$. \square

The lemma below shows that bisimilarity is closed with respect to switchings, a property needed to show that it is transitive.

Lemma B.8 (Variant) *(i) If $K \approx_{\Lambda} L$ then $K^{a\smallfrown b} \approx_{\Lambda} L^{a\smallfrown b}$. (ii) If $b \notin \text{fn}(K, L)$ then $K \approx_{\Lambda} L \implies K\{b/a\} \approx_{\Lambda} L\{b/a\}$.*

Proof. (i) Let $\rho =^{a\smallfrown b}$. We show that the family \approx with generic element $\approx_{\Delta} = \{(K\rho, L\rho) : K \approx_{\Delta} L\}$ is a domain bisimulation. Assume $K \approx_{\Delta} L$ for some Δ . Suppose $K\rho \xrightarrow{\alpha_l} K'$ and $l \notin \Delta$. By $K \approx_{\Delta} L$, $K \approx_{\Delta \cup \{l\}} L$, hence $(K\rho, L\rho) \in \approx_{\Delta \cup \{l\}}$. Suppose instead $l \in \Delta$ and $\text{rel}(\alpha_l, L\rho)$. By Lemma B.7, $K \xrightarrow{\alpha_l\rho} K'\rho$ with $\text{rel}(\alpha_l\rho, L)$. By bisimilarity, $L \xrightarrow{\alpha_l\rho} L'$ with $K'\rho \approx_{\Delta} L'$. By Lemma B.7, $L\rho \xrightarrow{\alpha_l} L'\rho$. By definition, $(K'\rho\rho, L'\rho) \in \approx_{\Delta}$. We conclude because, by Observation B.5, $K'\rho\rho = K'$. (ii) Follows from (i) by definition of switching. \square

Domain bisimilarity is an equivalence relation. This property is very important, because in the rest of the proofs in this section we will often rely on symmetry and transitivity.

Proposition B.9 (Equivalence) *Domain bisimilarity is an equivalence relation.*

Proof. Reflexivity and symmetry are immediate. Transitivity states that if $K \approx_{\Delta} M$ and $M \approx_{\Delta} L$ then $K \approx_{\Delta} L$. We show that the family \approx with generic element

$$\approx_{\Delta} = \{(K, L) : K \approx_{\Delta} M, M \approx_{\Delta} L\}$$

is a domain bisimulation. Let Δ be arbitrary and suppose $K \xrightarrow{\alpha_l} K'$ with $l \notin \Delta$. By Definition 3.14, $K \approx_{\Delta \cup \{l\}} M$ and $M \approx_{\Delta \cup \{l\}} L$, hence $(K, L) \in \approx_{\Delta \cup \{l\}}$. If $l \in \Delta$ and $rel(\alpha_l, L)$ then there are two cases, determined by the relevance of α_l to M . If $rel(\alpha_l, M)$, the proof is straightforward. If α_l is not relevant to M , the action α_l must necessarily have some bound names \tilde{c} such that $\{\tilde{c}\} \subseteq fn(M)$. By the second premise of (LTS PAR) used to derive the bound transition, $\{\tilde{c}\} \cap fn(K) = \emptyset$. Let \tilde{a} have the same length as \tilde{c} , and be such that $\{\tilde{a}\} \cap fn(K, L, M) = \emptyset$. By Lemma B.8, $K = K\{\tilde{a}/\tilde{c}\} \approx_{\Delta} M\{\tilde{a}/\tilde{c}\} = M'$. By the same argument, $M' \approx_{\Delta} L$. Since now $rel(\alpha_l, M')$, the proof is straightforward. \square

B.2 Congruence

Our next objective is to show that domain bisimilarity is a congruence. We already know that it is an equivalence relation (Proposition B.9), and that it preserves switchings (Lemma B.8). We also know how to relate labelled transitions with the syntactic structure of configurations (Lemma B.3). Using these tools, it is pretty easy to show that bisimilarity is closed under the restriction operator and, for processes, under prefixes. Most of the work in this section is dedicated to show *directly* closure under parallel composition. In contrast, Jeffrey and Rathke [17] for example show the soundness of their bisimilarity with respect to barbed congruence by using an auxiliary reduction-closed relation, which is closed under parallel composition. Showing the corresponding completeness result, they derive that also bisimilarity is closed under parallel composition. We cannot use their approach due to the inherent incompleteness of domain bisimilarity (see Section 3.2.2).

Lemma B.10 (Restriction) *Bisimilarity is closed under restriction: $K \approx_{\Delta} L \implies (\nu \tilde{c})K \approx_{\Delta} (\nu \tilde{c})L$.*

Proof. The family \approx with generic element

$$\approx_{\Delta} = \{(K_1, L_1) : K_1 \equiv (\nu \tilde{c})K, L_1 \equiv (\nu \tilde{c})L, K \approx_{\Delta} L\}$$

is a domain bisimulation. Suppose $K_1 \xrightarrow{\alpha_l} K'_1$. The proof is by cases on α_l using Lemma B.3. We show the case for $\alpha_l = l \cdot c(\tilde{v})$ which is the most interesting. If $l \notin \Delta$, the proof is easy. Suppose $l \in \Delta$. By Lemma B.3, $K'_1 \equiv K_1 | \overline{l \cdot c}(\tilde{v})$. By hypothesis, $K_1 \equiv (\nu \tilde{c})K$, $K \approx_{\Delta} L$ and $L_1 \equiv (\nu \tilde{c})L$. By α -conversion, $K_1 \equiv$

Figure 25: Merge operator for configurations

$\langle\langle P \rangle\rangle = P$	(MERGE PROC)
$\langle\langle (\nu c)K \rangle\rangle = (\nu c)\langle\langle K \rangle\rangle$	(MERGE RES)
$\langle\langle K \mid \langle k \Leftarrow A \rangle \rangle\rangle = \langle\langle K \langle k \Leftarrow A \rangle \rangle\rangle$	(MERGE DEF)

$(\nu \tilde{c}')K\{\tilde{c}'/\tilde{c}\}$ for a fresh tuple of names \tilde{c}' . By (LTS STRUCT), (LTS PAR) and (LTS IN), $(\nu \tilde{c}')K\{\tilde{c}'/\tilde{c}\} \xrightarrow{l \cdot c(\tilde{v})} (\nu \tilde{c}')(K\{\tilde{c}'/\tilde{c}\} \mid \overline{l \cdot c}(\tilde{v}))$ and $K\{\tilde{c}'/\tilde{c}\} \xrightarrow{l \cdot c(\tilde{v})} K\{\tilde{c}'/\tilde{c}\} \mid \overline{l \cdot c}(\tilde{v})$. By Lemma B.8, $K\{\tilde{c}'/\tilde{c}\} \approx_{\Delta} L\{\tilde{c}'/\tilde{c}\}$, hence $L\{\tilde{c}'/\tilde{c}\} \xrightarrow{l \cdot c(\tilde{v})} \Delta L' \approx_{\Delta} K\{\tilde{c}'/\tilde{c}\} \mid \overline{l \cdot c}(\tilde{v})$. By (LTS RES) and freshness of \tilde{c}' , $(\nu \tilde{c}')(L\{\tilde{c}'/\tilde{c}\}) \xrightarrow{l \cdot c(\tilde{v})} \Delta (\nu \tilde{c}')L'$. By (LTS STRUCT), $(\nu \tilde{c}')L' \xrightarrow{l \cdot c(\tilde{v})} \Delta L'_1 \equiv (\nu \tilde{c}')L'$. By α -conversion and freshness of \tilde{c}' , $K'_1 \equiv (\nu \tilde{c}')(K\{\tilde{c}'/\tilde{c}\} \mid \overline{l \cdot c}(\tilde{v}))$, hence $K'_1 \approx_{\Delta} L'_1$. \square

Following Jeffrey and Rathke [17], we define in Figure 25 a *merge* operator $\langle\langle - \rangle\rangle$ to reconstruct processes from configurations. The merge operator of [17] though is partial, due to a potential circularity of references between trigger names and definitions. Since scripted processes in definitions cannot contain triggers, our merge operator is total. This operator plays a substantial rôle in showing that \approx_{Λ} is closed under parallel composition.

Before showing the properties of the merge operator, we illustrate three simple properties of the extraction function: it does not remove trigger names; it associates a definition to each trigger name it introduces; and we can recover the initial term by substituting the new definitions in the result term.

Lemma B.11 (Extraction Properties) *Suppose $\mathfrak{X}(\tilde{v}) = (\tilde{v}'; \Theta^{\tilde{k}})$, The following properties hold:*

1. *if $k \in \text{fn}(\tilde{v})$ then $k \in \text{fn}(\tilde{v}')$;*
2. *if $k \in \text{fn}(\tilde{v}') \setminus \{\tilde{k}\}$ then $k \in \text{fn}(\tilde{v})$;*
3. $\tilde{v} = \tilde{v}'^{\Theta^{\tilde{k}}}$.

Proof. By induction on the structure of \tilde{v} , using Definition 3.9. \square

Since definitions can appear only at the top level and scripts cannot contain free private channel names, we can always use structural equivalence to factor any configuration into the parallel composition of a process and a group of definitions. We will make substantial use of this property to show closure under parallel composition of \approx_{Λ} . Merging a configuration corresponds to substituting the script in each definition for the corresponding trigger names in the process

term of the configuration. Hence, the merge operator preserves the transitions that do not involve trigger names for which there is a corresponding definition. Moreover, if two configurations are bisimilar, then they must define the same trigger names.

Lemma B.12 (Merge Properties) *The merge operator satisfies the following properties:*

1. *factorization: for any well-formed K , there exist a process P and a configuration $\Theta^{\tilde{k}}$ such that $K \equiv P | \Theta^{\tilde{k}}$ and $\langle\langle K \rangle\rangle \equiv P^{\Theta^{\tilde{k}}}$;*
2. *transition preservation: for any $\Theta^{\tilde{k}}, \Theta^{\tilde{j}}$, if $P \xrightarrow{\alpha_l} P'$ and $\{\tilde{k}, \tilde{j}\} \cap n(\alpha_l) = \emptyset$ then $\langle\langle P | \Theta^{\tilde{k}} \rangle\rangle | \Theta^{\tilde{j}} \xrightarrow{\alpha_l} \langle\langle P' | \Theta^{\tilde{k}} \rangle\rangle | \Theta^{\tilde{j}}$;*
3. *If $K \approx_{\Lambda} L$ then $K \equiv P | \Theta^{\tilde{k}}$ and $L \equiv Q | \Theta^{\tilde{k}}$, and each pair of corresponding definitions contains scripts with the same patterns: that is, $\langle k \leftarrow (\tilde{\pi})P_k \rangle$ in $\Theta^{\tilde{k}}$ implies $\langle k \leftarrow (\tilde{\pi})Q_k \rangle$ in $\Theta^{\tilde{k}}$, and viceversa.*

Proof.

1. By induction on the structure of K .
2. Follows from syntactic reasoning using Lemma B.3. We show the case for request transitions, the other ones are similar. If $P \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} P'$, by point (4) of Lemma B.3, $P \equiv (\nu a)(Q | l \cdot \text{req}_{p'}(c))$ and $P' \equiv (\nu a)(Q | \overline{l \cdot c}(T) | \Theta^{\tilde{h}})$ for some p' such that $\mathfrak{X}_Q(p') = (p, \Theta^{\tilde{h}})$ and some appropriate T . By hypothesis, $\{\tilde{k}, \tilde{j}\} \cap n(\alpha_l) = \emptyset$, hence $\{\tilde{k}, \tilde{j}\} \cap \text{fn}(p) = \emptyset$. By point (1) of Lemma B.11, $\{\tilde{k}, \tilde{j}\} \cap \text{fn}(p') = \emptyset$. Since definitions do not contain free trigger names, $\{\tilde{k}, \tilde{j}\} \cap \text{fn}(\Theta^{\tilde{h}}) = \emptyset$. By definition of merge, $\langle\langle P | \Theta^{\tilde{k}} \rangle\rangle \equiv (\nu a)(\langle\langle Q | \Theta^{\tilde{k}} \rangle\rangle | l \cdot \text{req}_{p'}(c))$ and $\langle\langle P' | \Theta^{\tilde{k}} \rangle\rangle \equiv (\nu a)(\langle\langle Q | \Theta^{\tilde{k}} \rangle\rangle | \overline{l \cdot c}(T) | \Theta^{\tilde{h}})$. By definition of lts, $\langle\langle P | \Theta^{\tilde{k}} \rangle\rangle \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} \langle\langle P' | \Theta^{\tilde{k}} \rangle\rangle$ and $\langle\langle P | \Theta^{\tilde{k}} \rangle\rangle | \Theta^{\tilde{j}} \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} \langle\langle P' | \Theta^{\tilde{k}} \rangle\rangle | \Theta^{\tilde{j}}$.
3. By point (1) above, $K \equiv P | \Theta^{\tilde{k}}$ and $L \equiv Q | \Theta^{\tilde{j}}$. The argument is by contradiction. Suppose that $\Theta^{\tilde{k}} \equiv \Theta^{\tilde{h}} | \langle k \leftarrow (\tilde{\pi})P_k \rangle$ and $k \notin \tilde{j}$. By definition of lts, $K \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} K'$ but, since $k \notin \tilde{j}$, it is not possible to derive a corresponding (weak) transition for L , which contradicts $K \approx_{\Lambda} L$. The case for a different pattern π is analogous.

□

The lemma below analyzes the relationship between bisimilarity and definitions. We start noting that if we remove from the two configurations the

definitions for the same set of names, bisimilarity is preserved. Then, we note that the configurations obtained by duplicating existing definitions, using arbitrary fresh trigger names, remain bisimilar. These properties will be useful for showing that bisimilarity is closed under parallel composition.

Lemma B.13 (Bisimilarity and Definitions) *Let K and L be well-formed configurations.*

1. If $K \mid \Theta^{\tilde{k}} \approx_{\Delta} L \mid \Omega^{\tilde{k}}$ then $K \approx_{\Delta} L$.
2. If $K \mid \Theta^{\tilde{k}} \approx_{\Delta} L \mid \Omega^{\tilde{k}}$ then $K \mid \Theta^{\tilde{k}} \mid \Theta^{\tilde{k} \curvearrowright \tilde{j}} \approx_{\Delta} L \mid \Omega^{\tilde{k}} \mid \Omega^{\tilde{k} \curvearrowright \tilde{j}}$.

Proof.

1. The family \approx with generic element

$$\approx_{\Delta} = \left\{ (K, L) : K \mid \Theta^{\tilde{k}} \approx_{\Delta} L \mid \Omega^{\tilde{k}} \right\}$$

is a domain bisimulation. Follows by analyzing the transitions of K . The intuition is every transition by $K \mid \Theta^{\tilde{k}}$ originating from K must be matched by $L \mid \Omega^{\tilde{k}}$ using a (weak) transition originating from L alone, since $\Theta^{\tilde{k}}$ and $\Omega^{\tilde{k}}$ can perform only (LTS DEF) transitions. We show the most interesting case, for (LTS REQ). Suppose $K \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} K'$. By hypothesis, $K \mid \Theta^{\tilde{k}} \approx_{\Delta} L \mid \Omega^{\tilde{k}}$ for some \tilde{k} such that $\{\tilde{h}\} \cap \{\tilde{k}\} = \emptyset$. By definition of lts, $K \mid \Theta^{\tilde{k}} \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} K' \mid \Theta^{\tilde{k}}$. By bisimilarity, $L \mid \Omega^{\tilde{k}} \xrightarrow{\tau} L' \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} L'' \xrightarrow{\tau} L'''$ and $K' \mid \Theta^{\tilde{k}} \approx_{\Delta} L'''$. By definition of lts and by induction on the number of tau transitions, we have that $L \xrightarrow{\tau} L_1$ where $L' = L_1 \mid \Omega^{\tilde{k}}$. The transition $L_1 \mid \Omega^{\tilde{k}} \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} L''$ must follow by repeatedly applying (LTS PAR), starting from the premise $L_1 \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} L_2$, hence $L'' = L_2 \mid \Omega^{\tilde{k}}$. Again, by definition of lts and by induction on the number of tau transitions, we have that $L_2 \xrightarrow{\tau} L_3$ where $L''' = L_3 \mid \Omega^{\tilde{k}}$. By composing the transitions, we obtain $L \xrightarrow{(\tilde{h})l \cdot \text{req}(p)(T)} L_3$, and we conclude because $K' \mid \Theta^{\tilde{k}} \approx_{\Delta} L_3 \mid \Omega^{\tilde{k}}$.

2. The family \approx with generic element

$$\approx_{\Delta} = \left\{ (K \mid \Theta^{\tilde{k}} \mid \Theta^{\tilde{k} \curvearrowright \tilde{j}}, L \mid \Omega^{\tilde{k}} \mid \Omega^{\tilde{k} \curvearrowright \tilde{j}}) : K \mid \Theta^{\tilde{k}} \approx_{\Delta} L \mid \Omega^{\tilde{k}} \right\}$$

is a domain bisimulation. Follows by analysis of the transitions of $K \mid \Theta^{\tilde{k}} \mid \Theta^{\tilde{k} \curvearrowright \tilde{j}}$, by syntactic reasoning using Lemma B.7. The transitions by K and by L are treated like in point (1) above. The intuition for the (LTS DEF) transitions originated from $\Theta^{\tilde{k} \curvearrowright \tilde{j}}$ is that since $\Theta^{\tilde{k}}$ and $\Theta^{\tilde{k} \curvearrowright \tilde{j}}$ have

the same transitions up-to renaming of triggers, every process generated by trigger transitions from $\Theta^{\tilde{k} \curvearrowright \tilde{j}}$ could also be generated by $\Theta^{\tilde{k}}$ alone. All that is needed is to match the transition of $\Theta^{\tilde{k} \curvearrowright \tilde{j}}$ with a corresponding one by $\Omega^{\tilde{k} \curvearrowright \tilde{j}}$, which exists because $K \mid \Theta^{\tilde{k}} \approx_{\Lambda} L \mid \Omega^{\tilde{k}}$ and $\Omega^{\tilde{k}}$ can match $\Theta^{\tilde{k}}$.

□

Lemma B.14 (Parallel Composition) *Bisimilarity is closed under parallel composition: $K \approx_{\Lambda} L \implies K \mid M \approx_{\Lambda} L \mid M$.*

Proof. In order to show closure under parallel composition, we will identify a domain bisimulation \approx containing all the pairs of the form $(K \mid M, L \mid M)$ such that K is bisimilar to L , plus any other pair of terms generated by the labelled transition system. In particular, we must handle with care the terms generated by a communication steps between K (or L) and M involving scripts. The idea is that we represent explicitly, using the merge operators, the definitions corresponding to the communicated scripts. More in detail, by point 3 of Lemma B.12 we know that, since $K \approx_{\Lambda} L$, then $K \equiv \Theta^{\tilde{k}} \mid P'$ and $L \equiv \Omega^{\tilde{k}} \mid Q'$ for some appropriate $P', \Theta^{\tilde{k}}, Q', \Omega^{\tilde{k}}$. Moreover, using point 1 of Lemma B.12 we can rewrite $P' \equiv \langle\langle P \mid \Theta^{\tilde{m}} \rangle\rangle$ and $Q' \equiv \langle\langle Q \mid \Omega^{\tilde{m}} \rangle\rangle$ for some appropriate $P, \Theta^{\tilde{m}}, Q, \Omega^{\tilde{m}}$. That is, $K \equiv \Theta^{\tilde{k}} \mid \langle\langle P \mid \Theta^{\tilde{m}} \rangle\rangle$ and $L \equiv \Omega^{\tilde{k}} \mid \langle\langle Q \mid \Omega^{\tilde{m}} \rangle\rangle$. Again by point 1 of Lemma B.12, we also know that $M \equiv \langle\langle \Phi^{\tilde{n}} \mid R \rangle\rangle \mid \Phi^{\tilde{i}}$ for some $R, \Phi^{\tilde{i}}, \Phi^{\tilde{n}}$. The terms $\Theta^{\tilde{m}}, \Omega^{\tilde{m}}$ and $\Phi^{\tilde{n}}$ represent definitions corresponding to the scripts that respectively K, L or M may communicate in a future transition.

Our candidate bisimulation is the family \approx with generic element \approx_{Δ} defined (up to \equiv) by the pairs

$$((\nu \tilde{c})(\Theta^{\tilde{k}} \mid \langle\langle P \mid \Theta^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle \mid \langle\langle \Theta^{\tilde{j}} \mid \Phi^{\tilde{n}} \mid R \rangle\rangle \mid \Phi^{\tilde{i}}), (\nu \tilde{c})(\Omega^{\tilde{k}} \mid \langle\langle Q \mid \Omega^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle \mid \langle\langle \Omega^{\tilde{j}} \mid \Phi^{\tilde{n}} \mid R \rangle\rangle \mid \Phi^{\tilde{i}}))$$

(where all the $\tilde{h}, \tilde{i}, \tilde{j}, \tilde{k}, \tilde{m}, \tilde{n}$ are distinct) such that

$$\Theta^{\tilde{k}} \mid \Theta^{\tilde{m}} \mid \Theta^{\tilde{j}} \mid P \approx_{\Delta} \Omega^{\tilde{k}} \mid \Omega^{\tilde{m}} \mid \Omega^{\tilde{j}} \mid Q$$

The extra terms $\Phi^{\tilde{h}}, \Theta^{\tilde{j}}, \Omega^{\tilde{j}}$ represent the definitions corresponding to the scripts that respectively M, K or L may have communicated using a labelled transition to either K or L , or to M . Note that $\{(K \mid M, L \mid M) : K \approx_{\Delta} L\}$ is contained in \approx_{Δ} (up to \equiv), by choosing $\Phi^{\tilde{h}} = \Theta^{\tilde{j}} = \Phi^{\tilde{j}} = \mathbf{0}$ and \tilde{c} empty.

We now proceed to show that \approx is a domain bisimulation. For readability, we omit the subscript Δ on weak transitions $\xrightarrow{\alpha}_{\Delta}$, and we use the abbreviations

$$\begin{aligned} K_* &= (\nu \tilde{c})(\Theta^{\tilde{k}} \mid \langle\langle P \mid \Theta^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle \mid \langle\langle \Theta^{\tilde{j}} \mid \Phi^{\tilde{n}} \mid R \rangle\rangle \mid \Phi^{\tilde{i}}) \\ L_* &= (\nu \tilde{c})(\Omega^{\tilde{k}} \mid \langle\langle Q \mid \Omega^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle \mid \langle\langle \Omega^{\tilde{j}} \mid \Phi^{\tilde{n}} \mid R \rangle\rangle \mid \Phi^{\tilde{i}}) \\ K_1 &= \Theta^{\tilde{k}} \mid \Theta^{\tilde{m}} \mid \Theta^{\tilde{j}} \mid P \\ L_1 &= \Omega^{\tilde{k}} \mid \Omega^{\tilde{m}} \mid \Omega^{\tilde{j}} \mid Q \end{aligned}$$

Suppose $K_* \xrightarrow{\alpha_l} K'$. The proof is by cases on α_l , where we only consider the subcases with $l \in \Delta$ as the others follow from the definition of domain bisimulation.

For each case we use Lemma B.3, and pattern matching between the syntax of the terms above and of the terms in the lemma. We start with the case for input transitions, which is the easiest.

- $(K_* \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} K')$: We aim to bring the script instantiated by the transition inside the leftmost merge operator in K_* . In order to do so, we need to avoid both the capture of private channel names in σ by \tilde{c} , and clashes between trigger names in sigma and the vectors \tilde{h}, \tilde{m} . We split the proof in two cases, depending on whether $k \in \tilde{k}$ or $k \in \tilde{i}$.

- $\Theta^{\tilde{k}} \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} \Theta^{\tilde{k}} | \mathbf{P}_k \sigma$, where $\langle k \leftarrow (\tilde{\pi})\mathbf{P}_k \rangle$ is in $\Theta^{\tilde{k}}$. To avoid clashes between trigger names, we choose for some fresh \tilde{h}', \tilde{m}' and, using standard properties of substitution, rewrite

$$\langle\langle P | \Theta^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle = \langle\langle P \{ \tilde{m}' / \tilde{m} \} \{ \tilde{h}' / \tilde{h} \} | \Theta^{\tilde{m} \frown \tilde{m}'} | \Phi^{\tilde{h} \frown \tilde{h}'} \rangle\rangle$$

Let $\rho = \{ \tilde{m}' / \tilde{m} \} \{ \tilde{h}' / \tilde{h} \} \{ \tilde{c}' / \tilde{c} \}$ for a fresh \tilde{c}' , and recall that the private channel names \tilde{c} cannot appear free in definitions (by well-formedness of scripts). By α -conversion,

$$K' \equiv (\nu \tilde{c}') (\Theta^{\tilde{k}} | \langle\langle \mathbf{P}_k \sigma | P \rho | \Theta^{\tilde{m} \frown \tilde{m}'} | \Phi^{\tilde{h} \frown \tilde{h}'} \rangle\rangle | \langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R \{ \tilde{c}' / \tilde{c} \} \rangle\rangle | \Phi^{\tilde{i}})$$

By $K_1 \approx_\Delta L_1$ and Lemma B.8, $K_2 = K_1 \rho \approx_\Delta L_1 \rho = L_2$. Since $\Theta^{\tilde{k}}$ occurs in K_2 , $K_2 \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} K_2 | \mathbf{P}_k \sigma = K'_2$. By $K_2 \approx_\Delta L_2$, $L_2 \xrightarrow{\tau} \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} \xrightarrow{\tau} \Omega^{\tilde{k}} | \Omega^{\tilde{m} \frown \tilde{m}'} | \Omega^{\tilde{j}} | Q' = L'_2$ with $L'_2 \approx_\Delta K'_2$. We will now use the transition between L_2 and L'_2 to derive an appropriate one between L_* and L' . Since the action $\xrightarrow{l \cdot k(\tilde{\pi}\sigma)}$ necessarily originated by $\Omega^{\tilde{k}}$, which contains the definition $\langle k \leftarrow (\tilde{\pi})\mathbf{Q}_k \rangle$, we can reorder the reduction obtaining

$$L_2 \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} \Omega^{\tilde{k}} | \Omega^{\tilde{m} \frown \tilde{m}'} | \Omega^{\tilde{j}} | Q \rho | \mathbf{Q}_k \sigma \xrightarrow{\tau} \Omega^{\tilde{k}} | \Omega^{\tilde{m} \frown \tilde{m}'} | \Omega^{\tilde{j}} | Q'$$

By α -conversion, properties of substitutions and definition of lts,

$$L_* \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} L''$$

$$L'' = (\nu \tilde{c}') (\Omega^{\tilde{k}} | \langle\langle \mathbf{Q}_k \sigma | Q \rho | \Omega^{\tilde{m} \frown \tilde{m}'} | \Phi^{\tilde{h} \frown \tilde{h}'} \rangle\rangle | \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R \{ \tilde{c}' / \tilde{c} \} \rangle\rangle | \Phi^{\tilde{i}})$$

where we have brought \mathbf{Q}_k inside the leftmost merge operator, in order to preserve the general structure that we have imposed on

terms in \approx_{Δ} . By syntactical reasoning, it must be the case that $\mathbf{Q}_k \sigma | Q \rho \xrightarrow{\tau} Q'$. By point 2 of Lemma B.12,

$$L'' \xrightarrow{\tau} (\nu \tilde{c}')(\Omega^{\tilde{k}} | \langle\langle Q' | \Omega^{\tilde{m}, \tilde{h}'} | \Phi^{\tilde{h}, \tilde{h}'} \rangle\rangle | \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R\{\tilde{c}'/\tilde{c}\} \rangle\rangle | \Phi^{\tilde{i}}) = L'$$

and we conclude because, since $K'_2 \approx_{\Lambda} L'_2$, we have $(K', L') \in \approx_{\Delta}$.

- $\Phi^{\tilde{i}} \xrightarrow{l \cdot k(\tilde{\pi}\sigma)} \Phi^{\tilde{i}} | \mathbf{R}_k \sigma$: similar to the previous case but simpler, since, instead of using the hypothesis $K_1 \approx_{\Delta} L_1$, it is enough to use syntactical reasoning.
- $(K_* \xrightarrow{l \cdot a(\tilde{v})} K')$: Similar to the previous case.
- $(K_* \xrightarrow{(\tilde{b}_*, \tilde{k}_*) \overline{l \cdot a(\tilde{v})}} K')$: We distinguish two cases, depending on whether the output transition is originated by R or P .
 - Suppose $\langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle \xrightarrow{(\tilde{b}', \tilde{k}', \tilde{i}') \overline{l \cdot a(\tilde{v})}} K_0$, where $\tilde{c} = \tilde{c}', \tilde{b}$ and $\tilde{b}_* = \tilde{b}, \tilde{b}'$. We assume that the trigger names \tilde{k}' come from $\Theta^{\tilde{j}}$, whereas the \tilde{i}' come from $\Phi^{\tilde{n}}$ or R . Unfolding the definition of merge, by Lemma B.3 we have that $\langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle = R^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}} \equiv (\nu \tilde{b}')(\overline{l \cdot a(\tilde{v}_1^{\Theta^{\tilde{j}}})} | R_1^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}})$ with $\{\tilde{b}'\} \subseteq \text{fn}(\tilde{v}_1)$, where $\tilde{v}_1 = \tilde{v}_*^{\Phi^{\tilde{n}}}$ for some appropriate \tilde{v}_* . Moreover, we have $K_0 \equiv R_1^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}} | \Theta^{\tilde{k}'} | \Phi^{\tilde{i}'}$, where $\mathfrak{X}(\tilde{v}_1^{\Theta^{\tilde{j}}}) = (\tilde{v}; \Theta^{\tilde{k}'} | \Phi^{\tilde{i}'})$. To find out how to split the definitions produced by the extraction into $\Theta^{\tilde{k}'}$ and $\Phi^{\tilde{i}'}$, we assume to have first applied the extraction $\mathfrak{X}(\tilde{v}_1) = (\tilde{v}'; \Phi^{\tilde{i}'})$, which ensures that the definitions in $\Phi^{\tilde{i}'}$ come from R or $\Phi^{\tilde{n}}$. Then, applying $\mathfrak{X}(\tilde{v}_1^{\Theta^{\tilde{j}}}) = (\tilde{v}; \Theta^{\tilde{k}'} | \Phi^{\tilde{i}'})$ we can infer that for each $\langle k' \leftarrow P_0 \rangle$ in $\Theta^{\tilde{k}'}$ there is $\langle j \leftarrow P_0 \rangle$ in $\Theta^{\tilde{j}}$, since, by points 1 and 2 of Lemma B.11, we have that $\tilde{v}\{\tilde{k}'/\tilde{j}'\} = \tilde{v}'$, where \tilde{j}' are the triggers in \tilde{j} occurring also in v_1 . With this information, by definition of lts, we can rearrange K' to respect our general pattern

$$K' \equiv (\nu \tilde{c}')(\Theta^{\tilde{k}} | \Theta^{\tilde{k}'} | \langle\langle P | \Theta^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R_1 \rangle\rangle | \Phi^{\tilde{i}} | \Phi^{\tilde{i}'})$$

By applying the same argument to L_* , $\langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle = R^{\Omega^{\tilde{j}} \Phi^{\tilde{n}}}$ and

$$R^{\Omega^{\tilde{j}} \Phi^{\tilde{n}}} \equiv (\nu \tilde{b}')(\overline{l \cdot a(\tilde{v}_1^{\Omega^{\tilde{j}}})} | R_1^{\Omega^{\tilde{j}} \Phi^{\tilde{n}}}) \xrightarrow{(\tilde{b}', \tilde{k}', \tilde{i}') \overline{l \cdot a(\tilde{v})}} R_1^{\Omega^{\tilde{j}} \Phi^{\tilde{n}}} | \Omega^{\tilde{k}'} | \Phi^{\tilde{i}'}$$

for $\mathfrak{X}(\tilde{v}_1^{\Omega^{\tilde{j}}}) = (\tilde{v}; \Omega^{\tilde{k}'} | \Phi^{\tilde{i}'})$, where for each $\langle k' \leftarrow Q_0 \rangle$ in $\Omega^{\tilde{k}'}$ there is a $\langle j \leftarrow Q_0 \rangle$ in $\Omega^{\tilde{j}}$. By definition of lts,

$$L_* \xrightarrow{(\tilde{b}, \tilde{b}', \tilde{k}', \tilde{i}') \overline{l \cdot a(\tilde{v})}} L' \equiv (\nu \tilde{c}')(\Omega^{\tilde{k}} | \Omega^{\tilde{k}'} | \langle\langle Q | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R_1 \rangle\rangle | \Phi^{\tilde{i}} | \Phi^{\tilde{i}'})$$

By $K_1 \approx_\Delta L_1$ and point 2 of Lemma B.13,

$$\Theta^{\tilde{k}} | \Theta^{\tilde{k}'} | \Theta^{\tilde{m}} | \Theta^{\tilde{j}} | P \approx_\Delta \Omega^{\tilde{k}} | \Omega^{\tilde{k}'} | \Omega^{\tilde{m}} | \Omega^{\tilde{j}} | Q$$

and we conclude because $(K', L') \in \approx_\Delta$.

- Suppose $\langle\langle P | \Theta^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle \xrightarrow{(\tilde{b}', \tilde{k}', \tilde{m}', \tilde{i}') \overline{l \cdot a} \langle \tilde{v} \rangle} K_0$, where $\tilde{c} = \tilde{c}', \tilde{b}$ and $\tilde{b}_* = \tilde{b}, \tilde{b}'$. We assume \tilde{i}' are the new trigger names from $\Phi^{\tilde{h}}$, \tilde{m}' the ones from $\Theta^{\tilde{m}}$ and \tilde{k}' the ones from P . Differently from the previous case, we need to keep track explicitly of the triggers coming from $\Theta^{\tilde{m}}$, which will correspond in L_* to triggers coming from $\Omega^{\tilde{m}}$. We have that

$$\langle\langle P | \Theta^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle = P^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} \equiv (\nu \tilde{b}') (\overline{l \cdot a} \langle \tilde{v}_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} \rangle | P_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}})$$

with $\{\tilde{b}'\} \subseteq \text{fn}(\tilde{v})$, and $K_0 \equiv P_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} | \Theta^{\tilde{k}'} | \Theta^{\tilde{m}' } | \Phi^{\tilde{i}'}$, where we assume that \tilde{k}', \tilde{m}' and \tilde{i}' are disjoint from \tilde{k}, \tilde{m} and \tilde{i} . Moreover, we have

$$\begin{aligned} \mathfrak{X}(\tilde{v}_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}}) &= (\tilde{v}; \Theta^{\tilde{k}'} | \Theta^{\tilde{m}' } | \Phi^{\tilde{i}'}) \\ \mathfrak{X}(\tilde{v}_1^{\Theta^{\tilde{m}}}) &= (\tilde{v}'; \Theta^{\tilde{k}'} | \Theta^{\tilde{m}'}) \\ \mathfrak{X}(\tilde{v}_1) &= (\tilde{v}'; \Theta^{\tilde{k}'}) \end{aligned}$$

Hence, by Lemma B.11, for each $\langle i' \leftarrow R_0 \rangle$ in $\Phi^{\tilde{i}'}$ there is $\langle h \leftarrow R_0 \rangle$ in $\Phi^{\tilde{h}}$, and for each $\langle m' \leftarrow P_0 \rangle$ in $\Theta^{\tilde{m}'}$ there is $\langle m \leftarrow P_0 \rangle$ in $\Theta^{\tilde{m}}$. By definition of lts,

$$K' \equiv (\nu \tilde{c}') (\Theta^{\tilde{k}} | \Theta^{\tilde{k}'} | \Theta^{\tilde{m}' } | \langle\langle P_1 | \Theta^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle | \Phi^{\tilde{i}} | \Phi^{\tilde{i}'})$$

where we have used the information gathered above on $\Theta^{\tilde{k}'}$, $\Theta^{\tilde{m}'}$ and $\Phi^{\tilde{i}'}$ to decide how to rearrange K' , in order to fit our general pattern. By definition of lts,

$$K_1 \xrightarrow{(\tilde{b}', \tilde{k}') \overline{l \cdot a} \langle \tilde{v}' \rangle} P_1 | \Theta^{\tilde{k}} | \Theta^{\tilde{k}'} | \Theta^{\tilde{m}} | \Theta^{\tilde{j}} = K'_1$$

where $\{\tilde{b}'\} \subseteq \text{fn}(\tilde{v}')$ and $\mathfrak{X}(\tilde{v}_1) = (\tilde{v}'; \Theta^{\tilde{k}'})$. By $K_1 \approx_\Delta L_1$ and point 3 of Lemma B.12,

$$L_1 \xrightarrow{(\tilde{b}', \tilde{k}') \overline{l \cdot a} \langle \tilde{v}' \rangle} Q_1 | \Omega^{\tilde{k}} | \Omega^{\tilde{k}'} | \Omega^{\tilde{m}} | \Omega^{\tilde{j}} = L'_1$$

and $K'_1 \approx_\Delta L'_1$. We now derive a corresponding transition for L_* . Since none of the transitions above can be generated by a definition, we can deduce

$$Q \xrightarrow{\tau} (\nu \tilde{b}') (\overline{l \cdot a} \langle \tilde{v}_2 \rangle | Q_2) = Q_3, \quad Q_3 \xrightarrow{(\tilde{b}', \tilde{k}') \overline{l \cdot a} \langle \tilde{v}' \rangle} Q_2 | \Omega^{\tilde{k}'} \xrightarrow{\tau} Q_1 | \Omega^{\tilde{k}'}$$

where $\mathfrak{X}(\tilde{v}_2) = (\tilde{v}'; \Omega^{\tilde{k}'})$. By Lemma B.11, \tilde{v}_2 has exactly the same occurrences of trigger names in \tilde{h} as \tilde{v}' , which are the same of \tilde{v}_1 . By point 2 of Lemma B.12, $\langle\langle Q | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle \xrightarrow{\tau} \langle\langle Q_3 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle$. By syntactical reasoning

$$\langle\langle Q_3 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle \xrightarrow{(\tilde{b}', \tilde{k}', \tilde{m}', \tilde{i}') \overline{l \cdot a}(\tilde{v})} \langle\langle Q_2 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \Omega^{\tilde{k}'} | \Omega^{\tilde{m}'} | \Phi^{\tilde{i}'}$$

where $\mathfrak{X}(\tilde{v}_2^{\Omega^{\tilde{m}} \Phi^{\tilde{h}}}) = (\tilde{v}; \Omega^{\tilde{k}'} | \Omega^{\tilde{m}'} | \Phi^{\tilde{i}'})$. By point 2 of Lemma B.12 and by definition of lts,

$$\langle\langle Q_2 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \Omega^{\tilde{k}'} | \Omega^{\tilde{m}'} | \Phi^{\tilde{i}'} \xrightarrow{\tau} \langle\langle Q_1 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \Omega^{\tilde{k}'} | \Omega^{\tilde{m}'} | \Phi^{\tilde{i}'}$$

By definition of lts, $L_* \xrightarrow{(\tilde{b}, \tilde{b}', \tilde{k}', \tilde{m}', \tilde{i}') \overline{l \cdot a}(\tilde{v})} L'$ and

$$L' \equiv (\nu \tilde{c}')(\Omega^{\tilde{k}} | \Omega^{\tilde{k}'} | \Omega^{\tilde{m}'} | \langle\langle Q_1 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle | \Phi^{\tilde{i}} | \Phi^{\tilde{i}'})$$

By $K'_1 \approx_{\Delta} L'_1$ and point 2 of Lemma B.13,

$$\Theta^{\tilde{k}} | \Theta^{\tilde{k}'} | \Theta^{\tilde{m}'} | \Theta^{\tilde{m}} | \Theta^{\tilde{j}} | P_1 \approx_{\Delta} \Omega^{\tilde{k}} | \Omega^{\tilde{k}'} | \Omega^{\tilde{m}'} | \Omega^{\tilde{m}} | \Omega^{\tilde{j}} | Q_1$$

and we conclude because $(K', L') \in \approx_{\Delta}$.

- $(K_* \xrightarrow{(\tilde{b}_*, \tilde{k}_*) \overline{l \cdot j}(\tilde{v})} K')$: Analogous to the case for output.
- $(K_* \xrightarrow{(\tilde{k}) \overline{l \cdot \text{req}(p)}(T)} K')$: By combining the argument for input and output.
- $(K_* \xrightarrow{l \cdot \tau} K')$: First we analyze the case where the transition is determined by the interaction of R and P , then the case where the transition is derived by R or P in isolation.

Interaction. We analyze the transitions resulting from an interaction between R and P . We must distinguish four cases depending on whether R or P receives the value, and whether the value is received by a replicated input. Suppose $\langle\langle P | \Theta^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle \xrightarrow{\tau} K_0$.

Replicated input by R : By Lemma B.3, we have that

$$R^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}} \equiv (\nu \tilde{b})(R_1^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}} | !l \cdot a(\tilde{\pi}) \cdot R_2^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}})$$

where \tilde{b} is fresh with respect to P and \tilde{c} , and $a \notin \{\tilde{b}\}$. We also have that $P^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} \equiv (\nu \tilde{c}')(\overline{l \cdot a}(\tilde{v}^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}}) | P_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}})$, where \tilde{c}' is fresh with respect to R and \tilde{c} , $\{\tilde{c}'\} \subseteq \text{fn}(v^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}})$, and $\tilde{v}^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} = \tilde{\pi} \sigma'$. Moreover,

$$K_0 \equiv (\nu \tilde{c}')(P_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} | (\nu \tilde{b})(R_1^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}} | !l \cdot a(\tilde{\pi}) \cdot R_2^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}} | R_2^{\Theta^{\tilde{j}} \Phi^{\tilde{n}}} \sigma'))$$

Since scripts cannot contain free private names, $\{\tilde{c}'\} \subseteq \text{fn}(v^{\Theta^{\tilde{m}}\Phi^{\tilde{h}}})$ implies $\{\tilde{c}'\} \subseteq \text{fn}(v)$. Since patterns cannot contain scripts (or trigger names), there exists σ such that $\tilde{v} = \tilde{\pi}\sigma$ and $\sigma^{\Theta^{\tilde{m}}\Phi^{\tilde{h}}} = \sigma'$. We want to derive a configuration K' of the right form. By definition of merge, $P_1^{\Theta^{\tilde{m}}\Phi^{\tilde{h}}} = \langle\langle P_1 \mid \Theta^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle$. Before rewriting

$$R_* = (\nu \tilde{b})(R_1^{\Theta^{\tilde{j}}\Phi^{\tilde{n}}} \mid !l \cdot a(\tilde{\pi}).R_2^{\Theta^{\tilde{j}}\Phi^{\tilde{n}}} \mid R_2^{\Theta^{\tilde{j}}\Phi^{\tilde{n}}} \sigma^{\Theta^{\tilde{m}}\Phi^{\tilde{h}}})$$

in terms of the merge operator, we want to be explicit about the scripts occurring in \tilde{v} . Suppose $\mathfrak{X}(\tilde{v}) = (\tilde{v}'; \Theta^{\tilde{k}'})$ for fresh \tilde{k}' . Since $\tilde{v} = \tilde{\pi}\sigma$ and \tilde{v}' differs from \tilde{v} only for having triggers replacing scripts, there exists ρ such that $\tilde{v}' = \tilde{\pi}\rho$. By Lemma B.11, $\tilde{v} = \tilde{v}'^{\Theta^{\tilde{k}'}}$. Since $\tilde{v} = \tilde{\pi}\sigma$, $\tilde{v}' = \tilde{\pi}\rho$ and $\tilde{\pi}$ cannot contain trigger names, we have that $\sigma = \rho^{\Theta^{\tilde{k}'}}$, and both \tilde{v} and \tilde{v}' have the same occurrences of triggers in \tilde{h} and \tilde{m} . Without loss of generality, we assume that $\{\tilde{j}, \tilde{n}\} \cap \text{fn}(\rho) = \emptyset$. By standard properties of substitution, $R_2^{\Theta^{\tilde{j}}\Phi^{\tilde{n}}} \sigma^{\Theta^{\tilde{m}}\Phi^{\tilde{h}}} = (R_2\rho\{\tilde{m}'/\tilde{m}\})^{\Theta^{\tilde{m} \curvearrowright \tilde{m}'}\Phi^{\tilde{h}}\Theta^{\tilde{k}'}} \Theta^{\tilde{j}\Phi^{\tilde{n}}}$, where the vector \tilde{m}' is fresh. By definition of merge,

$$R_* = \langle\langle \Theta^{\tilde{j}} \mid \Theta^{\tilde{k}'} \mid \Theta^{\tilde{m} \curvearrowright \tilde{m}'} \mid \Phi^{\tilde{n}} \mid (\nu \tilde{b})(R_1 \mid !l \cdot a(\tilde{\pi}).R_2 \mid R_2\rho\{\tilde{m}'/\tilde{m}\}^{\Phi^{\tilde{h}}}) \rangle\rangle$$

By definition of lts,

$$K' \equiv (\nu \tilde{c}, \tilde{c}')(\Theta^{\tilde{k}} \mid \langle\langle P_1 \mid \Theta^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle \mid R_* \mid \Phi^{\tilde{i}})$$

By definition of lts,

$$K_1 \xrightarrow{(\tilde{c}', \tilde{k}')\overline{l \cdot a}\langle \tilde{v}' \rangle} P_1 \mid \Theta^{\tilde{j}} \mid \Theta^{\tilde{m}} \mid \Theta^{\tilde{k}} \mid \Theta^{\tilde{k}'} = K'_1$$

where, as noted above, $\mathfrak{X}(\tilde{v}) = (\tilde{v}'; \Theta^{\tilde{k}'})$. By $K_1 \approx_{\Lambda} L_1$,

$$L_1 \xrightarrow{\tau} (\nu \tilde{c}')(Q_1 \mid \overline{l \cdot a}\langle \tilde{v}_2 \rangle) \mid \Omega^{\tilde{j}} \mid \Omega^{\tilde{m}} \mid \Omega^{\tilde{k}} = L_2$$

$$L_2 \xrightarrow{(\tilde{c}', \tilde{k}')\overline{l \cdot a}\langle \tilde{v}' \rangle} Q_1 \mid \Omega^{\tilde{j}} \mid \Omega^{\tilde{m}} \mid \Omega^{\tilde{k}} \mid \Omega^{\tilde{k}'} = L_3$$

where $\mathfrak{X}(\tilde{v}_2) = (\tilde{v}'; \Omega^{\tilde{k}'})$, and

$$L_3 \xrightarrow{\tau} Q' \mid \Omega^{\tilde{j}} \mid \Omega^{\tilde{m}} \mid \Omega^{\tilde{k}} \mid \Omega^{\tilde{k}'} = L'_1$$

with $K'_1 \approx_{\Lambda} L'_1$. By point 2 of Lemma B.12, since $Q \xrightarrow{\tau} (\nu \tilde{c}')(Q_1 \mid \overline{l \cdot a}\langle \tilde{v}_2 \rangle)$,

$$\langle\langle Q \mid \Omega^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle \xrightarrow{\tau} \langle\langle (\nu \tilde{c}')(Q_1 \mid \overline{l \cdot a}\langle \tilde{v}_2 \rangle) \mid \Omega^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle\rangle = (\nu \tilde{c}')(Q_1^{\Phi^{\tilde{h}}} \mid \overline{l \cdot a}\langle \tilde{v}_2^{\Omega^{\tilde{m}}\Phi^{\tilde{h}}} \rangle)$$

By syntactical reasoning, $L_* \xrightarrow{\tau} L''$, where

$$L'' = (\nu \tilde{c}, \tilde{c}')(\Omega^{\tilde{k}} | \langle\langle Q_1 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | R'_* | \Phi^{\tilde{i}})$$

where, using an argument similar to that used for R_* ,

$$R'_* = \langle\langle \Omega^{\tilde{j}} | \Omega^{\tilde{k}'} | \Omega^{\tilde{m} \wedge \tilde{m}'} | \Phi^{\tilde{n}} | (\nu \tilde{b})(R_1 | !l \cdot a(\tilde{\pi}).R_2 | R_2 \rho \{\tilde{m}' / \tilde{m}\}^{\Phi^{\tilde{h}}}) \rangle\rangle$$

Note that R'_* is essentially R_* where each Θ is replaced by an Ω . By point 2 of Lemma B.12, since $Q_1 \xrightarrow{\tau} Q'$ we have that $L'' \xrightarrow{\tau} L'$, where

$$L' \equiv (\nu \tilde{c}, \tilde{c}')(\Omega^{\tilde{k}} | \langle\langle Q' | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | R'_* | \Phi^{\tilde{i}})$$

By $K'_1 \approx_{\Delta} L'_1$ and point 2 of Lemma B.13,

$$\Theta^{\tilde{k}} | \Theta^{\tilde{k}'} | \Theta^{\tilde{m} \wedge \tilde{m}'} | \Theta^{\tilde{m}} | \Theta^{\tilde{j}} | P_1 \approx_{\Delta} \Omega^{\tilde{k}} | \Omega^{\tilde{k}'} | \Omega^{\tilde{m} \wedge \tilde{m}'} | \Omega^{\tilde{m}} | \Omega^{\tilde{j}} | Q_1$$

and we conclude because $(K', L') \in \approx_{\Delta}$.

Input by R : analogous to the previous case.

Input by P : By Lemma B.3, we have $R^{\Phi^{\tilde{n}} \Theta^{\tilde{j}}} \equiv (\nu \tilde{c}') (R_1^{\Phi^{\tilde{n}} \Theta^{\tilde{j}}} | \overline{l \cdot a}(\tilde{v}^{\Theta^{\tilde{j}}}))$, where \tilde{c}' is fresh and $\{\tilde{c}'\} \subseteq \text{fn}(\tilde{v})$. Moreover, $P^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} \equiv (\nu \tilde{b})(P_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} | l \cdot a(\tilde{\pi}).P_2^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}})$, where \tilde{b} is fresh and $\tilde{v} = \tilde{\pi} \sigma$ (since scripts and trigger names cannot appear in patterns). Additionally, $K_0 \equiv (\nu \tilde{c}')((\nu \tilde{b})(P_1^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} | P_2^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} \sigma^{\Theta^{\tilde{j}}}) | R_1^{\Phi^{\tilde{n}} \Theta^{\tilde{j}}})$. In order to derive a K' of a suitable form, we follow a strategy similar to the one used in the case of replicated input by R . Let $\tilde{v}' = \tilde{v} \{ \tilde{j}' / \tilde{j} \}$ for some vector of fresh triggers \tilde{j}' . Since $\tilde{v} = \tilde{\pi} \sigma$ and $\tilde{\pi}$ cannot contain triggers, $\tilde{v}' = \tilde{\pi} \sigma \{ \tilde{j}' / \tilde{j} \}$. By standard properties of substitution, $\tilde{v}^{\Theta^{\tilde{j}}} = \tilde{v}'^{\Theta^{\tilde{j} \wedge \tilde{j}'}}$. Suppose $\mathfrak{X}(\tilde{v}') = (\tilde{v}_1; \Phi^{\tilde{h}'})$ for some fresh \tilde{h}' such that $\tilde{v}_1 = \tilde{\pi} \rho$. By Lemma B.11 and by freshness of \tilde{j}' , \tilde{h}' , we have $\tilde{v}' = \tilde{v}_1^{\Phi^{\tilde{h}'}}$, $\rho = \sigma \{ \tilde{j}' / \tilde{j} \}^{\Phi^{\tilde{h}'}}$ and $P_2^{\Theta^{\tilde{m}} \Phi^{\tilde{h}}} \sigma^{\Theta^{\tilde{j}}} = P_2 \rho^{\Theta^{\tilde{j} \wedge \tilde{j}'} \Phi^{\tilde{h}'}} \Theta^{\tilde{m}} \Phi^{\tilde{h}}$. By definition of lts,

$$K' \equiv (\nu \tilde{c}, \tilde{c}')(\Theta^{\tilde{k}} | \langle\langle (\nu \tilde{b})(P_1 | P_2 \rho) | \Theta^{\tilde{m}} | \Theta^{\tilde{j} \wedge \tilde{j}'} | \Phi^{\tilde{h}} | \Phi^{\tilde{h}'} \rangle\rangle | \langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R_1 \rangle\rangle | \Phi^{\tilde{i}})$$

By definition of lts,

$$K_1 \xrightarrow{l \cdot a(\tilde{v}_1)} \xrightarrow{\tau} (\nu \tilde{b})(P_1 | P_2 \rho) | \Omega^{\tilde{j}} | \Theta^{\tilde{m}} | \Theta^{\tilde{k}} = K'_1$$

By $K_1 \approx_{\Lambda} L_1$ and composing the weak actions of L_1 ,

$$L_1 \xrightarrow{l \cdot a(\tilde{v}_1)} \xrightarrow{\tau} Q' | \Omega^{\tilde{j}} | \Omega^{\tilde{m}} | \Omega^{\tilde{k}} = L'_1$$

and $K'_1 \approx_{\Lambda} L'_1$. By syntactical reasoning,

$$Q \xrightarrow{\tau} Q_1 \xrightarrow{l \cdot a(\tilde{v}_1)} Q_1 | \overline{l \cdot a}(\tilde{v}_1) \xrightarrow{\tau} Q'$$

By syntactical reasoning and by point 2 of Lemma B.12,

$$L_* \xrightarrow{\tau} (\nu \tilde{c})(\Omega^{\tilde{k}} | \langle\langle Q_1 | \Omega^{\tilde{m}} | \Phi^{\tilde{h}} \rangle\rangle | \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle | \Phi^{\tilde{i}}) = L''$$

By structural congruence, by $\tilde{v}' = \tilde{v}_1 \Phi^{\tilde{h}'}$ and $\tilde{v}^{\Omega^{\tilde{j}}\tilde{h}'}$, and since $R^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} \equiv (\nu \tilde{c}')(R_1^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} | \overline{l \cdot a}(\tilde{v}^{\Omega^{\tilde{j}}}))$,

$$L'' \equiv (\nu \tilde{c}, \tilde{c}')(\Omega^{\tilde{k}} | \langle\langle Q_1 | \overline{l \cdot a}(\tilde{v}_1) | \Omega^{\tilde{m}} | \Omega^{\tilde{j}\tilde{h}'} | \Phi^{\tilde{h}} | \Phi^{\tilde{h}'} \rangle\rangle | \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R_1 \rangle\rangle | \Phi^{\tilde{i}})$$

By point 2 of Lemma B.12, $L'' \xrightarrow{\tau} L'$ where

$$L' = (\nu \tilde{c}, \tilde{c}')(\Omega^{\tilde{k}} | \langle\langle Q' | \Omega^{\tilde{m}} | \Omega^{\tilde{j}\tilde{h}'} | \Phi^{\tilde{h}} | \Phi^{\tilde{h}'} \rangle\rangle | \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R_1 \rangle\rangle | \Phi^{\tilde{i}})$$

By point 2 of Lemma B.13, $K'_1 | \Theta^{\tilde{j}\tilde{h}'} \approx_{\Delta} L'_1 | \Omega^{\tilde{j}\tilde{h}'}$. and we conclude because $(K', L') \in \approx_{\Delta}$.

Replicated input by P : similar to the previous case.

Isolation. We show the case for R , as the case for P is similar. There are four ways to derive the sub-transition $\langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle \xrightarrow{l \cdot \tau} M_1$.

$$- R^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | l \cdot c(\tilde{\pi}) \cdot R_2^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | \overline{l \cdot c}(\tilde{\pi}\sigma^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}})) \text{ and}$$

$$M_1 \equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | (R_2^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}})\sigma^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}}) \equiv \langle\langle \Theta^{\tilde{j}} | \Phi^{\tilde{n}} | (\nu \tilde{a})(R_1 | R_2\sigma) \rangle\rangle$$

By syntactical reasoning, we can also derive

$$\langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | R \rangle\rangle \equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} | l \cdot c(\tilde{\pi}) \cdot R_2^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} | \overline{l \cdot c}(\tilde{\pi}\sigma^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}}))$$

$$\xrightarrow{l \cdot \tau} (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} | (R_2^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}})\sigma^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}}) = \langle\langle \Omega^{\tilde{j}} | \Phi^{\tilde{n}} | (\nu \tilde{a})(R_1 | R_2\sigma) \rangle\rangle$$

and we conclude because, by definition of lts, we can use this transition to derive a transition for L_* matching the one of K_* , with the resulting states K' and L' still in \approx_{Δ} .

$$- R^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | !l \cdot c(\tilde{\pi}) \cdot R_2^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | \overline{l \cdot c}(\tilde{\pi}\sigma^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}}))$$
: analogous to the previous case.

$$- R^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | m \cdot \text{go } l \cdot R_2^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}})$$
: similar to the previous cases, although the reasoning on $R_2^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}}$ is carried on at location l .

$$- R^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | (x, \tilde{\pi})R_2 \circ \langle l, \tilde{\pi}\sigma^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \rangle)$$
 and

$$M_1 \equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} | R_2\{l/x\}\sigma^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}})$$

where we have used the equation $R_2^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} = R_2$, which holds because scripts cannot contain trigger names. In order for $R^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}}$ to have the form given above, R itself must have one of the three forms given below.

1. If $R \equiv (\nu \tilde{a})(R_1 \mid (x, \tilde{\pi})R_2 \circ \langle l, \tilde{\pi}\sigma \rangle)$ then, by definition of script, $((x, \tilde{\pi})R_2 \circ \langle l, \tilde{\pi}\sigma \rangle)^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} = (x, \tilde{\pi})R_2 \circ \langle l, \tilde{\pi}\sigma^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \rangle$ and the reasoning is similar to the cases above.
2. If $R \equiv (\nu \tilde{a})(R_1 \mid k \circ \langle l, \tilde{\pi}\sigma \rangle)$ where $\Phi^{\tilde{n}} = \Phi^{\tilde{n}_1} \mid \langle k \Leftarrow (x, \tilde{\pi})R_2 \rangle \mid \Phi^{\tilde{n}_2}$ with $\tilde{n}_1, k, \tilde{n}_2 = \tilde{n}$, then we have

$$\begin{aligned} R^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} &\equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \mid (x, \tilde{\pi})R_2 \circ \langle l, \tilde{\pi}\sigma^{\Phi^{\tilde{n}}\Theta^{\tilde{j}}} \rangle) \\ R^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} &\equiv (\nu \tilde{a})(R_1^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} \mid (x, \tilde{\pi})R_2 \circ \langle l, \tilde{\pi}\sigma^{\Phi^{\tilde{n}}\Omega^{\tilde{j}}} \rangle) \end{aligned}$$

and the reasoning is once again analogous to case 1.

3. The last and most interesting case arises if $R \equiv (\nu \tilde{a})(R_1 \mid k \circ \langle l, \tilde{\pi}\sigma \rangle)$ and $\Theta^{\tilde{j}} = \Theta^{\tilde{j}_1} \mid \langle k \Leftarrow (x, \tilde{\pi})R_2 \rangle \mid \Theta^{\tilde{j}_2}$, where $\tilde{j}_1, k, \tilde{j}_2 = \tilde{j}$. In this case, a tau transition by K_* corresponds to a (LTS DEF) transition by K_1 . Without loss of generality, we assume that the names \tilde{a} are fresh. In order to derive an appropriate K' , we need to be explicit about the scripts in $\tilde{\pi}\sigma$. Hence, suppose that $\mathfrak{X}(\tilde{\pi}\sigma) = (\tilde{\pi}\rho; \Phi^{\tilde{n}'})$ where \tilde{n}' is fresh. By definition of lts, $K_* \xrightarrow{\tau} K'$ where

$$K' = (\nu \tilde{c})(\Theta^{\tilde{k}} \mid \langle \langle P \mid \Theta^{\tilde{m}} \mid \Phi^{\tilde{h}} \rangle \rangle \mid \langle (\nu \tilde{a})(R_1 \mid R_2\{l/x\}\rho) \mid \Phi^{\tilde{n}'} \mid \Phi^{\tilde{n}} \mid \Theta^{\tilde{j}} \rangle \rangle \mid \Phi^{\tilde{i}})$$

Since R_2 comes from the definition $\langle k \Leftarrow (x, \tilde{\pi})R_2 \rangle$ which is part of $\Theta^{\tilde{j}}$ (hence was originated by some previous transition of P), we need to move $R_2R_2\{l/x\}\rho$ inside the leftmost merge operator.

$$K' \equiv (\nu \tilde{c}, \tilde{a})(\Theta^{\tilde{k}} \mid \langle \langle P \mid R_2\{l/x\}\rho \mid \Theta^{\tilde{m}} \mid \Phi^{\tilde{h}} \mid \Phi^{\tilde{n}'} \rangle \rangle \mid \langle \langle R_1 \mid \Phi^{\tilde{n}} \mid \Theta^{\tilde{j}} \rangle \rangle \mid \Phi^{\tilde{i}})$$

With K' of this form, we will derive a transition from K_1 to a suitable K'_1 , and use the bisimilarity hypothesis to derive a matching transition for L_* . By definition of lts,

$$K_1 \xrightarrow{l \cdot k(\tilde{\pi}\rho)} \Theta^{\tilde{k}} \mid \Theta^{\tilde{m}} \mid \Theta^{\tilde{j}} \mid P \mid R_2\{l/x\}\rho = K'_1$$

By $K_1 \approx_{\Lambda} L_1$ and by point 3 of Lemma B.12 we know that $\Omega^{\tilde{j}} = \Omega^{\tilde{j}_1} \mid \langle k \Leftarrow (x, \tilde{\pi}')R_3 \rangle \mid \Omega^{\tilde{j}_2}$. Using our standard reasoning on the hypothesis $K_1 \approx_{\Lambda} L_1$, we can derive the transition

$$L_1 \xrightarrow{l \cdot k(\tilde{\pi}\rho)} \Omega^{\tilde{k}} \mid \Omega^{\tilde{m}} \mid \Omega^{\tilde{j}} \mid Q' = L'_1 \approx_{\Lambda} K'_1$$

Breaking down the transition into $\xrightarrow{\tau} \xrightarrow{l \cdot k(\tilde{\pi}\rho)} \xrightarrow{\tau}$ and using point 2 of Lemma B.12, we obtain

$$L_* \xrightarrow{\tau} L' = (\nu \tilde{c}, \tilde{a})(\Omega^{\tilde{k}} \mid \langle \langle Q' \mid \Omega^{\tilde{m}} \mid \Phi^{\tilde{h}} \mid \Phi^{\tilde{n}'} \rangle \rangle \mid \langle \langle R_1 \mid \Phi^{\tilde{n}} \mid \Omega^{\tilde{j}} \rangle \rangle \mid \Phi^{\tilde{i}})$$

and hence $(K', L') \in \approx_{\Delta}$.

□

We can finally prove that domain bisimilarity is a congruence on both configurations and, more importantly, open processes. The extension to open processes does not involve significant difficulties because, by definition, bisimilarity for open processes is closed under arbitrary substitutions.

Theorem B.15 (Congruence) *Domain bisimilarity is a congruence:*

1. for all configuration contexts $C \in \mathcal{K}_W$ (Figure 11), if $K \approx_\Lambda L$ then $C[K] \approx_\Lambda C[L]$;
2. for all extended process contexts $C \in \mathcal{K}_f$ (Definition 3.5), if $P \approx_\Lambda Q$ then $C[P] \approx_\Lambda C[Q]$.

Proof.

1. By Lemma B.10, $K \approx_\Lambda L \implies (\nu \tilde{c})K \approx_\Lambda (\nu \tilde{c})L$. By Lemma B.14, $K \approx_\Lambda L \implies K | M \approx_\Lambda L | M$. By Proposition B.1, $K | M \approx_\Lambda L | M \implies M | K \approx_\Lambda M | L$.
2. We need to show that
 - (a) $P \approx_\Lambda Q \implies (\nu \tilde{c})P \approx_\Lambda (\nu \tilde{c})Q$;
 - (b) $P \approx_\Lambda Q \implies P | R \approx_\Lambda Q | R$;
 - (c) $P \approx_\Lambda Q \implies l \cdot c(\tilde{\pi}).P \approx_\Lambda l \cdot c(\tilde{\pi}).Q$;
 - (d) $P \approx_\Lambda Q \implies !l \cdot c(\tilde{\pi}).P \approx_\Lambda !l \cdot c(\tilde{\pi}).Q$;
 - (e) $P \approx_\Lambda Q \implies l \cdot \text{go } m.P \approx_\Lambda l \cdot \text{go } m.Q$.

By definition $P \approx_\Lambda Q$ if and only if $P\sigma \approx_\Lambda Q\sigma$ for all closing substitutions σ .

- (a) Consider an arbitrary closing substitution σ for P, Q . Since we assume substitutions to be capture avoiding, $((\nu \tilde{c})P)\sigma \equiv (\nu \tilde{c}')$ $(P\{c'/c\}\sigma)$ and $((\nu \tilde{c})Q)\sigma \equiv (\nu \tilde{c}')$ $(Q\{c'/c\}\sigma)$. By hypothesis, $P\{c'/c\}\sigma \approx_\Lambda Q\{c'/c\}\sigma$. By point 1 above, $(\nu \tilde{c}')(P\{c'/c\}\sigma) \approx_\Lambda (\nu \tilde{c}')(Q\{c'/c\}\sigma)$.
- (b) Similar to point 2a, using point 1.
- (c) Let the family \approx have the generic element

$$\approx_\Delta = \approx_\Delta \cup \{(l \cdot c(\tilde{\pi}).P\sigma | M, l \cdot c(\tilde{\pi}).Q\sigma | M) : P\sigma \approx_\Delta Q\sigma\}$$

where σ is a closing substitution, $M = \prod_{n \geq 0} \overline{l_i \cdot c_i}(\tilde{v}_i)$ and $\text{scripts}(\tilde{v}_i) = \emptyset$.

The thesis follows by showing that \approx is a domain bisimulation, using point 1.

- (d) Similar to point 2c, using also transitivity of \approx_Δ (Lemma B.9).

(e) Let the family \approx have the generic element \approx_Δ given by

$$\begin{aligned} & \text{if } m \notin \Delta: \approx_\Delta = \approx_\Delta \cup \{(l \cdot \text{go } m \cdot P\sigma \mid M, l \cdot \text{go } m \cdot Q\sigma \mid M) : P\sigma \approx_\Delta Q\sigma\} \\ & \text{if } m \in \Delta: \approx_\Delta = \approx_\Delta \cup \left\{ \begin{array}{l} (l \cdot \text{go } m \cdot P\sigma \mid M, l \cdot \text{go } m \cdot Q\sigma \mid M), \\ (P\sigma \mid M, l \cdot \text{go } m \cdot Q\sigma \mid M), \\ (l \cdot \text{go } m \cdot P\sigma \mid M, Q\sigma \mid M) \end{array} : P\sigma \approx_\Delta Q\sigma \right\} \end{aligned}$$

where $\text{dom}(M) \subseteq \Delta$, σ is a closing substitution, $M = \prod_{n \geq 0} \overline{l_i \cdot c_i}(\tilde{v}_i)$ and $\text{scripts}(\tilde{v}_i) = \emptyset$. The family \approx is a domain bisimulation. □

B.3 Soundness

In this section, we show *soundness*: if two processes are domain bisimilar with respect to Λ , then they are request congruent with respect to Λ . Our strategy for proving the soundness of \approx_Λ consists of three main steps. First, we define an auxiliary relation \succ on Core Xd π networks such that two networks are in the relation if the corresponding processes, in parallel with the definitions extracted from the scripts in the corresponding stores, are Λ -bisimilar. Second we show that \succ is included in \sim^Λ , and third we use \succ as a stepping stone to relate \approx_Λ with \sim^Λ .

We begin comparing reductions and transitions. If a configuration K can perform a tau transition to become K' , then the process $\langle\langle K \rangle\rangle$ obtained by merging the configuration can reduce to $\langle\langle K' \rangle\rangle$, for any store compatible with K . On the other hand, if a process does a reduction step then, according to the lts, it can either perform a request transition or a tau transition, depending on whether (CRED REQUEST) was used in the derivation.

Lemma B.16 (Reductions) *Tau transitions between configurations imply reductions between the corresponding networks. For all D, K such that $\text{dom}(K) \subseteq \text{dom}(D)$*

1. if $K \xrightarrow{l\tau} K'$ then $(D, \langle\langle K \rangle\rangle) \rightarrow (D, \langle\langle K' \rangle\rangle)$;
2. if $K \xrightarrow{\tau} K'$ then $(D, \langle\langle K \rangle\rangle) \rightarrow^* (D, \langle\langle K' \rangle\rangle)$.

Proof. Point 1 follows from point 2 of Lemma B.12. Point 2 follows from point 1 noticing that tau transitions do not increase the domain of a configuration. □

Lemma B.17 (Transitions) *Reductions between networks imply tau or request transitions between the corresponding configurations. If $(D, P) \longrightarrow (D_1, P_1)$ then one of the following holds:*

1. $P \xrightarrow{l\tau} P_1$ and $D = \{l \mapsto T\} \uplus E = D_1$;

$$2. P \xrightarrow{\langle \tilde{k} \rangle l \cdot \text{req} \langle p' \rangle (T')} P_2 | \Theta^{\tilde{k}} \text{ and } \begin{cases} P_1 \equiv \langle\langle P_2 | \Theta^{\tilde{i}} \rangle\rangle, \\ D = \{l \mapsto T\} \uplus E, \\ D_1 = \{l \mapsto T_1\} \uplus E, \end{cases} \text{ where there exists a}$$

$$\text{path } p \text{ such that } \begin{cases} \mathfrak{E}(p, T) = (T_1, U_1 | \dots | U_n | \emptyset), \\ \mathfrak{X}(p) = (p'; \Theta^{\tilde{k}}), \\ \mathfrak{X}(\mathfrak{r}[U_1] | \dots | \mathfrak{r}[U_n] | \emptyset) = (T'; \Theta^{\tilde{i}}). \end{cases}$$

Proof. Both points follow by induction on the depth of the derivation tree of $(D, P) \rightarrow (D_1, P_1)$, using points 3 and 4 of Lemma B.3 to derive the labelled transition from the structure of the processes as revealed by the reduction step. \square

We know by Definition 3.8 that a script-independent query language, starting from queries and input trees which are equivalent up-to substitutions of scripts for trigger names, gives equivalent output trees and results. The lemma below relates this notion to extraction.

Lemma B.18 (Extraction) *Consider an arbitrary script-independent query language. Suppose $\mathfrak{X}(T) = (T_0; \Theta^{\tilde{k}})$, $\mathfrak{X}(p) = (p'; \Theta^{\tilde{j}})$, $\mathfrak{E}(p, T) = (T_1, U_1 | \dots | U_n | \emptyset)$, $\mathfrak{X}(\mathfrak{r}[U_1] | \dots | \mathfrak{r}[U_n] | \emptyset) = (T'; \Theta^{\tilde{i}})$ and $\mathfrak{X}(T_1) = (T'_1; \Theta^{\tilde{h}})$.*

1. for any definition $\langle k \Leftarrow A \rangle$ occurring in $\Theta^{\tilde{i}}$ or $\Theta^{\tilde{h}}$ there must be a definition $\langle k' \Leftarrow A \rangle$ occurring in $\Theta^{\tilde{k}}$ or $\Theta^{\tilde{j}}$.
2. if $\mathfrak{X}(S) = (T_0; \Omega^{\tilde{k}})$ and $\mathfrak{X}(q) = (p'; \Omega^{\tilde{j}})$, then $\mathfrak{E}(q, S) = (S_1; V_1 | \dots | V_n | \emptyset)$, $\mathfrak{X}(\mathfrak{r}[V_1] | \dots | \mathfrak{r}[V_n] | \emptyset) = (T'; \Omega^{\tilde{i}})$ and $\mathfrak{X}(S_1) = (T'_1; \Omega^{\tilde{h}})$.

Proof. Both points follow easily by Definition 3.8 and Observation 3.10. \square

We need to compare domain bisimilarity, which is defined on configurations without taking the store into account, with domain congruence, which is defined using reduction congruence (a relation on networks). We can do this because bisimilarity requires a correspondence between matching actions of two configurations, which implies that the stores in the networks corresponding to the configurations can diverge, after each reduction step, only up-to equivalent scripts. To formalize this intuition, we introduce the relation \asymp on networks.

Definition B.19 (Candidate Relation) *We define the candidate relation \asymp by*

$$\asymp \stackrel{\text{def}}{=} \left\{ ((D, P), (B, Q)) : \mathfrak{X}(D) = (D'; \Theta^{\tilde{k}}), \mathfrak{X}(B) = (D'; \Omega^{\tilde{k}}), P | \Theta^{\tilde{k}} \approx_{\Lambda} Q | \Omega^{\tilde{k}} \right\}$$

where $\text{dom}(B) = \text{dom}(D) = \Lambda$ and $(\text{fn}(P) \cup \text{fn}(Q)) \cap \mathcal{Y} = \emptyset$.

We can show now that the candidate relation \asymp is sound with respect to \simeq , the relation on networks inducing request congruence.

Lemma B.20 *The candidate relation is contained in the reduction congruence induced by request observables: $\asymp \subseteq \simeq$.*

Proof. By definition of \simeq , we need to show that \asymp is (1) observation preserving, (2) contextual, and (3) reduction-closed.

1. Follows from the definition of request observables, the hypothesis $P \mid \Theta^{\tilde{k}} \approx_{\Lambda} Q \mid \Omega^{\tilde{k}}$ and Lemma B.16, noticing that $\Theta^{\tilde{k}}$ and $\Omega^{\tilde{k}}$ cannot perform tau or request transitions.
2. Consider a generic reduction context $(E \uplus -, (\nu \tilde{c})(R \mid -))$. By definition of \mathfrak{X} , $\mathfrak{X}(E \uplus D) = (E' \uplus D'; \Phi^{\tilde{j}} \mid \Theta^{\tilde{k}})$ and $\mathfrak{X}(E \uplus B) = (E' \uplus D'; \Phi^{\tilde{j}} \mid \Omega^{\tilde{k}})$. By hypothesis, $P \mid \Theta^{\tilde{k}} \approx_{\Lambda} Q \mid \Omega^{\tilde{k}}$. By Theorem B.15, $(\nu \tilde{c})(R \mid P \mid \Theta^{\tilde{k}} \mid \Phi^{\tilde{j}}) \approx_{\Lambda} (\nu \tilde{c})(R \mid Q \mid \Omega^{\tilde{k}} \mid \Phi^{\tilde{j}})$. Since scripts have no private channel names, by structural congruence we conclude that $(\nu \tilde{c})(R \mid P) \mid \Theta^{\tilde{k}} \mid \Phi^{\tilde{j}} \approx_{\Lambda} (\nu \tilde{c})(R \mid Q) \mid \Omega^{\tilde{k}} \mid \Phi^{\tilde{j}}$.
3. Suppose $(D, P) \asymp (B, Q)$ and $(D, P) \longrightarrow (D_1, P_1)$. We need to show that $(B, Q) \xrightarrow{*} (B_1, Q_1) \asymp (D_1, P_1)$. For convenience, we report below what $(D, P) \asymp (B, Q)$ means:

$$\mathfrak{X}(D) = (D'; \Theta^{\tilde{k}}) \quad (1)$$

$$\mathfrak{X}(B) = (D'; \Omega^{\tilde{k}}) \quad (2)$$

$$P \mid \Theta^{\tilde{k}} \approx_{\Lambda} Q \mid \Omega^{\tilde{k}} \quad (3)$$

$$\text{dom}(B) = \text{dom}(D) = \Lambda \quad (4)$$

By Lemma B.17, there are two cases:

1. $P \xrightarrow{l\tau} P_1$ and $D = \{l \mapsto T\} \uplus E = D_1$.

By definition of lts and by equation (3) above,

$$\begin{array}{ccc} P \mid \Theta^{\tilde{k}} & \xrightarrow{l\tau} & P_1 \mid \Theta^{\tilde{k}} \\ \approx_{\Lambda} \downarrow & & \uparrow \approx_{\Lambda} \\ Q \mid \Omega^{\tilde{k}} \xrightarrow{\tau} \Lambda & \xrightarrow{l\tau} & \xrightarrow{\tau} \Lambda Q_1 \mid \Omega^{\tilde{k}} \end{array}$$

By syntactical reasoning, $Q \xrightarrow{l\tau} \Lambda Q_1$.

By Lemma B.16, $(B, Q) \xrightarrow{*} (B, Q_1)$. By (1), (2) and (4) we conclude with

$$\begin{array}{ccc} (D, P) & \longrightarrow & (D, P_1) \\ \asymp \downarrow & & \uparrow \asymp \\ (B, Q) & \xrightarrow{*} & (B, Q_1) \end{array}$$

2. $P \xrightarrow{(\tilde{j})^{l\text{-req}}(p')(T')} P_2 | \Theta^{\tilde{j}}$ where

$$\begin{aligned} P_1 &\equiv \langle\langle P_2 | \Theta^{\tilde{i}} \rangle\rangle, \\ D &= \{l \mapsto T\} \uplus E, \\ D_1 &= \{l \mapsto T_1\} \uplus E \end{aligned}$$

for some p such that

$$\begin{aligned} \mathfrak{E}(p, T) &= (T_1, U_1 | \dots | U_n | \emptyset), \\ \mathfrak{X}(p) &= (p'; \Theta^{\tilde{j}}), \\ \mathfrak{X}(\mathfrak{r}[U_1] | \dots | \mathfrak{r}[U_n] | \emptyset) &= (T'; \Theta^{\tilde{i}}). \end{aligned}$$

By definition of lts and by (3),

$$\begin{array}{ccc} P | \Theta^{\tilde{k}} & \xrightarrow{(\tilde{j})^{l\text{-req}}(p')(T')} & P_2 | \Theta^{\tilde{j}} | \Theta^{\tilde{k}} \\ \approx_\Lambda \downarrow & & \uparrow \approx_\Lambda \\ Q | \Omega^{\tilde{k}} \xrightarrow{\tau} \Lambda & \xrightarrow{(\tilde{j})^{l\text{-req}}(p')(T')} & \xrightarrow{\tau} \Lambda Q_2 | \Omega^{\tilde{j}} | \Omega^{\tilde{k}} \end{array} \quad (5)$$

By syntactical reasoning and point 4 of Lemma B.3,

$$Q \xrightarrow{\tau} \Lambda (\nu a)(Q_3 | l\text{-req}_q \langle c \rangle) = Q_M \quad (6)$$

$$Q_M \xrightarrow{(\tilde{j})^{l\text{-req}}(p')(T')} (\nu a)(Q_3 | \overline{l\text{-c}} \langle T' \rangle) | \Omega^{\tilde{j}} \xrightarrow{\tau} \Lambda Q_2 | \Omega^{\tilde{j}} \quad (7)$$

for some q such that $\mathfrak{X}(q) = (p'; \Omega^{\tilde{j}})$. By (6) and Lemma B.16,

$$(B, Q) \xrightarrow{*} (B, Q_M).$$

By (1) and definition of \mathfrak{X} ,

$$D' = \{l \mapsto T_0\} \uplus E_0 \text{ and } \Theta^{\tilde{k}} = \Theta^{\tilde{k}'} | \Theta^{\tilde{k}''}$$

where $\mathfrak{X}(\{l \mapsto T\}) = (\{l \mapsto T_0\}; \Theta^{\tilde{k}'})$ and $\mathfrak{X}(E) = (E_0; \Theta^{\tilde{k}''})$.

By (2) and a similar argument,

$$B = \{l \mapsto S\} \uplus E' \text{ and } \Omega^{\tilde{k}} = \Omega^{\tilde{k}'} | \Omega^{\tilde{k}''}$$

where $\mathfrak{X}(\{l \mapsto S\}) = (\{l \mapsto T_0\}; \Omega^{\tilde{k}'})$ and $\mathfrak{X}(E') = (E_0; \Omega^{\tilde{k}''})$.

Suppose $\mathfrak{X}(T_1) = (T'_1; \Theta^{\tilde{h}})$. By point 2 of Lemma B.18,

$$\begin{aligned} \mathfrak{E}(q, S) &= (S_1, V_1 | \dots | V_n | \emptyset), \\ \mathfrak{X}(\mathfrak{r}[V_1] | \dots | \mathfrak{r}[V_n] | \emptyset) &= (T'; \Omega^{\tilde{i}}), \\ \mathfrak{X}(S_1) &= (T'_1; \Omega^{\tilde{h}}). \end{aligned}$$

By definition of reduction,

$$(B, Q_M) \longrightarrow (B_1, (\nu a)(Q_3 \mid \overline{l \cdot c}(\mathbf{r}[V_1] \mid \dots \mid \mathbf{r}[V_n] \mid \emptyset)))$$

By point 2 of Lemma B.17,

$$(\nu a)(Q_3 \mid \overline{l \cdot c}(\mathbf{r}[V_1] \mid \dots \mid \mathbf{r}[V_n] \mid \emptyset)) \equiv \langle\langle (\nu a)(Q_3 \mid \overline{l \cdot c}(T')) \mid \Omega^{\tilde{i}} \rangle\rangle$$

By syntactical reasoning on (7),

$$(\nu a)(Q_3 \mid \overline{l \cdot c}(T')) \xrightarrow{\tau} {}_{\Lambda} Q_2.$$

By point 2 of Lemma B.12,

$$\langle\langle (\nu a)(Q_3 \mid \overline{l \cdot c}(T')) \mid \Omega^{\tilde{i}} \rangle\rangle \xrightarrow{\tau} {}_{\Lambda} \langle\langle Q_2 \mid \Omega^{\tilde{i}} \rangle\rangle.$$

By Lemma B.16,

$$(B_1, \langle\langle (\nu a)(Q_3 \mid \overline{l \cdot c}(T')) \mid \Omega^{\tilde{i}} \rangle\rangle) \xrightarrow{*} (B_1, \langle\langle Q_2 \mid \Omega^{\tilde{i}} \rangle\rangle)$$

By two applications of point 1 of Lemma B.18, for any definition $\langle i \Leftarrow A \rangle$ occurring in $\Theta^{\tilde{i}}$ or $\Theta^{\tilde{h}}$ there must be a definition $\langle k' \Leftarrow A \rangle$ occurring in $\Theta^{\tilde{k}}$ or $\Theta^{\tilde{j}}$.

By (5) and by Lemma B.13,

$$P_2 \mid \Theta^{\tilde{h}} \mid \Theta^{\tilde{k}'''} \mid \Theta^{\tilde{i}} \approx_{\Lambda} Q_2 \mid \Omega^{\tilde{h}} \mid \Omega^{\tilde{k}'''} \mid \Omega^{\tilde{i}}.$$

By using an appropriate instance of the candidate bisimulation in the proof of Lemma B.14,

$$\langle\langle P_2 \mid \Theta^{\tilde{i}} \rangle\rangle \mid \Theta^{\tilde{h}} \mid \Theta^{\tilde{k}'''} \approx_{\Lambda} \langle\langle Q_2 \mid \Omega^{\tilde{i}} \rangle\rangle \mid \Omega^{\tilde{h}} \mid \Omega^{\tilde{k}'''},$$

and we conclude with

$$\begin{array}{ccc} (D, P) & \longrightarrow & (D_1, \langle\langle P_2 \mid \Theta^{\tilde{i}} \rangle\rangle) \\ \simeq \downarrow & & \uparrow \simeq \\ (B, Q) & \xrightarrow{*} & (B_1, \langle\langle Q_2 \mid \Omega^{\tilde{i}} \rangle\rangle) \end{array}$$

□

We have now all the tools necessary to show the soundness of domain bisimilarity with respect to request congruence.

Theorem B.21 (Soundness) *Domain bisimilarity is a sound approximation of the domain congruence induced by request observables: for all $\Lambda, \mathbf{P}, \mathbf{Q}$ where $(fn(\mathbf{P}) \cup fn(\mathbf{Q})) \cap \mathcal{Y} = \emptyset$, if $\mathbf{P} \approx_{\Lambda} \mathbf{Q}$ then $\mathbf{P} \sim^{\Lambda} \mathbf{Q}$.*

Proof. By definition, $\mathbf{P} \sim^\Lambda \mathbf{Q}$ if and only if $(D, C[\mathbf{P}]) \simeq (D, C[\mathbf{Q}])$ for all $D, C[-]$ such that $\Lambda \subseteq \text{dom}(D)$ and $C[-]$ does not contain trigger names and is closing for both \mathbf{P} and \mathbf{Q} . Suppose $\mathbf{P} \approx_\Lambda \mathbf{Q}$ and consider some arbitrary $D, C[-]$ respecting the conditions above. Suppose $\mathfrak{X}(D) = (D'; \Theta^{\tilde{k}})$. By point 2 of Theorem B.15, $C[\mathbf{P}]|_{\Theta^{\tilde{k}}} \approx_\Lambda C[\mathbf{Q}]|_{\Theta^{\tilde{k}}}$. By Definition B.19, $(D, C[\mathbf{P}]) \asymp (D, C[\mathbf{Q}])$. By Lemma B.20, $(D, C[\mathbf{P}]) \simeq (D, C[\mathbf{Q}])$. \square