



A coordination language for databases

Li, Ximeng; Wu, Xi; Lluch Lafuente, Alberto; Nielson, Flemming; Nielson, Hanne Riis

Published in:
Logical Methods in Computer Science

Link to article, DOI:
[10.23638/LMCS-13\(1:10\)2017](https://doi.org/10.23638/LMCS-13(1:10)2017)

Publication date:
2017

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Li, X., Wu, X., Lluch Lafuente, A., Nielson, F., & Nielson, H. R. (2017). A coordination language for databases. *Logical Methods in Computer Science*, 13(1). [https://doi.org/10.23638/LMCS-13\(1:10\)2017](https://doi.org/10.23638/LMCS-13(1:10)2017)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A COORDINATION LANGUAGE FOR DATABASES

XIMENG LI^a, XI WU^b, ALBERTO LLUCH LAFUENTE^c, FLEMMING NIELSON^d,
AND HANNE RIIS NIELSON^e

^{a,c,d,e} Technical University of Denmark
e-mail address: {ximl, albl, fnie, hrni}@dtu.dk

^b The University of Queensland, East China Normal University
e-mail address: xi.wu@uq.edu.au

ABSTRACT. We present a coordination language for the modeling of distributed database applications. The language, baptized Klaim-DB, borrows the concepts of localities and nets of the coordination language Klaim but re-incarnates the tuple spaces of Klaim as databases. It provides high-level abstractions and primitives for the access and manipulation of structured data, with integrity and atomicity considerations. We present the formal semantics of Klaim-DB and develop a type system that avoids potential runtime errors such as certain evaluation errors and mismatches of data format in tables, which are monitored in the semantics. The use of the language is illustrated in a scenario where the sales from different branches of a chain of department stores are aggregated from their local databases. Raising the abstraction level and encapsulating integrity checks in the language primitives have benefited the modeling task considerably.

INTRODUCTION

Today’s data-intensive applications are becoming increasingly distributed. Multi-national collaborations on science, economics, military etc., require the communication and aggregation of data extracted from databases that are geographically dispersed. Distributed applications including websites frequently adopt the Model-View-Controller (MVC) design pattern in which the “model” layer is a database. Fault tolerance and recovery in databases also favors the employment of distribution and replication. The programmers of distributed database applications are faced with not only the challenge of writing good queries, but also that of dealing with the coordination of widely distributed databases. It is commonly accepted in the formal methods community that the *modeling* of complex systems in design can reduce implementation errors considerably [23, 1, 18].

Klaim [7] is a kernel language facilitating the specification of distributed and coordinated processes. In Klaim, processes and information repositories exist at different localities. The information repositories are in the form of tuple spaces, that can hold data and code. The processes can read tuples from (resp. write tuples to) local or remote tuple spaces, or spawn other processes to be executed at certain localities. Many distributed programming paradigms can be modeled in Klaim [7].

Key words and phrases: Coordination Language, Database, Distribution.

Klaim thus provides an ideal ground for the modeling of networked applications in general. However, the unstructured nature of tuple spaces and fine-grained operations mostly targeting individual tuples create difficulties in the description of the data-manipulation tasks that are usually performed using a high-level language such as SQL. A non-exhaustive list of the disadvantages of staying at the original abstraction level of Klaim are given below.

- A considerable amount of meta-data needed by databases has to be maintained as particular tuples or components of tuples.
- The sanity checks associated with database operations have to be borne in mind and performed manually by the programmer.
- Atomicity guarantees are difficult to deliver when batch operations are performed.

To support the modeling of applications operating on distributed, structured data, we propose the coordination language Klaim-DB. The language is inspired by Klaim in which it allows the distribution of *nodes*, and remote operations on data. Its succinct syntax facilitates the precise formulation of a structural operational semantics, giving rigor to high-level formal specification and reasoning of distributed database applications. The language also provides structured data organized as databases and tables, and high-level actions that accomplish the data-definition and data-manipulation tasks ubiquitous in these applications.

Since it is desirable to be able to predict and rule out certain classes of runtime errors statically, we develop a type system that scrutinizes specifications for potential evaluation errors and format mismatches in manipulating structured data. These errors are monitored in the semantics, and the safety of the type system (in the sense of subject reduction [19]) and its soundness with respect to the monitoring are proved. The efficiency of type checking is also evidenced by a formal result.

We will use database operations involved in the management of a large-scale chain of department stores as our running example. Each individual store in the chain has its local database containing information about the current stock and sales of each kind of product. The semantic rules for the core database operations will be illustrated by the local query and maintenance of these individual databases, and our final case study will be concerned with data aggregation across multiple local databases needed to generate statistics on the overall product sales.

Compared to the conference version [24], our new contribution is twofold.

- (1) The language design is made more succinct and uniform, and is augmented with two important features: code mobility and join operations on tables (which were only discussed as a potential extension previously).
- (2) The type system mentioned above and its theoretical properties were not developed in [24]. To clarify the guarantees of well-typedness, we have also adapted the semantics of [24] to signal run-time errors explicitly, rather than get stuck, in case certain classes of abnormalities arise.


This paper is structured as follows. In Section 1, we briefly review the common database operations and the Klaim language, and describe the design of Klaim-DB. In Section 2, the syntax of Klaim-DB is presented, which is followed by the structural operational semantics specified in Section 3. In Section 4, we present the type system and establish its theoretical properties. Our case study is then presented in Section 5. We discuss potential alternatives at several points in our development and related work in Section 6, and conclude in Section 7. Multiset notations are explained in Appendix A and the proofs of theoretical results can be found in Appendix B.

1. BACKGROUND

In this section we give a brief overview of common database operations and of the coordination language Klaim. On this basis, we introduce our ideas about the design of Klaim-DB.

Our discussion hereafter will be related to the following scenario. Consider the management of a large-scale chain of department stores, in which the head office can manage the sales of different imaginary brands (e.g., KLD, SH,...) in its individual branches, as shown in Figure 1. The localities l_0, \dots, l_n represent the distinct places where the database systems of the head office and the branches are located.

Figure 1: Running Example



Suppose a table **Stores** exists in the database of the head office (Figure 2). Each row is a record of the information of one branch: the city of the branch, its address and name, the brands sold in the branch and the locality corresponding to the branch.

Figure 2: The Table **Stores**

<i>City</i>	<i>Address</i>	<i>Shop_Name</i>	<i>Brands</i>	<i>Locality</i>
CPH	ABC DEF 2, 1050	Shop1	{KLD, SH, ...}	l_1
CPH	DEF HIJ 13, 2800	Shop2	{KLD, SH, ...}	l_2
CPH	HIJ KLM 26, 1750	Shop3	{KLD, SH, ...}	l_3
AAL	KLM NOP 3, 3570	Shop4	{LAM, IMK, ...}	l_4
AAL	NOP QUW 18, 4500	Shop5	{LAM, IMK, ...}	l_5
...

Similarly, each database of a branch has several tables identified by brand names. Each such table records the stock and sales of all types of shoes of its corresponding brand. The table **KLD** in one of the branches, about the shoe brand KLD, is shown in Figure 3, where “HB” stands for “High Boot” and “SB” for “Short Boot”.

A Brief Overview of Database Operations. We now sketch the typical database operations involved in the management of the chain of department stores.

Insertion. Suppose a new shoe type of the brand KLD is added to stock, or a new branch has been opened up. Then a new row is supposed to be inserted into either the table **KLD** or the table **Stores**.

Figure 3: The Table **KLD** in one branch

<i>Shoe_ID</i>	<i>Shoe_type</i>	<i>Year</i>	<i>Color</i>	<i>Size</i>	<i>In-stock</i>	<i>Sales</i>
001	HB	2015	red	38	5	2
001	HB	2015	red	37	8	5
001	HB	2015	red	36	3	1
001	HB	2015	black	38	3	2
001	HB	2015	black	37	5	2
002	SB	2015	green	38	2	0
002	SB	2015	brown	37	4	3
...

Deletion. When a shoe type has become outdated and removed from stock, a row corresponding to this shoe type needs to be deleted from the table **KLD**.

Selection. From time to time, a manager at the head office may want to know the stock and sales figures of certain types of KLD shoes. Selection operations are well-suited to the manager’s needs. A concrete example is “select the color, size, and sales of the high boots that are not red”.

Update. After each pair of KLD shoes of some type is sold, the row in table **KLD** corresponding to that shoe type should be updated, deducing one from the *In-stock* field. After a new brand is introduced into the inventory at certain branches, the rows corresponding to these branches need to be updated, adding the brand name to the *Brand* field.

Aggregation. Generating sales statistics is an important task, especially when the amount of data gets large. Tasks such as totaling the number of KLD shoes colored black sold can be performed using aggregation operations. In general, the aggregation operation plays a key role in developing business analytics.

Table Creation. After a new brand is introduced into the inventory at certain branches, we need not only to update the records for these branches in the table **Stores**, but also to create a new table for the brand at each of these branches.

Table Deletion. If the company producing a shoe brand has gone out of business, and all the shoes of that brand has been sold or disposed of at some branch, then it would be desirable to be able to remove the table corresponding to that brand at that branch.

The Essentials of Klaim. Before touching upon the design rationale of Klaim-DB, we discuss the essential elements of the coordination language Klaim, which is a fundamental inspiration for Klaim-DB.

Klaim [7] (Kernel Language for Agents Interaction and Mobility) is a core language derived from the Linda language [12]. Like Linda, Klaim features *generative communication* [11] between processes: processes share “tuple spaces” and can generate tuples in them for other processes to retrieve via pattern matching. This allows a higher level of abstraction and expressiveness compared with communication via shared variables or message passing.

Klaim supports programming with *explicit use of localities*. The action $\text{out}(t)@l$ generates the tuple t in the tuple space *located at* l . The action $\text{in}(T)@l$ looks for tuples t matching the template T that potentially contains formal fields, in the tuple space *located*

at l . If such tuples exist, the formal fields of T are bound to the actual counterparts in t for later use, and t is subsequently removed from the tuple space. The action $\text{read}(T)@l$ is a non-destructive version of $\text{in}(T)@l$ — it does not remove matched tuples.

As an alternative to addressing a remote locality l directly, a process P can be spawned at l using the action $\text{eval}(P)@l$ to perform the tasks locally — this is a form of *code mobility*.

The Design of Klaim-DB. In Klaim-DB, we conceive tuple spaces as being structured into tables. Each located tuple space in Klaim is now a located database consisting of a collection of tables. A table has the structure (I, R) , where I is an *interface* that contains the table identifier and describes the data format of the table, and R is a multiset containing the data records.

Correspondingly, the basic operations of Klaim-DB are the high-level database operations such as insertion, selection, aggregation, etc., targeting whole tables, rather than plain outputs and inputs of tuples. Each of these database operations is executed in an *atomic* step and makes sure when necessary that the data involved conforms to the format specified in the interface of the table targeted. An example on atomic operations is an update operation concerned with the *Brands* field of the table **Stores**, that adds the brand **PKL** to the set of brands sold in all shops. This operation is performed as one holistic step, rather than in a record-by-record fashion. Hence if it is performed simultaneously with another update that removes the brand **PKL** from all shops, we will either find **PKL** among the brands sold in all the shops, or none of them, after both operations are completed.

In the operations of Klaim-DB, we often use Klaim-style pattern matching, combined with predicates, to provide considerable expressive power. Suppose we need to select from the table **KLD** (of Figure 3) the color, size and sales of the types of high boots with ID “001”, which are not red. We can achieve this with the **select** action of Klaim-DB, using the template $(!id, !tp, !yr, !cr, !sz, !is, !ss)$ to match against the 7-field rows of the table, specifying the criterion with the predicate $id = \text{“001”} \wedge tp = \text{“HB”} \wedge cr \neq \text{“red”}$, and signaling that the only fields needed in the result are for “color”, “size” and “sales” with the tuple (cr, sz, ss) .

Joining tables is supported in Klaim-DB, and we can in fact select data from the join of tables located at different places. In addition, the different kinds of operations can be combined to form more powerful queries. An example is guiding selections with the statistics generated in aggregation operations — consider the selection of shoes whose sales figures are above average, where the average number is obtained via a preceding aggregation.

Klaim-DB inherits the code mobility primitive from Klaim; so instead of operating on remote tables directly, processes can be spawned onto remote localities, and collect and manipulate data from tables at different places. Code mobility can yield typical parallel processing patterns, for example combining the **eval** primitive with **foreach** loops to process all tuples in a table in parallel.

Klaim-DB versus Standard Klaim. One may be prompted to think about whether standard Klaim is already sufficiently expressive to model the databases and their corresponding management issues in the aforementioned scenario. The answer is positive, but without defeating the purpose of Klaim-DB:

- (1) By providing data abstractions like tables, and language primitives such as selection and aggregation, Klaim-DB makes a designer *think* in terms of the management of

distributed databases, rather than the manipulation of individual tuples. For example, data selection is naturally a batch operation in Klaim-DB, where the system designer can specify complicated criteria via predicates, rather than a replication/recursion in which basic input operations are performed, as in Klaim.

- (2) Related to (1), Klaim-DB enables us to maintain certain meta-data in a more natural and efficient way. In standard Klaim, the membership of the tuple $(001, HB, 2015, red, 38, 5, 2)$ in the table **KLD** can be represented by augmenting the tuple with the component **KLD**, e.g., as $(KLD, 001, HB, 2015, red, 38, 5, 2)$. Such a mixture of structural information and content information is not beneficial for addressing the meta-data by specified fields of language primitives — instead of adhering to one single convention imposed by a language, one has to follow potentially different conventions each specific to a particular system. In addition, meta-data such as the table identifier need to be repeated in all tuples, reducing space efficiency.
- (3) Klaim-DB actions are atomic, which avoids race conditions existing in situations like updating a table over which a selection is being performed, or concurrently updating the same table, and reduces unexpected behaviors under concurrency in general. The ability to have fine-grained interleaving of tuple-by-tuple operations can be recovered using **foreach** loops, *only when efficiency is important*.
- (4) Access control types are developed in standard Klaim and datatypes are used in X-Klaim. We need *dedicated* typing disciplines to ease database-oriented thinking and design. Supposing an action inserting the tuple $(001, HB, 2015)$ into the table **KLD** is placed in a piece of specification, it would be desirable to raise a typing error. In Klaim-DB tables, the separation of structural information in interfaces I and content information in data sets R facilitates maintaining the types of *all* the rows of each table in its interface, thereby facilitating the detection of typing errors including but are not limited to the aforementioned kind.

2. SYNTAX

This section focuses on the syntactic features of Klaim-DB. We start presenting the syntax of Klaim-DB in Section 2.1. We then present in Section 2.2 a first description of the types that we shall consider in our type system. We also provide some additional remarks on Klaim (Section 2.3) and examples based on our case study (Section 2.4).

2.1. Klaim-DB syntax. The syntax of Klaim-DB is presented in Figure 4. In the following we discuss all syntactic ingredients of the language following bottom-up order: from the bottom-most syntactic category e for expressions to the top-most syntactic category N for networks.

Expressions. We assume a set \mathcal{L} of localities, and a set \mathcal{U} of locality variables. An *expression* e can be an integer num , a string str , a locality $l \in \mathcal{L}$, a table identifier tid , a variable for data (x) or for locality ($u \in \mathcal{U}$), the concatenation $e_1 \cdot e_2$ of (string) expressions e_1 and e_2 , an arithmetic operation aop applied on sub-expressions e_1 and e_2 , or a multiset $\{e_1, \dots, e_n\}$ of expressions e_1, \dots, e_n which are themselves multiset-free. We will follow [7] to write ℓ for either a locality l or a locality variable u .

Figure 4: The Syntax of Klaim-DB

$N ::= \text{nil}_N \mid ERR \mid N_1 \parallel N_2 \mid (\nu l)N \mid l :: C$	(networks)
$C ::= P \mid (I, R) \mid C_1 \mid C_2$	(components)
$P ::= \text{nil}_P \mid a.P \mid A(\tilde{e}) \mid \text{foreach}(TB, T, \psi, \leq) : P \mid P_1; P_2$	(processes)
$a ::= \text{insert}(t, tid@l) \mid \text{delete}(tid@l, T, \psi) \mid \text{select}(\widetilde{TB}, T, \psi, t, !tbv) \mid$ $\text{update}(tid@l, T, \psi, t) \mid \text{aggr}(tid@l, T, \psi, f, T') \mid$ $\text{create}(tid@l, sk) \mid \text{drop}(tid@l) \mid$ $\text{eval}(P)@l$	(actions)
$TB ::= tid@l \mid tbv \mid (I, R)$	(tables)
$I ::= (tid, sk)$	(interfaces)
$T ::= W_1, \dots, W_n \quad (n > 0)$	(templates)
$W ::= !x \mid !u$	(fields)
$t ::= e_1, \dots, e_n \quad (n > 0)$	(tuples)
$\psi ::= \text{true} \mid e_1 \text{ cop } e_2 \mid e_1 \in e_2 \mid \neg \psi_1 \mid \psi_1 \wedge \psi_2$	(predicates)
$e ::= \text{num} \mid \text{str} \mid tid \mid l \mid x \mid u \mid e_1 \cdot e_2 \mid e_1 \text{ aop } e_2 \mid \{e_1, \dots, e_n\} \quad (n \geq 0)$	(expressions)

Predicates. A *predicate* ψ can be **true** for logical truth, a comparison $e_1 \text{ cop } e_2$ of two expressions e_1 and e_2 , a membership test of expression e_1 in e_2 , the negation of some predicate ψ_1 , or a conjunction of two predicates ψ_1 and ψ_2 . For a comparison $e_1 \text{ cop } e_2$, *cop* ranges over equality between integers, strings, or localities, numerical ordering of integers, and alphabetical ordering of strings. For the membership test $e_1 \in e_2$ to be carried out, e_2 needs to evaluate to a multiset M .

Tuples and Templates. We distinguish between *tuples* t and *templates* T . The components of tuples are arbitrary expressions. On the other hand, a template T contains only formal fields $!x$ and $!u$ where x is a variable that can be bound to data, and u to localities as introduced above. We assume templates T are linear, i.e., each bound variable can only appear once in T .

Interfaces. An *interface* I of a table has the structure (tid, sk) where tid is the table identifier and sk is a schema (or a type) describing the data format of the table. We will use the syntax $I.tid$ and $I.sk$ to refer to the corresponding components in I , i.e., $(tid', sk').tid = tid'$ and $(tid', sk').sk = sk'$.

Tables. *Tables* of different forms constitute the syntactical category TB . In more detail, a table can be a reference $tid@l$, with identifier tid and locality l , a table variable tbv , or a concrete table (I, R) with interface I and data set R . For a concrete table (I, R) , the data set R is a multiset of tuples whose components are constant values. Each tuple in R corresponds to one row of data in the table, and is supposed to have $I.sk$ as its type.

Actions. There are eight different kinds of *actions*, which are explained in detail below. To help the readers familiar with the SQL language, we will discuss when put in a SQL-like syntax, what the DB-oriented actions would be. These *SQL-like* queries used for the analogy differ from standard SQL queries in two ways: located table references are used instead of plain table references, and bound variables of templates are used (in predicates and tuples) instead of field names (or attribute names). The motivation for this deviation is that a detailed encoding of Klaim-DB actions into SQL would require, for example, to deal with the conversion from position-based indexing in Klaim-DB schemas to name-based indexing in SQL. Since such encoding would hamper readability and the aim of the SQL analogy is to provide intuition rather than formal results, we prefer to keep a lighter notation.

Insertion: The action $\text{insert}(t, \text{tid}@\ell)$ is used to insert a new row t into a table with identifier tid inside the database at ℓ . In SQL-like syntax, the insertion action can be written as `INSERT INTO $\text{tid}@\ell$ VALUES t .`

Example 2.1 (Adding New Shoes). The following action inserts into the table `KLD` at \mathbf{l}_1 an entry for KLD high boots of a new color, white, sized “37”, produced in 2015, with 6 in stock:

$$\text{insert}((\text{“001”}, \text{“HB”}, \text{“2015”}, \text{“white”}, \text{“37”}, 6, 0), \text{KLD}@\mathbf{l}_1).$$

Deletion: The action $\text{delete}(\text{tid}@\ell, T, \psi)$ deletes all rows matching the template T , resulting in bindings that satisfy the predicate ψ , from the table referenced by tid in the database located at ℓ . This action corresponds to the SQL-like query `DELETE FROM $\text{tid}@\ell$ AS T WHERE ψ .`¹

Example 2.2 (Deleting Existing Shoes). The following action deletes all entries for white high boots of the brand KLD, sized “37”, from the table `KLD` at \mathbf{l}_1 :

$$\text{delete}(\text{KLD}@\mathbf{l}_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), tp = \text{“HB”} \wedge cr = \text{“white”} \wedge sz = \text{“37”}).$$

Selection: The action $\text{select}(\widetilde{TB}, T, \psi, t, !tbv)$ picks from the Cartesian product of the data sets of tables identified by the list elements of \widetilde{TB} all rows that match the template T , and resulting in bindings that satisfy the predicate ψ . The instantiations of t with these bindings form a new table that is bound to tbv . Forming Cartesian products and specifying constraints with ψ capture certain join operations of the Relational Algebra [6] and SQL. The selection action corresponds to the SQL-like query `SELECT t INTO $!tbv$ FROM TB_1, \dots, TB_n AS T WHERE ψ .`

Example 2.3 (Selection of Shoes in a Certain Color). The following action selects the color, size, and sales of the types of high boots of the brand KLD that are not red, from the database at \mathbf{l}_1 , with the resulting table substituted into tbv :

$$\begin{aligned} &\text{select}(\text{KLD}@\mathbf{l}_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), \\ &\quad id = \text{“001”} \wedge tp = \text{“HB”} \wedge cr \neq \text{“red”}, (cr, sz, ss), !tbv). \end{aligned}$$

¹Readers familiar with SQL will notice our abuse of notation using SQL’s `AS` clause in combination with a template T . As we explained above, we prefer to keep a lighter notation and resort to the reader’s intuition. This remark applies to all subsequent references to SQL-like operations, where we perform similar abuses of notation.

Update: The action $\text{update}(tid@l, T, \psi, t)$ replaces each row matching T and satisfying ψ in table tid (at l) with a new row t , while leaving the rest of the rows unchanged. This action corresponds to the SQL-like query $\text{UPDATE } tid@l \text{ AS } T \text{ SET } t \text{ WHERE } \psi$.

Example 2.4 (Update of Shoes Information). Suppose two more red KLD high boots sized 37 are sold. The following action informs the database at l_1 of the corresponding changes in the stock and sales figures.

$$\text{update}(\text{KLD}@l_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), tp = \text{"HB"} \wedge cr = \text{"red"} \wedge sz = \text{"37"}, \\ (id, tp, yr, cr, sz, is - 2, ss + 2)).$$

Aggregation: The action $\text{aggr}(tid@l, T, \psi, f, T')$ applies the aggregator function f on the multiset of all rows matching T and satisfying ψ in table tid (at l) and binds the aggregation result to the pattern T' . In SQL, aggregation is achieved via selection, and our aggregation action can be compared to the SQL-like query $\text{SELECT } f(t) \text{ INTO } T' \text{ FROM } tid@l \text{ AS } T \text{ WHERE } \psi$.

Example 2.5 (Aggregation of Sales Figures). The following action produces the total sales (substituted into res) of shoes of the brand KLD with ID “001” at l_1 :

$$\text{aggr}(\text{KLD}@l_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), id = \text{"001"}, sum_7, !res),$$

where $sum_7 = \lambda R. (sum(\{v_7 | (v_1, \dots, v_7) \in R\}))$, i.e., sum_7 is a function from multisets R to unary tuples containing the summation results of the 7-th components of the tuples in R .

Table Creation and Removal: The action $\text{create}(tid@l, sk)$ creates a table with interface (tid, sk) at locality l . The SQL-like version is $\text{CREATE TABLE } tid@l (sk)$. The action $\text{drop}(tid@l)$ drops the table with identifier tid at locality l . In a SQL-like syntax this could be written $\text{DROP TABLE } tid@l$.

Example 2.6 (Creating Table for Turnovers). The following action creates a table with identifier **Turnover** at the head office, to record the yearly turnovers of the whole chain of department stores:

$$\text{create}(\text{Turnover}@l_0, \text{String} \times \text{Int}).$$

The table only has two columns — one for the year, represented by strings, and the other for the amount, represented by integers.

Code Mobility: As in Klaim, the action $\text{eval}(P)@l$ spawns the process P at the locality l . For our chain of department stores, it can be useful to create a mobile process from the system of the head office and to make it migrate among certain relevant branches to accomplish tasks including but are not limited to the collection of data. In general, the use of $\text{eval}(P)@l$ enhances the ability to parallelize and distribute the operations on databases, in particular when combined with the construct to perform iterations, as we shall see later on.

We will use $fv(-)$ and $bv(-)$, respectively, to refer to the set of free variables and the set of bound variables of their arguments, which can be expressions, predicates, tables, tuples, templates or processes. We will also use $fl(N)$ to refer to the set of free localities of net N . We globally assume that bound variables and bound localities all have distinct names.

Processes. A *process* P can be an inert process nil_P , an action-prefixed process $a.P$, a parameterized procedure invocation $A(\tilde{e})$, a looping process $\text{foreach}(TB, T, \psi, \leq) : P$, or a sequential composition $P_1; P_2$. Looping is introduced in addition to recursion via process invocation, to ease the task of traversing tables or data selection results in a natural way. In more detail, the process $\text{foreach}(TB, T, \psi, \leq) : P$ traverses the rows of TB that match the template T , resulting in bindings that satisfy the predicate ψ , and execute the process P once for each such binding produced. The parameter \leq is a partial order on tuples: for t_1 and t_2 as rows of TB such that $t_1 \leq t_2$ and $t_1 \neq t_2$, t_1 is processed before t_2 is. Sequential composition is allowed in addition to action-prefixing, to support imperative-programming-style thinking. We have chosen that prefixing binds more tightly than sequential composition, i.e., $a.P_1; P_2$ represents $(a.P_1); P_2$, and that in $P_1; P_2$, the scopes of the variables declared in P_1 are local to P_1 . Note that the explicit bracketing $a.(P_1; P_2)$ allows variables bound in a to be used in both processes P_1 and P_2 , unlike what is the case with $(a.P_1); P_2$. Each procedure A that is invoked with the list \tilde{e} of arguments needs to be defined with $A(\text{var}_1 : \tau_1, \dots, \text{var}_n : \tau_n) \triangleq P$, where $\text{var}_1 : \tau_1, \dots, \text{var}_n : \tau_n$ is a list of formal parameters x or u , associated with their types (the types admitted in our language will be explained shortly), and P is the process acting as the body of the procedure.

Components. A *component* C can be a process P , a table of the form (I, R) , or a parallel composition $C_1|C_2$ of two components.

Networks. A *net* N models a distributed database system that may contain several databases situated at different localities. An empty net is represented by nil_N . The construct $l :: C$ represents a node of the net, which captures the ensemble C of processes and tables at l . All the tables at a locality constitute the database at that locality. The special net ERR signals that an error has happened in the execution. The parallel composition of different nodes is represented using the \parallel operator. With the restriction operator $(\nu l)N$, the scope of l is restricted to N .

Systems. The specification of a Klaim-DB system takes the following form:

$$\begin{aligned} &\text{let } A_1(\text{var}_{11} : \tau_{11}, \dots, \text{var}_{1k_1} : \tau_{1k_1}) \triangleq P_1 \\ &\quad \dots \\ &\quad A_n(\text{var}_{n1} : \tau_{n1}, \dots, \text{var}_{nk_n} : \tau_{nk_n}) \triangleq P_n \\ &\text{in } N \end{aligned}$$

where each procedure used in N has its own definition. We focus on *closed* systems: i.e. systems where all variables to be used in any execution must have been previously bound, and there are no unknown components in the systems.

2.2. A taste of types. We introduce the types that our language admits — they have already appeared together with the formal parameters of procedures, and will play a fundamental role in our type system to be developed in Section 4. The following types are allowed:

$$\begin{aligned} \tau_d &::= \text{Int} \mid \text{String} \mid \text{Loc} \mid \text{Id} \\ \tau_m &::= \tau_d \text{ mset} \mid \tau_d \\ \tau_p &::= \tau_m^1 \times \dots \times \tau_m^n \\ \tau &::= \tau_p \mid \text{Bool} \mid \tau_p \text{ mset} \rightarrow \tau_p. \end{aligned}$$

A datatype τ_d can be *Int* for integers, *String* for strings, *Id* for table identifiers, and *Loc* for localities. A multiset type τ_m is either of the form τ_d *mset* for multisets containing elements of type τ_d , or just a data type. A product type τ_p is of the form $\tau_m^1 \times \dots \times \tau_m^n$, and is the type for tuples. Finally, a type τ can be a product type τ_p for structured data in tables, a boolean type *Bool* for predicates, or an arrow type τ_p *mset* $\rightarrow \tau_p$ for aggregator functions. Note that multisets are involved in two different ways. A component of a tuple can be a multiset, which is the case for the fourth component of all the tuples in Figure 2. In this situation, we restrict the types of the elements to be data types τ_d . On the other hand, each data set of a table is a multiset of tuples; hence the elements have product types τ_p , as is reflected in the signature of aggregator functions.

Throughout our presentation, a list of objects o_1, \dots, o_n will be denoted by \tilde{o} , or $\langle o(i) \rangle_{i \leq n}$ when the objects are structured. We also admit $\langle o(i) \rangle_i$ as a simpler version of the latter when the length is unimportant. For example, a list of formal parameters and types for a procedure definition can be $\langle var_i : \tau_i \rangle_i$. A list \widetilde{TB} of tables is a shorthand for $\langle TB_i \rangle_i$. For an object o , $|o|$ will denote its size — the number of elements of o when o is a list, and the number of components of o when o is a tuple or a template, etc.

2.3. On the syntax of Klaim and Klaim-DB. We comment here on two additional aspects in which Klaim-DB deviates from standard Klaim.

First, in Klaim, there is a two-layer stratification of localities — logical ones and physical ones. If a logical locality is thought of as the URL of a site, then its corresponding physical locality can be thought of as the IP address. Allocation environments are needed at each physical site to record which logical localities are locally mapped to which physical ones. Compared to the conference version ([24]) of this paper, we have replaced this two-layer structure with a single notion of “locality”, to be able to focus on the exposition of the semantics and typing of database management constructs and primitives. Recovering this missing mechanism would not create any technical difficulties for the development.

Second, in [7], tuples can contain both formal fields and actual ones. The distinction between templates that can contain formal fields and tuples that cannot, as made in [8], [10] and the current paper, has the effect of explicating certain constraints such as formal fields are not allowed in the data to be output into a tuple space or inserted into a table, when a tuple is used instead of a template. The further constraint used in the current development, that templates can only contain formal fields, is no real restriction — combining templates with *predicates* provides more flexibility and expressiveness than the matching of tuples/templates in [7], [8] and [10].

2.4. Modeling the Chain of Department Stores. Going back to the example with a chain of department stores introduced in Section 1, we now formally model the network of databases in Klaim-DB. Recall that the head office maintains a database at locality l_0 , and the branches maintain their own databases at localities l_1 to l_n .

The table **Stores** at the head office (Figure 2) can be represented as (I_0, R_0) , where $I_0.tid = \text{Stores}$. The header describes the purpose of each column of the table and suggests its appropriate type that shows up in $I_0.sk$, and the subsequent rows constitute the multiset R_0 , contain information of the different branches. We have $I_0.sk = \text{String} \times \text{String} \times \text{String} \times (\text{Id } mset) \times \text{Loc}$, where each component type describes one column of the table. Similarly, the table **KLD** at locality l_j ($1 \leq j \leq n$) can be represented as (I_j, R_j) , where $I_j.tid = \text{KLD}$

and $I_j.sk = \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{Int} \times \text{Int}$. We assume that the data stored in these tables indeed have the types specified in the schemas (quotation marks around strings are omitted in Figure 2 and Figure 3).

To sum up, the databases and operating processes used by the chain of department stores constitute the following net N_{DC} :

$$1_0 :: ((I_0, R_0)|C'_0) \parallel 1_1 :: ((I_1, R_1)|C'_1) \parallel \dots \parallel 1_n :: ((I_n, R_n)|C'_n),$$

where for $j \in \{1, \dots, n\}$, (I_j, R_j) describes the local table for the brand KLD inside its database at 1_j , and for each $k \in \{0, \dots, n\}$, C'_k stands for the remaining processes and tables at 1_j .

Remark 2.7. The reader may have found that there is no formal counterpart of the field names (City, Address, etc.) in the headers of the tables in Figure 2 and Figure 3. These field names were part of the schemas in the conference version [24], and getting rid of them in the current development is a deliberate choice — all the operations on schemas can be realized *positionally*, and the existence of field names in the schemas would be purely symbolic.

3. SEMANTICS

We devise a structural operational semantics [20] for Klaim-DB. The semantics is defined with the help of a structural congruence — the smallest congruence relation satisfying the rules in Figure 5, where the α -equivalence of N and N' is denoted by $N \equiv_\alpha N'$.

Figure 5: The Structural Congruence

$N_1 \parallel N_2 \equiv N_2 \parallel N_1$	$(\nu l_1)(\nu l_2)N \equiv (\nu l_2)(\nu l_1)N$
$(N_1 \parallel N_2) \parallel N_3 \equiv N_1 \parallel (N_2 \parallel N_3)$	$N_1 \parallel ((\nu l)N_2) \equiv (\nu l)(N_1 \parallel N_2)$ (if $l \notin fl(N_1)$)
$N \equiv N'$ (if $N \equiv_\alpha N'$)	$l :: (C_1 C_2) \equiv l :: C_1 \parallel l :: C_2$
$N \parallel \text{nil}_N \equiv N$	$l :: (P \text{nil}_P) \equiv l :: P$

We start by covering some preliminaries. The evaluation of expressions, predicates and tuples using the function $\mathcal{E}[\![\!-\!]\!]$, is inductively defined in Figure 6. The evaluation result of an expression can be an integer, a string, a table identifier, a locality, or *err*, the evaluation result of a predicate is a value in $\{tt, ff, err\}$, and the evaluation result of a tuple is an evaluated tuple or *err*, where *err* signals an evaluation error in all three cases. We globally assume that arithmetic operations are ideal; thus there are no overflow errors. As a simplification, we also assume that division by zero yields the integer zero.

Pattern matching of evaluated tuples *et* against templates T , denoted et/T , is an operation that results in a substitution or *err*. The detailed definition is given in Figure 7, which contains several rule schemas that need to be instantiated by replacing *val* with one of *num*, *str* and *tid*. The replacement of *val_j* in the same multiset needs to be performed consistently, e.g., $\{num_1, num_2\}$ is allowed, but not $\{num_1, str_2\}$. In case the format of *et* does not match against that of T , the value *err* is resulted. For example, we have

$$(5, 7)/(!x, !y) = [5/x, 7/y]$$

$$(5, 7)/(!x, !y, !z) = err$$

We also write $\psi\sigma$ for the result of applying the substitution σ to the formula ψ .

Figure 6: The Evaluation of Expressions, Predicates and Tuples

Expressions	
$\mathcal{E}[\text{num}]$	$= \text{num}$
$\mathcal{E}[\text{str}]$	$= \text{str}$
$\mathcal{E}[\text{tid}]$	$= \text{tid}$
$\mathcal{E}[l]$	$= l$
$\mathcal{E}[e_1 \cdot e_2]$	$= \begin{cases} \mathcal{E}[e_1] \cdot \mathcal{E}[e_2] & \text{if } \mathcal{E}[e_1] \text{ and } \mathcal{E}[e_2] \text{ are strings} \\ err & \text{otherwise} \end{cases}$
$\mathcal{E}[e_1 \text{ aop } e_2]$	$= \begin{cases} \mathcal{E}[e_1] \text{ aop } \mathcal{E}[e_2] & \text{if } \mathcal{E}[e_1] \text{ and } \mathcal{E}[e_2] \text{ are integers} \\ err & \text{otherwise} \end{cases}$
$\mathcal{E}[\{e_1, \dots, e_n\}]$	$= \begin{cases} \{\mathcal{E}[e_1], \dots, \mathcal{E}[e_n]\} & \text{if for all } j, \mathcal{E}[e_j] \text{ are values of the same type } \tau_d \\ err & \text{otherwise} \end{cases}$
Predicates	
$\mathcal{E}[\text{true}]$	$= tt$
$\mathcal{E}[e_1 \text{ cop } e_2]$	$= \begin{cases} \mathcal{E}[e_1] \text{ cop } \mathcal{E}[e_2] & \text{if } \mathcal{E}[e_1] \text{ and } \mathcal{E}[e_2] \text{ are values of the same type } \tau_d \\ err & \text{otherwise} \end{cases}$
$\mathcal{E}[e_1 \in e_2]$	$= \begin{cases} \mathcal{E}[e_1] \in \mathcal{E}[e_2] & \text{if } \mathcal{E}[e_2] \text{ is a multiset of values of the type of } \mathcal{E}[e_1] \\ err & \text{otherwise} \end{cases}$
$\mathcal{E}[\neg \psi_1]$	$= \begin{cases} ff & \text{if } \mathcal{E}[\psi_1] = tt \\ tt & \text{if } \mathcal{E}[\psi_1] = ff \\ err & \text{if } \mathcal{E}[\psi_1] = err \end{cases}$
$\mathcal{E}[\psi_1 \wedge \psi_2]$	$= \begin{cases} tt & \text{if } \mathcal{E}[\psi_1] = tt \text{ and } \mathcal{E}[\psi_2] = tt \\ err & \text{if } \mathcal{E}[\psi_1] = err \text{ or } \mathcal{E}[\psi_2] = err \\ ff & \text{otherwise} \end{cases}$
Tuples	
$\mathcal{E}[e_1, \dots, e_n]$	$= \begin{cases} \mathcal{E}[e_1], \dots, \mathcal{E}[e_n] & \text{if } \forall j \in \{1, \dots, n\} : \mathcal{E}[e_j] \neq err \\ err & \text{otherwise} \end{cases}$

Figure 7: The Substitution et/T

$val/!x = [val/x] \quad \{val_1, \dots, val_n\}/!x = [\{val_1, \dots, val_n\}/x]$
$\{l_1, \dots, l_n\}/!x = [\{l_1, \dots, l_n\}/x] \quad l/!u = [l/u]$
$l/!x = err \quad val/!u = err \quad \{val_1, \dots, val_n\}/!u = err$
$(e_1, \dots, e_n)/(W_1, \dots, W_m) = \begin{cases} \sigma_1 \dots \sigma_m & \text{if } m = n \wedge \forall j : e_j/W_j = \sigma_j \wedge \sigma_j \neq err \\ err & \text{otherwise} \end{cases}$

The well-sortedness of tuples and templates is specified in Figure 8. The judgments $t \Vdash \tau_p$ and $T \Vdash \tau_p$ represent that the tuple t and the template T , respectively, are well-sorted under the sort τ_p , which is essentially a product type. In the case of templates, the

different kinds of binders $!x$ and $!u$ are treated separately — an $!x$ cannot have the sort Loc , which is the sort of any $!u$.

We will use \uplus , \cap and \setminus to represent the union, intersection and subtraction of multisets, and the detailed definitions are given in Appendix A.

Figure 8: Well-Sortedness of Evaluated Tuples and Templates

$num \Vdash Int$	$str \Vdash String$	$tid \Vdash Id$	$\frac{e_1 \Vdash \tau_d \dots e_n \Vdash \tau_d}{\{e_1, \dots, e_n\} \Vdash \tau_d \text{ mset}}$	$l \Vdash Loc$	$\frac{e_1 \Vdash \tau_1 \dots e_n \Vdash \tau_n}{e_1, \dots, e_n \Vdash \tau_1 \times \dots \times \tau_n}$
$!x \Vdash Int$	$!x \Vdash String$	$!x \Vdash Id$	$!x \Vdash \tau_d \text{ mset}$	$!u \Vdash Loc$	$\frac{W_1 \Vdash \tau_1 \dots W_n \Vdash \tau_n}{W_1, \dots, W_n \Vdash \tau_1 \times \dots \times \tau_n}$

We will use $I.sk \downarrow_t^T$ to represent the projection of the schema $I.sk$ according to the template T (matching the format requirements imposed by $I.sk$) and the tuple t . The projection result is a new schema that describes only the columns referred to by the components of t , obtainable according to Definition 3.1.

Definition 3.1 (Projection of Schema). Let W_1, \dots, W_k be a template and e_1, \dots, e_n be made of a constants or variables bound in some W_j , and such that $n \leq k$. The projection $(\tau_1 \times \dots \times \tau_k) \downarrow_{e_1, \dots, e_n}^{W_1, \dots, W_k}$ is defined as

$$(\tau_1 \times \dots \times \tau_k) \downarrow_{e_1, \dots, e_n}^{W_1, \dots, W_k} = \tau'_1 \times \dots \times \tau'_n \quad \text{where } \tau'_i = \begin{cases} \tau & \text{where } e_i \Vdash \tau \\ \tau_j & \text{where } W_j = !e_i \end{cases}$$

In all other cases the operation is undefined.

The following example demonstrates this projection operation on schemas.

Example 3.2. The table **KLD** of Figure 3 has the schema $\tau_{\text{KLD}} = String \times String \times String \times String \times String \times Int \times Int$. Suppose in addition that

$$T = (!id, !tp, !yr, !cr, !sz, !is, !ss)$$

and $t = (cr, sz, ss)$. Then $\tau_{\text{KLD}} \downarrow_t^T$ gives the types for tuples containing data from three of the columns, for the color, size and sales figures of KLD shoes. According to Definition 3.1, we have $\tau_{\text{KLD}} \downarrow_t^T = String \times String \times Int$. \square

To join tables, we define the operations $\otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$ and $\otimes_{\text{R}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$ that give the product of the schemas and data sets, respectively, of the list \widetilde{TB} of tables given in different forms. Elements of \widetilde{TB} that are table identifiers are regarded as references to tables that need to be provided in the list $\langle l_i :: (I_i, R_i) \rangle_i$.

In Definition 3.3, the operation $flatten_{\text{s}}(\cdot)$ eliminates one-level nesting in product types:

$$flatten_{\text{s}}((\tau_{\text{m}}^{11} \times \dots \times \tau_{\text{m}}^{1n_1}) \times \dots \times (\tau_{\text{m}}^{k1} \times \dots \times \tau_{\text{m}}^{kn_k})) = \tau_{\text{m}}^{11} \times \dots \times \tau_{\text{m}}^{1n_1} \times \dots \times \tau_{\text{m}}^{k1} \times \dots \times \tau_{\text{m}}^{kn_k}$$

In Definition 3.4, the operation $flatten_{\text{d}}(\cdot)$ takes a multiset of tuples as argument and eliminates one-level nesting in all the tuples in this argument set.

$$flatten_{\text{d}}(R) = \{(v_{11}, \dots, v_{1n_1}, \dots, v_{k1}, \dots, v_{kn_k}) \mid ((v_{11}, \dots, v_{1n_1}), \dots, (v_{k1}, \dots, v_{kn_k})) \in R\}$$

Definition 3.3 ($\otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$). Let \widetilde{TB} be a list of tables and $\langle l_i :: (I_i, R_i) \rangle_i$ be a list of localized tables.

$$\otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) = \begin{cases} \text{undef} & \text{if } (\exists tid, l, j : TB_j = tid@l \wedge \forall i : l_i \neq l \vee I_i.tid \neq tid) \vee \\ & \widetilde{TB} \text{ contains some } tbv \text{ or } tid@u \\ \text{flatten}_s(\tau_1 \times \dots \times \tau_n) \text{ where} & \text{otherwise} \\ n = |\widetilde{TB}| \text{ and} & \\ \tau_j = \begin{cases} I_k.sk & \text{if } \exists k : TB_j = I_k.tid@l_k \\ I_0.sk & \text{if } \exists R_0 : TB_j = (I_0, R_0) \end{cases} & \end{cases}$$

Definition 3.4 ($\otimes_R(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$). Let \widetilde{TB} be a list of tables and $\langle l_i :: (I_i, R_i) \rangle_i$ be a list of localized tables.

$$\otimes_R(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) = \begin{cases} \text{undef} & \text{if } (\exists tid, l, j : TB_j = tid@l \wedge \forall i : l_i \neq l \vee I_i.tid \neq tid) \vee \\ & \widetilde{TB} \text{ contains some } tbv \text{ or } tid@u \\ \text{flatten}_d(R_1 \times \dots \times R_n) \text{ where} & \text{otherwise} \\ n = |\widetilde{TB}| \text{ and} & \\ R_j = \begin{cases} R_k & \text{if } \exists k : TB_j = I_k.id@l_k \\ R_0 & \text{if } \exists I_0 : TB_j = (I_0, R_0) \end{cases} & \end{cases}$$

Example 3.5. Consider joining the table **KLD** at the locality l_1 and the table **KLD** at l_2 . Suppose the first table is given by $\text{KLD}@l_1$ and the second by (I_2, R_2) . We have $\otimes_{\text{sk}}(\langle \text{KLD}@l_1, (I_2, R_2) \rangle, \langle l_1 :: (I_1, R_1) \rangle) = \text{flatten}_s(I_1.sk \times I_2.sk) = \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{Int} \times \text{Int} \times \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{Int} \times \text{Int}$. Here $I_1.sk$ is obtained for the first table $\text{KLD}@l_1$ from the singleton list $\langle l_1 :: (I_1, R_1) \rangle$. We also have $\otimes_R(\langle \text{KLD}@l_1, (I_2, R_2) \rangle, \langle l_1 :: (I_1, R_1) \rangle) = \text{flatten}_d(R_1 \times R_2)$. We dispense with expanding the latter for its verbosity. \square

The semantics is designed to monitor evaluation errors and format mismatches such as between a tuple and a schema, that could arise at runtime. The side conditions of the transition rules are frequently structured as a case analysis — if an evaluation error or a format mismatch occurs, then the net *ERR* is resulted; otherwise a normal transition step will take place. In Section 4, we will develop a type system for which well-typedness implies the absence of situations where *ERR* is produced.

The semantic judgment is of the form $N_{\text{env}} \vdash N \rightarrow N'$, representing that the net N can make a transition into the net N' , with the net N_{env} as the environment. We allow the abbreviation of $\text{nil}_N \vdash N \rightarrow N'$ as $\vdash N \rightarrow N'$. The environment N_{env} is consulted when creating a table at some locality l , to check whether a table with the same identifier already exists at l .

We proceed with a detailed explanation of the semantic rules for Klaim-DB actions in Figure 9, Figure 10 and Figure 11, and the semantic rules for processes and nets in Figure 12. In the explanation, we will avoid reiterating that each table resides in a database located at some l , but directly state “table ... located at l ”. Note that an action that refers to a locality variable u , e.g., $\text{insert}(t, tid@u)$, cannot be executed until u is bound to some locality l . Hence in our semantic rules the actions (which are ready for execution) always come with concrete localities l , rather than u_0 , or ℓ .

Examples in the setting of the chain of department stores will be provided to illustrate these rules at work. The empty environment will be used for all the illustrations except that of **create**, since it is only with the **create** action that the environment is consulted for information.

Insertion and Deletion. Recall that the action $\text{insert}(t, \text{tid}@l)$ inserts a new row t into a table with identifier tid inside the database at l . In closed nets, t cannot contain variables when the insertion action is about to be executed, just as l cannot contain locality variables. The rule (INS) of Figure 9 describes the performance of this action from locality l_1 , with the parameter l already replaced with some locality l_2 . It is checked that the table identifier tid needs to agree with that of a destination table (I, R) already existing at l_2 . If such agreement exists, but the evaluation result of tuple t does not satisfy the format requirements imposed by the schema $I.sk$, then the net ERR is produced; otherwise the evaluated tuple is added into the data set R .

Figure 9: The Semantics for Insertion and Deletion

$$\begin{array}{l}
 \text{(INS)} \quad \frac{I.\text{tid} = \text{tid}}{N_{\text{env}} \vdash l_1 :: \text{insert}(t, \text{tid}@l_2).P \parallel l_2 :: (I, R) \rightarrow N'} \\
 \text{where } N' = \begin{cases} ERR & \text{if } \mathcal{E}[\![t]\!] \not\models I.sk \\ l_1 :: P \parallel l_2 :: (I, R \uplus \{\mathcal{E}[\![t]\!]\}) & \text{otherwise} \end{cases} \\
 \\
 \text{(DEL)} \quad \frac{I.\text{tid} = \text{tid}}{N_{\text{env}} \vdash l_1 :: \text{delete}(\text{tid}@l_2, T, \psi).P \parallel l_2 :: (I, R) \rightarrow N'} \\
 \text{where } N' = \begin{cases} ERR & \text{if } T \not\models I.sk \vee \exists t \in R : (t/T = \text{err} \vee \mathcal{E}[\![\psi(t/T)]\!] = \text{err}) \\ l_1 :: P \parallel l_2 :: (I, R') & \text{otherwise} \\ & \text{where } R' = \{t \in R \mid \mathcal{E}[\![\psi(t/T)]\!] \neq tt\} \end{cases}
 \end{array}$$

Example 3.6 (Adding New Shoes). By rule (INS), the insertion action of Example 2.1 can be executed locally at l_1 as follows.

$$\begin{aligned}
 & \vdash l_1 :: \text{insert}((\text{"001"}, \text{"HB"}, \text{"2015"}, \text{"white"}, \text{"37"}, 6, 0), \text{KLD}@l_1).nil_P \parallel l_1 :: (I_1, R_1) \\
 & \rightarrow l_1 :: nil_P \parallel l_1 :: (I_1, R'_1),
 \end{aligned}$$

where

$$R'_1 = R_1 \uplus \{(\text{"001"}, \text{"HB"}, \text{"2015"}, \text{"white"}, \text{"37"}, 6, 0)\}.$$

This is because $(\text{"001"}, \text{"HB"}, \text{"2015"}, \text{"white"}, \text{"37"}, 6, 0) \models I_1.sk$ holds. \square

Recall that the action $\text{delete}(\text{tid}@l, T, \psi)$ deletes from table tid at l all rows properly binding the pattern T such that the predicate ψ is satisfied. The rule (DEL) in Figure 9 describes the execution of this action from a locality l_1 with l already replaced with some concrete locality l_2 . It is checked in the premise that the specified table identifier tid needs to agree with that of the table (I, R) targeted. In addition, it is examined whether the template T satisfies the schema $I.sk$, and for each row t , whether the evaluation of the predicate ψ is erroneous under the substitution produced from the match. If the satisfaction of $I.sk$ fails, or the evaluation of ψ is sometimes erroneous, then the net ERR is produced;

otherwise the deletion operation is carried out normally, and the rows that do not match the pattern T or do not satisfy the condition ψ will constitute the resulting data set of the target. Note that it is most natural for the “where” clause in the rule (DEL) to follow the resulting net $l_1 :: P \parallel l_2 :: (I, R')$. However, we have placed it below “otherwise” for space reasons, and will follow the same style for some other semantic rules to come.

Example 3.7 (Deleting Existing Shoes). By rule (DEL), the deletion action of Example 2.2 can be executed locally at l_1 as follows, targeting the table resulting from the insertion action of Example 3.6.

$$\begin{aligned} & \vdash l_1 :: \text{delete}(\text{KLD}@l_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), \\ & \quad tp = \text{“HB”} \wedge cr = \text{“white”} \wedge sz = \text{“37”}).\text{nil}_P \parallel l_1 :: (I_1, R'_1) \\ & \rightarrow l_1 :: \text{nil}_P \parallel l_1 :: (I_1, R_1). \end{aligned}$$

This reflects that the original table is recovered after the deletion. \square

Selection, Update and Aggregation. Recall that the action $\text{select}(\widetilde{TB}, T, \psi, t, !tbv)$ picks from the Cartesian product of the data sets of tables identified by the list elements of \widetilde{TB} all rows that properly bind the template T such that the predicate ψ is satisfied, and produces a new table (bound to tbv) that contains the instantiations of t with these bindings. The rule (SEL) in Figure 10 describes the execution of this action at locality l_0 . It is checked that for each table reference of the form $tid@l$ in \widetilde{TB} with some table identifier tid and locality l , the referenced table is one of those in $\langle l_i :: (I_i, R_i) \rangle_i$. In addition, it is examined whether the template T satisfies the product of all the schemas of the tables in \widetilde{TB} , and whether for each tuple t' in the Cartesian product of the data sets of tables in \widetilde{TB} , the evaluation of both the predicate ψ and the tuple t under the substitution t'/T can be performed successfully. In the case of schema mismatches or evaluation errors, the net ERR is produced, signaling the erroneous situation; otherwise the selection is carried out properly, with the resulting table (I', R') substituted into each occurrence of the table variable tbv used in the continuation P . Here we have \perp as the value for $I'.tid$, representing that a temporary table is generated, and the projection of the product of the schemas of the source tables according to T and t as the value for $I'.sk$. On the other hand, R' consists of all the instantiations of t with the bindings produced by matching the tuples in the product data set against T , and satisfying the predicate ψ .

Example 3.8 (Selection of Shoes in a Certain Color). According to the rule (SEL), the selection action of Example 2.3 can be executed from the head office (at l_0) as follows:

$$\begin{aligned} & \vdash l_0 :: \text{select}(\text{KLD}@l_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), \\ & \quad id = \text{“001”} \wedge tp = \text{“HB”} \wedge cr \neq \text{“red”}, (cr, sz, ss), !tbv).\text{nil}_P \parallel l_1 :: (I_1, R_1) \\ & \rightarrow l_0 :: \text{nil}_P \parallel l_1 :: (I_1, R_1). \end{aligned}$$

The condition $I_1.tid = \text{KLD} \wedge l_1 = l_1$ as required by the premise of the rule is satisfied. The I' for the resulting table is such that $I'.tid = \perp$ and $I'.sk = \text{String} \times \text{Int} \times \text{Int}$. The table variable tbv is replaced by (I', R') , for some $R' = \{(\text{“black”}, \text{“38”}, 2), (\text{“black”}, \text{“37”}, 2)\}$. Note that the selection is completed in one *atomic* transition step, which implies, e.g., that no concurrent update to the table KLD at l_1 can be performed in the process. \square

Figure 10: The Semantics for Selection, Update and Aggregation

$$\begin{aligned}
(\text{SEL}) \quad & \frac{\forall tid, l : tid @ l \in \{\widetilde{TB}\} \Rightarrow (\exists j \in \{1, \dots, k\} : I_j.tid = tid \wedge l_j = l)}{N_{\text{env}} \vdash l_0 :: \text{select}(\widetilde{TB}, T, \psi, t, !tbv).P \parallel l_1 :: (I_1, R_1) \parallel \dots \parallel l_k :: (I_k, R_k) \rightarrow N'} \\
\text{where } N' = & \begin{cases} \text{ERR} & \text{if } T \not\models \otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) \vee \\ & \exists t' \in \otimes_{\text{R}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) : \\ & \forall \sigma = t'/T : (\sigma = \text{err} \vee \mathcal{E}[\psi\sigma] = \text{err} \vee \mathcal{E}[t\sigma] = \text{err}) \\ l_0 :: P\sigma' \parallel N'' & \text{otherwise} \\ \text{where} & \\ I' = (\perp, \otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) \downarrow_t^T) \wedge & \\ R' = \{\mathcal{E}[t\sigma] \mid \exists t' \in \otimes_{\text{R}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) : & \\ \quad t'/T = \sigma \wedge \mathcal{E}[\psi\sigma] = tt\} \wedge & \\ \sigma' = [(I', R')/tbv] \wedge & \\ N'' = l_1 :: (I_1, R_1) \parallel \dots \parallel l_k :: (I_k, R_k) & \end{cases} \\
(\text{UPD}) \quad & \frac{I.tid = tid}{N_{\text{env}} \vdash l_1 :: \text{update}(tid @ l_2, T, \psi, t).P \parallel l_2 :: (I, R) \rightarrow N'} \\
\text{where } N' = & \begin{cases} \text{ERR} & \text{if } T \not\models I.sk \vee \\ & \exists t' \in R : \\ & \forall \sigma = t'/T : (\sigma = \text{err} \vee \mathcal{E}[\psi\sigma] = \text{err} \vee \mathcal{E}[t\sigma] = \text{err} \vee \\ & \quad \mathcal{E}[\psi\sigma] = tt \wedge \mathcal{E}[t\sigma] \not\models I.sk) \\ l_1 :: P \parallel l_2 :: (I, R'_1 \uplus R'_2) & \text{otherwise} \\ \text{where } R'_1 = \{t' \in R \mid \mathcal{E}[\psi(t'/T)] \neq tt\} \wedge & \\ R'_2 = \{\mathcal{E}[t\sigma] \mid \exists t' : t' \in R \wedge t'/T = \sigma \wedge \mathcal{E}[\psi\sigma] = tt\} & \end{cases} \\
(\text{AGR}) \quad & \frac{I.tid = tid}{N_{\text{env}} \vdash l_1 :: \text{aggr}(tid @ l_2, T, \psi, f, T').P \parallel l_2 :: (I, R) \rightarrow N'} \\
\text{where } N' = & \begin{cases} \text{ERR} & \text{if } T \not\models I.sk \vee \\ & \exists t \in R : (t/T = \text{err} \vee \mathcal{E}[\psi(t/T)] = \text{err} \vee \\ & \quad \exists \tau_1, \tau_2 : (f : \tau_1 \text{ mset} \rightarrow \tau_2 \wedge (t \not\models \tau_1 \vee T' \not\models \tau_2))) \\ l_1 :: P\sigma' \parallel l_2 :: (I, R) & \text{otherwise} \\ \text{where } t' = f(\{t \in R \mid \mathcal{E}[\psi(t/T)] = tt\}) \wedge & \\ \sigma' = t'/T' & \end{cases}
\end{aligned}$$

In the example above, the substitution of the result table (I', R') into the table variable tbv does not affect the trivial continuation nil_P , in which tbv is not used. The use of result tables for selection actions will be demonstrated in our case study to be presented in Section 5. The following example illustrates selection from the join of multiple tables.

Example 3.9 (Comparing Inventories). The query below generates all pairs of shop names such that in each pair, both shops are in the same city and the brands sold in the first constitute a proper subset of the brands sold in the second.

$\text{select}(\{\text{Stores}@1_0, \text{Stores}@1_0\}, (!x, !y, !z, !w, !p, !x', !y', !z', !w', !p'), x = x' \wedge w \subset w', (z, z'), !tbv)$

This action needs to be executed with the parallel component $1_0 :: (I_0, R_0)$. The result of $\otimes_{sk}(\langle \text{Stores}@1_0, \text{Stores}@1_0 \rangle, \langle l_0 :: (I_0, R_0) \rangle)$ is $\text{flatten}_s(I_0.sk \times I_0.sk)$ and the result of $\otimes_R(\langle \text{Stores}@1_0, \text{Stores}@1_0 \rangle, \langle l_0 :: (I_0, R_0) \rangle)$ is $\text{flatten}_d(R_0 \times R_0)$. \square

The task accomplished in Example 3.9 can be achieved in SQL using selection with an *inner join* operation. Although this example only illustrates joining tables at the same locality, in Klaim-DB, the join of multiple tables at different localities is generally allowed.

Recall that the action $\text{update}(tid@l, T, \psi, t)$ replaces each row matching T and satisfying ψ in table tid at l with a new row t , while leaving the rest of the rows unchanged. The rule (UPD) in Figure 10 describes the execution of this action at l_1 with l already replaced with some constant locality l_2 . The premise checks that the specified table identifier matches the actual one. In addition, it is examined if the template T matches the schema $I.sk$, and for all tuples t' in R , producing substitution σ when matched against T , whether $\psi\sigma$ and $t\sigma$ evaluates properly to values, and whether the evaluation result of $t\sigma$ satisfies the schema $I.sk$, whenever $\psi\sigma$ evaluates to tt . In case of any evaluation error or schema mismatch, the net *ERR* is produced, signaling the abnormal situation; otherwise the update is carried out properly. In more detail, a row t' of R is updated only if it matches the template T , resulting in the substitution σ that makes the predicate ψ satisfied. The row t' is then updated by applying the substitution σ .

Example 3.10 (Update of Shoes Information). According to the rule (UPD), the update action of Example 2.4 can be executed from 1_1 as follows:

$$\begin{aligned} & \vdash 1_1 :: \text{update}(\text{KLD}@1_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), \\ & \quad tp = \text{"HB"} \wedge cr = \text{"red"} \wedge sz = \text{"37"}, (id, tp, yr, cr, sz, is - 2, ss + 2)).\text{nil}_P \parallel 1_1 :: (I, R) \\ & \rightarrow 1_1 :: \text{nil}_P \parallel 1_1 :: (I, R'_1 \uplus R'_2). \end{aligned}$$

The multiset R'_1 consists of all the entries that are intact — shoes that are not red high boots sized 37, while R'_2 contains all the updated items. \square

Recall that the action $\text{aggr}(tid@l, T, \psi, f, T')$ applies the aggregator function f on the multiset of all rows matching T and satisfying ψ in table tid (at l) and binds the aggregation result to the pattern T' . The rule (AGR) in Figure 10 describes the performance of this action from the locality l_1 , with l replaced with some constant locality l_2 . The aggregation is over the table (I, R) . The matching of localities and table identifiers is still required. In addition, it is checked whether the template T matches the schema $I.sk$, and for all rows t in R , whether t matches T , whether the evaluation of ψ can be performed properly under the substitution produced by t/T , and whether t and the template T' are respectively well-sorted under the domain and range types of the aggregator function f . If any of these additional checks fails, then the net *ERR* is produced, signaling the error that occurred; otherwise the aggregation is carried out properly by applying f to the multiset of all tuples matching T , resulting in a binding that satisfies the predicate ψ . The result t' is bound to the specified template T' , producing a substitution that is further applied to the continuation P .

Example 3.11 (Aggregation of Sales Figures). According to the rule (AGR), the aggregation action of Example 2.5 can be executed from the head office as follows:

$$\begin{aligned} & \vdash 1_1 :: \text{aggr}(\text{KLD}@1_1, (!id, !tp, !yr, !cr, !sz, !is, !ss), id = \text{"001"}, \text{sum}_7, !res).\text{nil}_P \parallel 1_1 :: (I_1, R_1) \\ & \rightarrow 1_1 :: \text{nil}_P \parallel 1_1 :: (I_1, R_1). \end{aligned}$$

The variable $!res$ is then bound to the integer value $12 (2 + 5 + 1 + 2 + 2)$. \square

Example 3.12 (Selection using Aggregation Results). Consider the query from l_0 that selects the colors, sizes and sales of all types of high boots whose sales figures are above average. This query can be modeled as a sequence of actions at l_0 , as follows.

$l_0 :: \text{aggr}(\text{KLD}@l_1, T_0, tp = \text{"HB"}, \text{avg}_7, !res). \text{select}(\text{KLD}@l_1, T'_0, ss' \geq res, (cr', sz', ss'), !tbv). \text{nil}_P$
 where $T_0 = (!id, !tp, !yr, !cr, !sz, !is, !ss)$, $T'_0 = (!id', !tp', !yr', !cr', !sz', !is', !ss')$, and

$$\text{avg}_7 = \lambda R. \frac{\text{sum}_7(R)}{|R|} = \lambda R. \left(\frac{\text{sum}(\{v_7 \mid (v_1, \dots, v_7) \in R\})}{|R|} \right).$$

In other words, avg_7 is a function from multisets R to unary tuples containing the average value of the 7-th components of the tuples in R . \square

Creating and Dropping Tables. Recall that the action $\text{create}(tid@l, sk)$ creates a table with interface (tid, sk) at locality l . The rule (CRT) in Figure 11 describes the execution of this action from the locality l_1 , with l already instantiated with some locality constant l_2 . It is ascertained that no table having the same identifier tid exists in the environment at the target locality. If this is the case, then a table with the interface $I = (tid, sk)$ and an empty data set is created at the specified locality; otherwise the creation is skipped. The aforementioned check for clashes of table identifiers at the same localities is realized with the help of the function $Lid(\dots)$ that gives the multiset of pairs of localities and table identifiers in nets and components. This function is overloaded on nets and components, and is defined inductively as follows.

$$\begin{aligned} Lid(\text{nil}_N) &= \emptyset & Lid(l::C) &= Lid(l, C) \\ Lid(ERR) &= \emptyset & Lid(l, P) &= \emptyset \\ Lid(N_1 || N_2) &= Lid(N_1) \uplus Lid(N_2) & Lid(l, (I, R)) &= \{(l, I.id)\} \\ Lid((\nu l)N) &= Lid(N) & Lid(l, C_1 | C_2) &= Lid(l, C_1) \uplus Lid(l, C_2) \end{aligned}$$

Recall that the action $\text{drop}(tid@l)$ drops the table with identifier tid at locality l . The rule (DRP) in Figure 11 describes the performance of this action from the locality l_1 , with l already replaced with some constant locality l_2 . It checks that a table with the specified identifier tid does exist at the specified locality l_2 . Then the table is dropped by replacing it with nil_P .

Figure 11: The Semantics for Creating and Dropping Tables

$$\begin{aligned} \text{(CRT)} \quad N_{\text{env}} \vdash l_1 :: \text{create}(tid@l_2, sk).P \parallel l_2 :: \text{nil}_P &\rightarrow N' \\ \text{where } N' &= \begin{cases} l_1 :: P \parallel l_2 :: ((tid, sk), \emptyset) & \text{if } (l_2, tid) \notin Lid(N_{\text{env}}) \\ l_1 :: P \parallel l_2 :: \text{nil}_P & \text{otherwise} \end{cases} \\ \text{(DRP)} \quad \frac{I.tid = tid}{N_{\text{env}} \vdash l_1 :: \text{drop}(tid@l_2).P \parallel l_2 :: (I, R) \rightarrow l_1 :: P \parallel l_2 :: \text{nil}_P} & \end{aligned}$$

Example 3.13 (Creating Table for Turnovers). Consider the net N_r , which is

$$l_0 :: ((I_0, R_0)|C''_0) || l_1 :: ((I_1, R_1)|C'_1) || \dots || l_n :: ((I_n, R_n)|C'_n),$$

where C''_0 is the residual component at l_0 — the locality of the head office. Suppose $(l_0, \text{Turnover}) \notin Lid(N_r)$, which indicates that there is no table with identifier **Turnover** currently existing at l_0 . According to the rule (CRT), the table creation action of Example 2.6

can be executed from the head office, with N_r as the environment, as follows:

$$\begin{aligned} N_r \vdash 1_0 :: \text{create}(\text{Turnover}@1_0, \text{String} \times \text{Int}).\text{nil}_P \parallel 1_0 :: \text{nil}_P \\ \rightarrow 1_0 :: \text{nil}_P \parallel 1_0 :: ((\text{Turnover}, \text{String} \times \text{Int}), \emptyset). \end{aligned}$$

□

Code Mobility. Recall that the action $\text{eval}(P)@l$ spawns the process P at the locality l . The rule (EVL) describes the execution of this action at the locality l_1 with l already replaced by the concrete locality constant l_2 . It is ensured that P is a closed process that does not contain *free variables* before the action can be performed.

$$(\text{EVL}) \quad N_{\text{env}} \vdash l_1 :: \text{eval}(P)@l_2.P' \parallel l_2 :: \text{nil}_P \rightarrow l_1 :: P' \parallel l_2 :: P \quad \text{if } fv(P) = \emptyset$$

In Klaim-DB, it is possible to spawn multiple processes that migrate among different sets of localities to accomplish diverse database query and management tasks. In this way, the parallelization and distribution of operations on databases is facilitated.

Semantics for Processes and Nets. Directing our attention to Figure 12, the rules (FOR^{tt}) and (FOR^{ff}) specify the continuation and completion of **foreach** loops, respectively. The rule (FOR^{tt}) says that if there are still more tuples in the multiset R matching the pattern specified by the template T , such that ψ holds under the resulting substitution, then we first execute one round of the loop, instantiating variables in T with a (non-deterministically chosen) minimal matching tuple t_0 , and then continue with the remaining rounds of the loop by removing t_0 from R . In the premise, $\text{Minimal}(R, \leq)$ gives the set of all tuples that are minimal in R with respect to the partial order \leq :

$$\text{Minimal}(R, \leq) = \{t \in R \mid \forall t' \in R : t' \leq t \Rightarrow t' = t\}.$$

The rule (FOR^{ff}) says that if there are no more tuples in R matching T , then the loop is completed.

The rules (SEQ^{tt}) and (SEQ^{ff}) describe the transitions that can be performed by sequential compositions $P_1; P_2$. In particular, the rule (SEQ^{tt}) accounts for the case where P_1 cannot finish after one more step; whereas the rule (SEQ^{ff}) accounts for the opposite case. Note that by our language design, the variables declared in P_1 have local scopes; hence there is no need to apply any substitution on P_2 when the execution embarks on it.

The rule (PAR) says that if a net N_1 can make a transition assuming a net N_2 as the environment, then the parallel composition of N_1 with N_2 can also make a transition, with an empty environment (nil_N). Having an empty environment in the conclusion indicates that the rule cannot be used multiple times with non-trivial effects. The idea is: to build a transition from a net with multiple parallel components, we use the structural congruence to extrude all the $(\nu l)(\cdot)$ to the outer level, and collect all the unaffected parallel nets into N_2 . We can then build the transition using (PAR) only once, combined with further uses of (RES), and of (EQUIV).

The rules (CALL), (RES) and (EQUIV) are self-explanatory.

It is a property of our semantics that no local repetition of table identifiers can be caused by a transition. Define $\text{no_rep}(A) = (A = \text{set}(A))$, which expresses that there is no repetition in the multiset A , thus A coincides with its underlying set. This property is formalized in Lemma 3.14 whose proof can be found in Appendix B.

Lemma 3.14. *If $\text{no_rep}(\text{Lid}(N) \uplus \text{Lid}(N_{\text{env}}))$ and $N_{\text{env}} \vdash N \rightarrow N'$, then it holds that $\text{no_rep}(\text{Lid}(N') \uplus \text{Lid}(N_{\text{env}}))$.*

Figure 12: The Semantics for Processes and Nets

$$\begin{array}{c}
\text{(FOR}^{\text{tt}}\text{)} \frac{t_0 \in \text{Minimal}(R, \leq) \quad t_0/T \neq \text{err} \quad \mathcal{E}[\![\psi(t_0/T)]\!] = tt}{N_{\text{env}} \vdash l :: \text{foreach}((I, R), T, \psi, \leq) : P \rightarrow l :: P(t_0/T); (\text{foreach}((I, R \setminus \{t_0\}), T, \psi, \leq) : P)} \\
\\
\text{(FOR}^{\text{ff}}\text{)} \frac{\neg(\exists t_0 \in R : \mathcal{E}[\![\psi(t_0/T)]\!] = tt)}{N_{\text{env}} \vdash l :: \text{foreach}((I, R), T, \psi, \leq) : P \rightarrow N'} \\
\text{where } N' = \begin{cases} \text{ERR} & \text{if } \exists t_0 \in R : (t_0/T = \text{err} \vee \mathcal{E}[\![\psi(t_0/T)]\!] = \text{err}) \\ l :: \text{nil}_P & \text{otherwise} \end{cases} \\
\\
\text{(SEQ}^{\text{tt}}\text{)} \frac{N_{\text{env}} \vdash l :: P_1 \parallel N \rightarrow l :: P'_1 \parallel N'}{N_{\text{env}} \vdash l :: P_1; P_2 \parallel N \rightarrow l :: P'_1; P_2 \parallel N'} \quad \text{if } P'_1 \neq \text{nil}_P \\
\\
\text{(SEQ}^{\text{ff}}\text{)} \frac{N_{\text{env}} \vdash l :: P_1 \parallel N \rightarrow l :: \text{nil}_P \parallel N'}{N_{\text{env}} \vdash l :: P_1; P_2 \parallel N \rightarrow l :: P_2 \parallel N'} \\
\\
\text{(CALL)} \quad N_{\text{env}} \vdash l :: A(\tilde{e}) \rightarrow l :: P[v_1/\text{var}_1] \dots [v_n/\text{var}_n] \\
\text{if } A(\langle \text{var}_i : \tau_i \rangle_{i \leq n}) \triangleq P \wedge \mathcal{E}[\![\tilde{e}]\!] = \tilde{v} \\
\\
\text{(PAR)} \quad \frac{N_2 \vdash N_1 \rightarrow N'_1}{\vdash N_1 \parallel N_2 \rightarrow N'_1 \parallel N_2} \\
\\
\text{(RES)} \quad \frac{N_{\text{env}} \vdash N \rightarrow N'}{N_{\text{env}} \vdash (\nu l)N \rightarrow (\nu l)N'} \quad \text{(EQUIV)} \quad \frac{N_1 \equiv N_2 \quad N_{\text{env}} \vdash N_2 \rightarrow N_3 \quad N_3 \equiv N_4}{N_{\text{env}} \vdash N_1 \rightarrow N_4}
\end{array}$$

Thus, imposing “non-existence of local repetition of table identifiers” as an integrity condition for the initial net will guarantee the satisfaction of this condition in all possible derivatives of the net.

The following example illustrates a transition of a net obtained from one of an action.

Example 3.15 (Transition of Nets). Continuing with Example 3.13, with the help of the rules (PAR), and (EQUIV), we can derive

$$\begin{aligned}
& \vdash \mathbf{1}_0 :: (I_0, R_0) | \text{create}(\text{Turnover} @ \mathbf{1}_0, \text{String} \times \text{Int}). \text{nil}_P | C''_0 \parallel \mathbf{1}_1 :: (I_1, R_1) | C'_1 \parallel \dots \parallel \mathbf{1}_n :: (I_n, R_n) | C'_n \\
& \rightarrow \mathbf{1}_0 :: (I_0, R_0) | ((\text{Turnover}, \text{String} \times \text{Int}), \emptyset) | C''_0 \parallel \mathbf{1}_1 :: (I_1, R_1) | C'_1 \parallel \dots \parallel \mathbf{1}_n :: (I_n, R_n) | C'_n. \quad \square
\end{aligned}$$

4. TYPE SYSTEM

It is desirable to eliminate certain inherent discrepancies in a piece of specification before ever executing it. In this section we develop a type system that scans specifications for discrepancies that lead to the run-time errors described by the semantics in Section 3. We now detail the types of such errors that we alluded to in the introduction:

- inconsistency between the format of a template T and a schema sk , such as when T is used in a selection from a table having the schema sk , T has three components while sk involves only two basic datatypes;

- inconsistency between the format of a tuple t and a schema sk , such as when t is to be inserted into a table having the schema sk , the second component of t is an integer, but is supposed to be a string according to sk ;
- mismatch of the signature of an aggregator function against the format of the data source;
- ill-formed expressions and predicates, such as a multiplication of two strings, or a predicate that cannot be evaluated to a boolean value;
- use of unbound (undefined) variables.

A typing environment Γ is a finite sequence of bindings $var : \tau$ of variables to types. For the variables in the domain of Γ (written \mathbf{D}_Γ), $\Gamma(x)$ is the data type bound to x , $\Gamma(tbv)$ is the schema bound to tbv , and $\Gamma(u)$ is Loc — the only type that can be bound to locality variables u . We assume that at most one binding can be present for each variable in Γ , which creates no essential restriction since bound variables can be renamed apart. We also write \square for the empty typing environment, $(\Gamma, var : \tau)$ for the *extension* of Γ with the binding $var : \tau$, and similarly Γ_1, Γ_2 for the *extension* of Γ_1 with (all bindings of) Γ_2 .

We assume that only one schema can be used for each table identifier at all localities, all the time. This assumption may cause slight peculiarities concerning the use of table identifiers, but avoids the potential explosion caused by over-approximating the sets of localities evaluated to by locality variables. Correspondingly, we introduce a partial function ∇ from table identifiers to their corresponding schemas.

The typing judgments for expressions, predicates and tuples are as follows.

$$\begin{aligned}\Gamma &\vdash e \triangleright \tau \\ \Gamma &\vdash \psi \triangleright Bool \\ \Gamma &\vdash t \triangleright \tau\end{aligned}$$

The environment Γ provides the types for the constituent variables of these syntactical entities, and the types τ derived from them are placed after the separator \triangleright . A well-typed predicate always has the type $Bool$.

The judgment for templates is:

$$\tau \vdash T \triangleright \Gamma$$

Due to the involvement of variable binding, a typing environment Γ is derived instead of a type. In more detail, τ is a product type and Γ records the association of each variable in T to the corresponding component type in τ . When this judgment is used, τ is often the schema of a table or the join of some tables, and deriving Γ from τ captures that the bound variables of T obtain their types from the actual types of the fields of the table.

The judgment for tables is:

$$\Gamma, \nabla \vdash TB \triangleright \tau$$

Both the variable environment Γ and the partial function ∇ are needed — depending on the particular form of TB , its schema τ is derived from either Γ or ∇ . A frequent usage pattern is that of deriving from a table its schema τ , and then using the derived τ to type a template T used in an operation to match against the rows of the table.

The judgment for actions is:

$$\Gamma, \nabla \vdash a \triangleright \Gamma'$$

Due to the potential involvement of tables in actions, the partial function ∇ is again needed. Actions such as aggregation can create bindings to be used in the continuation. Correspondingly, the derived environment Γ associates the variables in the domain of such bindings with their types.

Finally, the judgments for processes, components and nets are listed below.

$$\Gamma, \nabla \vdash P \quad \Gamma, \nabla \vdash C \quad \Gamma, \nabla \vdash N$$

Similar to the previous cases, Γ and ∇ are needed for obtaining the types of variables and tables referenced with identifiers, respectively. These judgments only assert that a process, component, or net is well-typed, without deriving a type.

4.1. Typing Rules. The typing rules for expressions, predicates, tuples, templates and tables are shown in Figure 13. The typing rules for actions are then given in Figure 14. Finally, the typing rules for processes, components, and nets are displayed in Figure 15.

Figure 13: Typing Rules for Expressions, Predicates, Tuples, Templates and Tables

Expressions

$$\begin{array}{l} \Gamma \vdash x \triangleright \tau \quad \text{if } \Gamma(x) = \tau \wedge \tau \neq Loc \quad \Gamma \vdash u \triangleright Loc \quad \text{if } \Gamma(u) = Loc \\ \Gamma \vdash num \triangleright Int \quad \Gamma \vdash str \triangleright String \quad \Gamma \vdash tid \triangleright Id \quad \Gamma \vdash l \triangleright Loc \\ \frac{\Gamma \vdash e_1 \triangleright String \quad \Gamma \vdash e_2 \triangleright String}{\Gamma \vdash e_1 \cdot e_2 \triangleright String} \quad \frac{\Gamma \vdash e_1 \triangleright Int \quad \Gamma \vdash e_2 \triangleright Int}{\Gamma \vdash e_1 \text{ aop } e_2 \triangleright Int} \\ \frac{\Gamma \vdash e_1 \triangleright \tau_d \dots \Gamma \vdash e_n \triangleright \tau_d}{\Gamma \vdash \{e_1, \dots, e_n\} \triangleright \tau_d \text{ mset}} \end{array}$$

Predicates

$$\begin{array}{l} \Gamma \vdash true \triangleright Bool \\ \frac{\Gamma \vdash e_1 \triangleright \tau_d \quad \Gamma \vdash e_2 \triangleright \tau_d}{\Gamma \vdash e_1 \text{ cop } e_2 \triangleright Bool} \quad \frac{\Gamma \vdash e_1 \triangleright \tau_d \quad \Gamma \vdash e_2 \triangleright \tau_d \text{ mset}}{\Gamma \vdash e_1 \in e_2 \triangleright Bool} \\ \frac{\Gamma \vdash \psi \triangleright Bool}{\Gamma \vdash \neg \psi \triangleright Bool} \quad \frac{\Gamma \vdash \psi_1 \triangleright Bool \quad \Gamma \vdash \psi_2 \triangleright Bool}{\Gamma \vdash \psi_1 \wedge \psi_2 \triangleright Bool} \end{array}$$

Tuples

$$\frac{\Gamma \vdash e_1 \triangleright \tau_1 \quad \dots \quad \Gamma \vdash e_n \triangleright \tau_n}{\Gamma \vdash e_1, \dots, e_n \triangleright \tau_1 \times \dots \times \tau_n}$$

Templates

$$\begin{array}{l} \tau_m \vdash !x \triangleright [x : \tau_m] \quad \text{if } \tau_m \neq Loc \\ Loc \vdash !u \triangleright [u : Loc] \quad \frac{\tau_1 \vdash W_1 \triangleright \Gamma_1 \quad \dots \quad \tau_n \vdash W_n \triangleright \Gamma_n}{\tau_1 \times \dots \times \tau_n \vdash W_1, \dots, W_n \triangleright \Gamma_1, \dots, \Gamma_n} \end{array}$$

Tables

$$\begin{array}{l} \frac{\Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash tid @ \ell \triangleright \nabla(tid)} \quad \text{if } tid \in \mathbf{D}_\nabla \quad \Gamma, \nabla \vdash tbv \triangleright \Gamma(tbv) \quad \text{if } tbv \in \mathbf{D}_\Gamma \\ \frac{\forall t \in R : \Gamma \vdash t \triangleright I.sk}{\Gamma, \nabla \vdash (I, R) \triangleright I.sk} \quad \text{if } I.tid \neq \perp \Rightarrow \nabla(I.tid) = I.sk \end{array}$$

Typing Expressions, Predicates, Tuples, Templates and Tables. Directing our attention to Figure 13, for expressions we obtain the types of data variables and locality variables directly from the typing environment Γ (first and second rule). For numerals, string literals, table identifiers and locality constants, we assume that it is clear which types they are from their literal appearance (next four rules). The concatenation of two expressions of type *String* has type *String*, and an arithmetic expression involving sub-expressions of type *Int* has type *Int* (next two rules). The type τ *mset* can be derived from a multiset only if all the elements have the same type τ (last rule for expressions).

For predicates, a comparison $e_1 \text{ cop } e_2$ has type *Bool* if its sub-expressions e_1 and e_2 have the same type. A membership test $e_1 \in e_2$ has type *Bool* if the type of e_2 reflects that it is a multiset of elements of the type that e_1 has. The composite predicates $\neg\psi$ and $\psi_1 \wedge \psi_2$ are well-typed if their sub-predicates are well-typed.

The type of a tuple e_1, \dots, e_n is the product of the types of its component expressions.

For templates, a singleton template $!x$ for a data variable x can be typed given a multiset type τ_m that is not *Loc*, with the environment $[x : \tau_m]$ derived. This reflects the fact that when the type of a bound variable is not declared with it in the syntax, we obtain the type of the variable from the type τ_m of its data source. For a locality variable, the only possibility is to derive the environment $[u : \text{Loc}]$, provided that the type *Loc* is given to the left of the turnstile. Finally, for a composite template W_1, \dots, W_n , a product type $\tau_1 \times \dots \times \tau_n$ should be given, each component W_j needs to be typed under τ_j , and the bindings from the derived typing environments $\Gamma_1, \Gamma_2 \dots$, and Γ_n constitute the typing environment derived from the template W_1, \dots, W_n .

Example 4.1. The template $(!id, !tp, !yr, !cr, !sz, !is, !ss)$ that can be used to match against the records of the table **KLD** (such as in the selection operation of Example 3.8) can be typed under the schema of the table (on the left of the turnstile) as follows.

$$\begin{aligned} &String \times String \times String \times String \times String \times Int \times Int \vdash (!id, !tp, !yr, !cr, !sz, !is, !ss) \triangleright \\ &[id : String, tp : String, yr : String, cr : String, sz : String, is : Int, ss : Int] \end{aligned} \quad \square$$

The rules for tables distinguish between three cases. If the table to be typed is a reference of the form $tid@l$, then its type is obtained from ∇ as $\nabla(tid)$, as long as tid is in the domain \mathbf{D}_∇ of ∇ , and l is well-typed under the same typing environment. The well-typedness of l ensures that when l is a locality variable u , u exists in the domain of the typing environment, which in turn indicates that u must have been bound previously. If the table is a variable tbv in the domain of the typing environment Γ , then a binding of tbv to a schema is supposed to exist in Γ , and the schema is obtained from Γ as the type of tbv . Finally, if the table is a concrete one, (I, R) , then its schema is directly obtained from I . However, it needs to be checked that each tuple in R indeed has the schema $I.sk$ as its type, and in case the table identifier $I.tid$ is not \perp , ∇ correctly records the schema $I.sk$ corresponding to this identifier.

Typing Actions. Turning to Figure 14, the rules always type check the localities occurring in the actions to be typed — we do not repeat this aspect in each individual rule that we explain below.

For an *insertion* action to be typed, the tuple t inserted into the table $tid@l$ needs to have the type $\nabla(tid)$ — the schema of all potential tables referenced by the identifier tid .

For a *deletion* action to be typed, the specified template T needs to match the structure of the target table, thus it needs to be well-typed under the schema $\nabla(tid)$ of that table. A

Figure 14: The Typing Rules for Actions

$$\begin{array}{c}
\frac{\Gamma \vdash t \triangleright \nabla(tid) \quad \Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash \text{insert}(t, tid@l) \triangleright []} \\
\\
\frac{\nabla(tid) \vdash T \triangleright \Gamma' \quad \Gamma, \Gamma' \vdash \psi \triangleright Bool \quad \Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash \text{delete}(tid@l, T, \psi) \triangleright []} \\
\\
\frac{\Gamma, \nabla \vdash TB_1 \triangleright \tau_1 \quad \dots \quad \Gamma, \nabla \vdash TB_n \triangleright \tau_n \quad \text{flatten}_s(\tau_1 \times \dots \times \tau_n) \vdash T \triangleright \Gamma' \quad \Gamma, \Gamma' \vdash \psi \triangleright Bool \quad \Gamma, \Gamma' \vdash t \triangleright \tau'}{\Gamma, \nabla \vdash \text{select}(\widetilde{TB}, T, \psi, t, !tbv) \triangleright [tbv : \tau']} \\
\\
\frac{\nabla(tid) \vdash T \triangleright \Gamma' \quad \Gamma, \Gamma' \vdash \psi \triangleright Bool \quad \Gamma, \Gamma' \vdash t \triangleright \nabla(tid) \quad \Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash \text{update}(tid@l, T, \psi, t) \triangleright []} \\
\\
\frac{\nabla(tid) \vdash T \triangleright \Gamma' \quad \Gamma, \Gamma' \vdash \psi \triangleright Bool \quad f : \nabla(tid) \text{ mset} \rightarrow \tau' \quad \tau' \vdash T' \triangleright \Gamma'' \quad \Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash \text{aggr}(tid@l, T, \psi, f, T') \triangleright \Gamma''} \\
\\
\frac{\Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash \text{create}(tid@l, \nabla(tid)) \triangleright []} \quad \frac{\Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash \text{drop}(tid@l) \triangleright []} \\
\\
\frac{\Gamma, \nabla \vdash P \quad \Gamma \vdash \ell \triangleright Loc}{\Gamma, \nabla \vdash \text{eval}(P)@l \triangleright []}
\end{array}$$

typing environment Γ' is then derived, associating the bound variables in T to their types. In addition, under the extension Γ, Γ' of the original typing environment Γ with Γ' , the predicate ψ needs to be well-typed. This extension captures that the free variables of ψ can be partly bound before the deletion action is ever performed and associated with types by Γ , and partly bound in T and associated with types by Γ' .

To type a *selection* action, the specified template T needs to be well-typed under the (flattened) product of the types of all the source tables, to ensure that T matches the structure of the Cartesian product of these source tables. A typing environment Γ' is then derived from typing T , binding variables of T to their types. Under the extension Γ, Γ' of Γ with Γ' , it is then checked that both the predicate ψ and the tuple t should be typable. In particular, the type τ' of t is supposed to be identical to the schema of the result table to be bound to tbv . Correspondingly the environment $tbv : \tau'$ is derived.

The typing rule for the *update* action can be understood similarly. Since the instantiations of the tuple t are put back in the target table with identifier tid that already exists, t is supposed to have the schema $\nabla(tid)$ as its type, and the empty environment $[]$ is derived since no variables bound in this action are further used in the continuation.

In the typing rule for *aggregation* actions, a similar pattern is involved in the typing of the specified template T and predicate ψ . Since the aggregator function f takes a multiset of tuples in the data set of the table tid as its argument, the type of f is supposed to be $\nabla(tid) \text{ mset} \rightarrow \tau'$, for some type τ' of the aggregation result. Since this result is further bound to T' , the template T' needs to be well-typed under τ' , producing a new typing environment Γ' that is also the environment derived in the typing of the whole action.

The action for *table creation* is typed by ensuring that the schema of the table to be created at ℓ with table identifier tid is in accordance with what ∇ prescribes.

The typing rule for the action to *drop* a table does not require any constraint to be satisfied: a finer analysis may be able to give an approximation of whether the table currently exists; for simplicity we stick to the scope of errors listed in the beginning of this section.

Finally, the typing rule for the action $\text{eval}(P)@l$ type checks P in the same environment Γ as for $\text{eval}(P)@l$, which corresponds to the static scoping of the free variables of P .

For examples on typing actions we refer the reader to Section 5.

Figure 15: The Typing Rules for Processes, Components, and Nets

Processes	$\Gamma, \nabla \vdash \text{nil}_P \quad \frac{\Gamma, \nabla \vdash a \triangleright \Gamma' \quad (\Gamma, \Gamma'), \nabla \vdash P}{\Gamma, \nabla \vdash a.P} \quad \frac{\Gamma, \nabla \vdash P_1 \quad \Gamma, \nabla \vdash P_2}{\Gamma, \nabla \vdash P_1; P_2}$ $\frac{\Gamma \vdash e_1 \triangleright \tau_1 \quad \dots \quad \Gamma \vdash e_n \triangleright \tau_n \quad (\text{var}_1 : \tau_1, \dots, \text{var}_n : \tau_n), \nabla \vdash P \quad \text{where } A(\langle \text{var}_i : \tau_i \rangle_{i \leq n}) \triangleq P}{\Gamma, \nabla \vdash A(\tilde{e})}$ $\frac{\Gamma, \nabla \vdash TB \triangleright \tau \quad \tau \vdash T \triangleright \Gamma' \quad \Gamma, \Gamma' \vdash \psi \triangleright \text{Bool} \quad (\Gamma, \Gamma'), \nabla \vdash P}{\Gamma, \nabla \vdash \text{foreach}(TB, T, \psi, \leq) : P}$
Components	$\frac{\forall t \in R : \Gamma \vdash t \triangleright I.sk}{\Gamma, \nabla \vdash (I, R)} \quad \text{if } \nabla(I.tid) = I.sk \quad \frac{\Gamma, \nabla \vdash C_1 \quad \Gamma, \nabla \vdash C_2}{\Gamma, \nabla \vdash C_1 C_2}$
Nets	$\Gamma, \nabla \vdash \text{nil}_N \quad \frac{\Gamma \vdash l \triangleright \text{Loc} \quad \Gamma, \nabla \vdash C}{\Gamma, \nabla \vdash l :: C}$ $\frac{\Gamma, \nabla \vdash N}{\Gamma, \nabla \vdash (\nu l)N} \quad \frac{\Gamma, \nabla \vdash N_1 \quad \Gamma, \nabla \vdash N_2}{\Gamma, \nabla \vdash N_1 N_2}$

Typing Processes, Components and Nets. Directing our attention to the typing rules for processes in Figure 15, the rule for the *inert process* nil_P is trivial. The rule for *prefixing* says that in order for $a.P$ to be typed, we need to type the action a , deriving some environment Γ' , and the process P needs to be well-typed under the original typing environment Γ extended with Γ' . On the other hand, to type a *sequential composition* $P_1; P_2$, each process only needs to be typed in the original environment Γ , which is because of the local scoping of the variables bound in P_1 . To type a *procedure call* $A(\tilde{e})$ where A has the procedure body P , the process P needs to be typed in the typing environment that binds all the formal parameters of A to their types obtained from the signature of A . Naturally, these types also need to agree with those of the actual arguments in \tilde{e} . The rule for the *foreach loop* features a similar pattern in the treatment of the source table TB , the template T and the predicate ψ , compared to that of certain typing rules for actions such as *select*. In addition, the body P of the loop needs to be typed with the original typing environment Γ extended

with Γ' that reflects the bindings produced by matching tuples from the source table against T .

The next group of typing rules in Figure 15 covers the typing of components. It is worth noting that the rule for typing a *concrete table* (I, R) as a component resembles the one for typing it as a table, except that the case where the table identifier is \perp is impossible and need not be dealt with for a table that stands as a component. Note also that the next rule for typing a *parallel composition* $C_1|C_2$ of components uses the same environments Γ and ∇ for the components C_1 and C_2 as for $C_1|C_2$: since there are no *linearity* issues of concern for the variables, no split operation is needed on Γ , and ∇ is simply a global constant.

The last group of typing rules in Figure 15 covers the typing of nets. The rule for the *empty net* nil_N is trivial. The rule for a *located component* $l :: C$ checks that l is indeed a locality ($\Gamma \vdash l \triangleright \text{Loc}$). The rule for *restriction* $(\nu l)N$ only ensures that the net N is itself typable. This check is performed with the same typing environment Γ since l is not a variable and we have assumed that the types of constants are clear for their literal appearance. Finally, the rule for parallel composition $N_1||N_2$ is conceived with a similar rationale to that of the rule for $C_1|C_2$.

4.2. Theoretical Results. It is usually required of type systems that they should be safe, in the sense of subject reduction [19] — well-typedness is preserved under transitions. In most cases it is important for type systems to be sound, i.e., the guarantees they aim to deliver should indeed be delivered at run time, for well-typed specifications and programs. It is also important for type checking to be efficient, to smooth the development process. In this section, we provide theoretical results demonstrating that our type system has all of the aforementioned properties. In the statements of these results, we will leave universal quantifiers implicit, and assume that it is clear from the use of symbols what kinds of entities they refer to.

The subject reduction result is given in Theorem 4.2. Its proof is structured as an induction on the semantic derivation, with details provided in Appendix B.

Theorem 4.2 (Subject Reduction). *For a closed net N , if $\text{bv}(N) \cap \mathbf{D}_\Gamma = \emptyset$, $\Gamma, \nabla \vdash N$, and $\vdash N \rightarrow N'$, then $\Gamma, \nabla \vdash N'$.*

We characterize a net N being non-erroneous as $\text{ok}(N)$, which says that one is unable to identify a parallel component that is *ERR* in N . The formal definition of $\text{ok}(N)$ is given below.

Definition 4.3. $\text{ok}(N) \triangleq \neg(\exists N' : N \equiv N' || \text{ERR})$

With a straightforward induction on the typing derivation, we can show that if a net N is typable, then it cannot be erroneous.

Lemma 4.4. *If $\Gamma, \nabla \vdash N$, then $\text{ok}(N)$ holds.*

Combining Theorem 4.2 and Lemma 4.4, we can easily obtain the following result that the well-typedness of closed nets implies its non-erroneous execution.

Theorem 4.5 (Soundness). *If $\Gamma, \nabla \vdash N$, and $\vdash N \rightarrow^* N'$, then $\text{ok}(N')$ holds.*

To be able to discuss the time complexity of type checking, a characterization of the sizes of nets is needed. We characterize the size of a net N as the number of ASCII characters

used for the full textual representation of N itself, as well as the definitions of procedures invoked in N :

$$size(N) = sz_asc(N) + \sum_{A(\{var_i:\tau_i\}_i) \triangleq P} sz_asc(P)$$

By “full textual representation”, it is meant that the contents of interfaces I and data sets R for tables need to be completely spelled out.

The following theorem states that the time complexity of type checking a net N is linear in the size of N . Its proof is sketched in Appendix B.

Theorem 4.6 (Efficiency of Type Checking). *With a given ∇ , the time complexity of type checking net N is linear in $size(N)$, provided that it takes $O(1)$ time to*

- *determine the types of all constant expressions,*
 - *decide the equality/inequality of table identifiers and types,*
 - *construct a singleton typing environment,*
 - *construct the extension (Γ_1, Γ_2) of a typing environment Γ_1 with Γ_2 , and*
 - *look up environments Γ (resp. ∇) for variables (resp. table identifiers),*
- and that it takes $O(n)$ time to*
- *form an n -ary product type, and*
 - *perform the operation $flatten_s(\tau_1 \times \dots \times \tau_n)$.*

5. CASE STUDY

Revisiting our scenario with the chain of department stores, we illustrate the modeling of data aggregation over multiple databases local to its different branches in Klaim-DB. In more detail, a manager of the head office wants statistics on the total sales of KLD high boots from the year 2015, in each branch operating in Copenhagen.

We will think of a procedure *stat* at the locality l_0 of the head office, carrying out the aggregation needed. Thus the net N_{DC} for the database systems of the department store chain, as considered in Section 2, specializes to the following, where C_0'' is the remaining tables and processes at l_0 apart from the table **Stores** and the procedure *stat*.

$$l_0 :: ((I_0, R_0)|stat|C_0'') \parallel l_1 :: ((I_1, R_1)|C_1') \parallel \dots \parallel l_n :: ((I_n, R_n)|C_n')$$

A detailed specification of *stat* is then given in Figure 16.

Figure 16: The Procedure for Distributed Data Aggregation

```

stat  $\triangleq$  create(SSResult@l0, String  $\times$  String  $\times$  Int).
  select(Stores@l0, (!x, !y, !z, !w, !p), KLD  $\in$  w  $\wedge$  x = “CPH”, (z, p), !tbv).
  foreach(tbv, (!q, !u), true, {}) :
    aggr(KLD@u, (!id, !tp, !yr, !cr, !sz, !is, !ss), tp = “HB”, sum7, !res).
    insert((q, “HB”, res), SSResult@l0).
  nilP;
...
drop(SSResult@l0).nilP

```

First of all, a result table with identifier **SSResult** and schema $String \times String \times Int$ is created. Then all the localities of the databases used by the branches in Copenhagen

that actually sell KLD shoes are selected, together with the shop names of such branches. This result set is then processed by a **foreach** loop. The number of KLD high boots from 2015 that are sold is counted at each of these localities (branches), and is inserted into the resulting table together with the corresponding shop name and the string “HB” describing the shoe type concerned. The resulting table, displayed in Figure 17, can still be queried/manipulated before being dropped.

Figure 17: The Table **SSResult**

<i>Shop_name</i>	<i>Shoe_type</i>	<i>Sales</i>
Shop1	HB	12
Shop2	HB	53
Shop3	HB	3
...

The typability of N_{DC} is formally stated in Fact 5.1.

Fact 5.1. Suppose

$$\begin{aligned} \nabla_{DC} = & [\text{Stores} \mapsto \text{String} \times \text{String} \times \text{String} \times (\text{Id mset}) \times \text{Loc}] \\ & [\text{KLD} \mapsto \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{Int} \times \text{Int}] \\ & [\text{SSResult} \mapsto \text{String} \times \text{String} \times \text{Int}]. \end{aligned}$$

We have $\square, \nabla_{DC} \vdash N_{DC}$, given $\forall i \in \{1, \dots, n\} : \square, \nabla_{DC} \vdash C'_i$ and $\square, \nabla_{DC} \vdash C''_0$.

Example 5.2. We briefly go through the typing of N_{DC} .

It can be established for all $j \in \{0, 1, \dots, n\}$ that $\square, \nabla_{DC} \vdash (I_j, R_j)$. First of all, with the schemas of these tables given at the end of Section 2, we have $\nabla_{DC}(I_j.tid) = I_j.sk$. Second of all, by the assumption made in the same section that the data stored in these tables have the types specified in the schemas, we have $\forall t \in R_j : \square \vdash t \triangleright I_j.sk$ for all j in range.

It is not difficult to see that under the assumed well-typedness of the components C'_1, \dots, C'_n and C''_0 , it boils down to establishing $\square, \nabla_{DC} \vdash stat$, in order to obtain $\square, \nabla_{DC} \vdash N_{DC}$.

For the action creating the result table we have

$$\frac{\square \vdash 1_0 \triangleright \text{Loc}}{\square, \nabla_{DC} \vdash \text{create}(\text{SSResult}@1_0, \text{String} \times \text{String} \times \text{Int}) \triangleright \square}$$

since it holds that $\text{String} \times \text{String} \times \text{Int} = \nabla_{DC}(\text{SSResult})$.

For the action selecting the pairs of shop names and localities, we have

$$\begin{aligned} & \square, \nabla_{DC} \vdash \text{Stores}@1_0 \triangleright \text{String} \times \text{String} \times \text{String} \times (\text{Id mset}) \times \text{Loc} \\ & \text{String} \times \text{String} \times \text{String} \times (\text{Id mset}) \times \text{Loc} \vdash (!x, !y, !z, !w, !p) \triangleright \\ & \quad [x : \text{String}, y : \text{String}, z : \text{String}, w : \text{Id mset}, p : \text{Loc}] \\ & [x : \text{String}, y : \text{String}, z : \text{String}, w : \text{Id mset}, p : \text{Loc}] \vdash \text{KLD} \in w \wedge x = \text{“CPH”} \triangleright \text{Bool} \\ & [x : \text{String}, y : \text{String}, z : \text{String}, w : \text{Id mset}, p : \text{Loc}] \vdash (z, p) \triangleright \text{String} \times \text{Loc} \\ \hline & \square, \nabla_{DC} \vdash \text{select}(\text{Stores}@1_0, (!x, !y, !z, !w, !p), \text{KLD} \in w \wedge x = \text{“CPH”}, (z, p), !tbv) \triangleright \\ & \quad [tbv : \text{String} \times \text{Loc}] \end{aligned}$$

For the aggregation action producing the sales figures, we have

$$\begin{array}{c}
\nabla_{\text{DC}}(\text{KLD}) \vdash (!id, !tp, !yr, !cr, !sz, !is, !ss) \triangleright \\
[id : \text{String}, tp : \text{String}, yr : \text{String}, cr : \text{String}, sz : \text{String}, is : \text{Int}, ss : \text{Int}] \\
[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}, id : \text{String}, tp : \text{String}, \\
yr : \text{String}, cr : \text{String}, sz : \text{String}, is : \text{Int}, ss : \text{Int}] \vdash tp = \text{"HB"} \triangleright \text{Bool} \\
sum_7 : (\text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{String} \times \text{Int} \times \text{Int}) \text{ mset} \rightarrow \text{Int} \\
\text{Int} \vdash !res \triangleright [res : \text{Int}] \quad [tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}] \vdash u \triangleright \text{Loc} \\
\hline
[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}], \nabla_{\text{DC}} \vdash \\
\text{aggr}(\text{KLD}@u, (!id, !tp, !yr, !cr, !sz, !is, !ss), tp = \text{"HB"}, sum_7, !res) \triangleright [res : \text{Int}]
\end{array}$$

For the action inserting the aggregation result, we have

$$\begin{array}{c}
[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}, res : \text{Int}] \vdash (q, \text{"HB"}, res) \triangleright \text{String} \times \text{String} \times \text{Int} \\
[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}, res : \text{Int}] \vdash \mathbf{1}_0 \triangleright \text{Loc} \\
\hline
[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}, res : \text{Int}], \nabla_{\text{DC}} \vdash \\
\text{insert}((q, \text{"HB"}, res), \text{SSResult}@1_0) \triangleright []
\end{array}$$

since it holds that $\text{String} \times \text{String} \times \text{Int} = \nabla_{\text{DC}}(\text{SSResult})$.

We can thus establish

$$[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}], \nabla_{\text{DC}} \vdash \underline{\text{aggr.ins.nil}}_P$$

where the underlined names abbreviate the actions under the same names in the procedure *stat*.

For the **foreach** construct we have

$$\begin{array}{c}
[tbv : \text{String} \times \text{Loc}], \nabla_{\text{DC}} \vdash tbv \triangleright \text{String} \times \text{Loc} \\
\text{String} \times \text{Loc} \vdash (!q, !u) \triangleright [q : \text{String}, u : \text{Loc}] \\
[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}] \vdash \text{true} \triangleright \text{Bool} \\
[tbv : \text{String} \times \text{Loc}, q : \text{String}, u : \text{Loc}], \nabla_{\text{DC}} \vdash \underline{\text{aggr.ins.nil}}_P \\
\hline
[tbv : \text{String} \times \text{Loc}], \nabla_{\text{DC}} \vdash \text{foreach}(tbv, (!q, !u), \text{true}, \emptyset) : \underline{\text{aggr.ins.nil}}_P
\end{array}$$

The typing of the action dropping the result table can be easily established, i.e.,

$$\begin{array}{c}
[] \vdash \mathbf{1}_0 \triangleright \text{Loc} \\
\hline
[], \nabla_{\text{DC}} \vdash \text{drop}(\text{SSResult}@1_0) \triangleright []
\end{array}$$

Using the typing rule for prefixing repeatedly, and the rule for sequential composition, we can obtain:

$$[], \nabla_{\text{DC}} \vdash \underline{\text{create.select.foreach; drop}}$$

Finally, for the procedure call to *stat* passing no arguments, we have

$$\begin{array}{c}
[], \nabla_{\text{DC}} \vdash \underline{\text{create.select.foreach; drop}} \\
\hline
[], \nabla_{\text{DC}} \vdash \text{stat}
\end{array}$$

□

Remark 5.3. The statistics-gathering task performed by the procedure *stat* of Figure 16 can also be accomplished using *code mobility* combined with local aggregations, instead of a (sequential) series of remote aggregations. In more detail, the head office can maintain a topology of all the branches, and spawn multiple agents that are in charge of different subsets of the localities $\mathbf{1}_1, \dots, \mathbf{1}_n$. These agents then travel from branch to branch according to the topology (potentially stored in tables) of which they are informed, aggregating the sales figures from each branch they arrive at, that does indeed sell shoes of the brand

KLD. The agents can either insert the result obtained from each locality remotely into a designated result table at l_0 , or return to l_0 in the end, when their results are gathered.

6. DISCUSSION AND RELATED WORK

This section provides a discussion of our work around several topics, where we provide additional motivations for our design choices, related to similar works and argue about possible future developments.

Tuple-based coordination languages. Our work studies what tuple-based coordination languages have to offer for distributed database systems. Our main sources of inspiration have been the family of such languages which includes Linda [11], Klaim [7], and SCEL, to mention a few. The suitability of Linda [11] to construct a software framework for distributed database systems was investigated in [21]. Their approach, however, was not concerned with formal language design. Instead, the authors focus on a low-level representation of individual records in tables as tuples in the Linda tuple spaces, and propose the strategy of translating SQL queries into basic Linda operations such as output and input. Correspondingly, they considered some fine-grained issues such as data replication, concurrency control, and fault tolerance.

A detailed discussion with respect to Klaim has been already provided in the Introduction. We would like to mention here an implementation-oriented extension of Klaim, namely X-Klaim [4]. X-Klaim is a full-fledged programming language targeting mobile applications. It provides language primitives for stronger forms of mobility — migrating not only code but also the current execution state. Later in this section, we will envision an implementation of Klaim-DB that builds on the X-Klaim implementation.

The Service Component Ensemble Language (SCEL) [10, 9] can also be seen as an evolution of Klaim, aimed at facilitating the programming of software component ensembles whose coordination patterns can adapt to environmental conditions. A salient feature of SCEL is that components publish their own attributes through interfaces and groups of components can be addressed in basic actions using predicates over these attributes in what the authors call *attribute-based communication* (see also [2]). The publication of table identifiers and schemas by interfaces of tables in Klaim-DB is analogous to the publication of attributes by interfaces of components in SCEL, although in our work we did not exploit such interface to provide richer attribute-based primitives. Further developments of Klaim-DB could indeed incorporate programming abstractions that may allow to operate on sets of databases without referring them by name but by their properties.

Interoperability. A closely related topic that we considered during the design of Klaim-DB is that of interoperability among multiple databases and/or user applications (possibly Klaim-based). In this regard, there are two aspects that we would like to mention. First, concerning homogeneous interoperability among multiple databases, we think that the support of join operations in selection actions paves the way for the general ability to operate on multiple databases by a single action, which is in line with the design philosophy of multi-database systems (e.g., [14]). Second, heterogeneous interoperability between database systems and user applications that are not data-centric can also be realized by bringing back the original Klaim primitives (such as $\text{out}(t)@l$ and $\text{in}(T)@l$) and allowing the co-existence of tables and plain tuples.

Pattern matching and predicates. Our preliminary work [24] considered templates featuring both formal and actual fields and was in line with Klaim templates and the database query language QBE (Query by Example [25]). In the present work templates are restricted to formal fields only, since the combined power of pattern matching with such templates and predicates provides a more powerful mechanism with respect to the bare pattern matching used, for instance, in Linda and Klaim variants (e.g. it is not possible to match tuples with distinct values in the fields). Moreover, *conditional* pattern matching is present in database query languages like SQL (through the `WHERE` clause).

Multiset semantics and temporary tables. Our use of multiset operations for the semantics of select actions has the flavor of the Domain Relational Calculus underlying QBE [25]. Concerning the treatment of the result of such operations, an alternative that we initially considered but did not adopt is the direct placement of the result in a separate table. This table would either be created automatically by the selection operation itself, with an identifier specified in the selection action, or a designated “result table” — one at each locality. However, a problem with this option is that the removal of the automatically created tables would need to be taken care of by the system designer making the specification, using `drop(tid@ℓ)` actions. Similar problems would arise with the maintenance of the designated “result tables” (e.g., the alteration of its schema, the cleaning of old results, etc.). To abstain from these low-level considerations, table variables were finally introduced and binding is used for the selection results.

Formal approaches to query languages. In Klaim-DB, the DB-oriented actions correspond tightly to SQL-like queries. In the literature, formalizations of the relational database model and SQL exist. One piece of work along this direction is [16]. In [16], the syntax and semantics of two different query formalisms, *relational algebra* and *conjunctive queries*, are specified, some existing methods of logical query optimization are dealt with, and the soundness and/or completeness of two existing procedures for inferring integrity constraints are proved. Another recent development, [3], focuses on a certified realization of core SQL operations. The relational algebra is used to define the semantics of certain key SQL operations. Some source-to-source optimizations of SQL queries are proven to be semantically preserving and crude accounts of correspondences in running times are given. Lastly, an efficient implementation of the SQL model considered is performed with the Ynot extension to Coq, and is proved to satisfy its abstract specification. The main differences with respect to our work is that [16] and [3] do not consider *distributed* databases, and are not concerned with the kind of properties our type system deals with.

Security. Security is an important concern, especially in a distributed setting. A *basic step* towards security is that of preventing the exploitation of *unsafe* language operations by a malicious user. An example of this kind with programming languages is unsafe C library functions giving rise to buffer overflow exploits. Providing safe operations needs to be considered along with the core language design. In our case this amounts to catching the insertion of mis-formatted records or the use of undefined variables, to mention a few, in the formal semantics of Klaim-DB and its type system. A *further step* towards security is the provision of abilities for access control and information flow control. For instance, when inserting data into a remote table, it is important to know whether the current locality trusts the remote locality with respect to confidentiality, and whether the remote locality trusts the current locality with respect to integrity. In addition, information should not be directly

or indirectly leaked to or influenced from, unauthorized localities. We expect that access control mechanisms can be devised for Klaim-DB using security type systems or flow logics following the paths taken by [7] and [8]. On the other hand, information flow mechanisms can be devised based on the insights in [22], which enforces locality-based information flow policies in Klaim, and [15], which is concerned with information flow security for SQL-like language primitives.

Implementation. While our focus was on developing a modelling language, it is worth discussing how the Klaim-DB could be turned into a domain-specific programming language following the inspiration of Klaim and its counterpart X-Klaim. As a matter of fact, we envision an implementation of the Klaim-DB language, building on the implementation [4] of X-Klaim. The latter consists of a compiler that translates X-Klaim programs into Java programs that make use of the Klava package [5] for tuple-based operations.

An important part of the implementation would be the Klaim-DB compiler, which would extend the X-Klaim compiler and turn a Klaim-DB specification into a Java program. The Java program would “execute” the specification; hence it would perform operations on *tables* according to the Klaim-DB semantics. The Klava package would need to be extended to support these operations. For instance, classes `Table`, `Schema`, etc., would need to be created to represent tables, schemas, etc., and the methods of these classes would correspond to the table operations. The side conditions of the semantic rules can be easily checked by generating appropriate Java code. Because of the correspondence between the DB-oriented actions in Klaim-DB and SQL queries described in Section 2, we would produce the query results by constructing SQL objects and queries via a Java API for database connectivity such as JDBC. Hence classes supporting the conversion between objects of the classes `Table`, `Schema`, etc., and SQL objects would be included in the extended Klava package. Since a SQL query is not always executed as an atomic operation, the locking mechanism of Java would be needed. To sum up, the Java code to be generated for executing a DB-oriented action would accomplish the following:

- Lock the table objects involved in the action;
- Check the side conditions specified in the semantics;
- Convert table objects involved in the action to SQL objects;
- Perform the query by interfacing with a SQL engine via JDBC;
- Convert the results back to internal objects;
- Unlock relevant objects.

The type checking of Klaim-DB specifications, on the other hand, can be performed directly on the AST (Abstract Syntax Tree) obtained by parsing the Klaim-DB specification in the Klaim-DB compiler.

Transactions. In supporting database operations, the Klaim-DB language is positioned between succinct core calculi such as the relational algebra, and full-fledged query languages such as SQL, to facilitate both the formalization² of its semantics, typing, etc., and its utilization in modeling tasks. Corresponding to its level of abstraction, *high-level atomic operations* are usually used in Klaim-DB to achieve tasks that need to be accomplished with transactions in other languages such as SQL. Nevertheless, Klaim-DB can be extended in the future with explicit transactions, to support use cases such as where a user creates a table, but does not want other users to have access to it before she has inserted all her

²Note that formalizations of SQL exist but often do not cover transactions [13, 17].

desired records into it. In general, an extension to provide full support of transactions needs not only to provide appropriate locking mechanisms (as language primitives), but also to support *rollbacks* to recover from inconsistent states.

7. CONCLUSION

In answer to the challenge of easily and correctly structuring, managing, and utilizing distributed databases, we have presented the design of a coordination language, Klaim-DB. The central notion of data in Klaim-DB is that of localized databases. Primitives for data definition and manipulation are provided at a high abstraction level, guarding against misuses of programming constructs and subtle concurrency errors in specifications. The correctness of data formats is checked in the semantics, and to a great extent statically guaranteed by a type system, which allows efficient type checking. We have modeled in the language a data aggregation task performed across geographically scattered databases, guided by a coordinator.

ACKNOWLEDGEMENT

We would like to thank the anonymous reviewers for their useful comments.

REFERENCES

- [1] Jean-Raymond Abrial. Formal methods: Theory becoming practice. *J. UCS*, 13(5):619–628, 2007.
- [2] Yehia Abd Alrahman, Rocco De Nicola, and Michele Loreti. On the power of attribute-based communication. In Elvira Albert and Ivan Lanese, editors, *Formal Techniques for Distributed Objects, Components, and Systems - 36th IFIP WG 6.1 International Conference, FORTE 2016, Held as Part of the 11th International Federated Conference on Distributed Computing Techniques, DisCoTec 2016, Heraklion, Crete, Greece, June 6-9, 2016, Proceedings*, volume 9688 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2016.
- [3] Véronique Benzaken, Évelyne Contejean, and Stefania Dumbrava. A Coq formalization of the relational data model. In *Programming Languages and Systems*, pages 189–208. Springer, 2014.
- [4] Lorenzo Bettini, Rocco De Nicola, and Michele Loreti. Implementing Mobile and Distributed Applications in X-Klaim. *Scalable Computing: Practice and Experience, Special Issue: Software Agent Mobility*, 7(4):13–35, 2006.
- [5] Lorenzo Bettini, Rocco De Nicola, and Rosario Pugliese. Klava: a Java Package for Distributed and Mobile Applications. *Software - Practice and Experience*, 32(14):1365–1394, 2002.
- [6] E. F. Codd. A relational model of data for large shared data banks. *Commun. ACM*, 13(6):377–387, 1970.
- [7] Rocco De Nicola, Gian Luigi Ferrari, and Rosario Pugliese. KLAIM: A kernel language for agents interaction and mobility. *IEEE Trans. Software Eng.*, 1998.
- [8] Rocco De Nicola, Daniele Gorla, René Rydhof Hansen, Flemming Nielson, Hanne Riis Nielson, Christian W. Probst, and Rosario Pugliese. From flow logic to static type systems for coordination languages. *Sci. Comput. Program.*, 75(6):376–397, 2010.
- [9] Rocco De Nicola, Diego Latella, Alberto Lluch Lafuente, Michele Loreti, Andrea Margheri, Mieke Massink, Andrea Morichetta, Rosario Pugliese, Francesco Tiezzi, and Andrea Vandin. The SCEL language: Design, implementation, verification. In Martin Wirsing, Matthias M. Hölzl, Nora Koch, and Philip Mayer, editors, *Software Engineering for Collective Autonomic Systems - The ASCENS Approach*, volume 8998 of *Lecture Notes in Computer Science*, pages 3–71. Springer, 2015.
- [10] Rocco De Nicola, Michele Loreti, Rosario Pugliese, and Francesco Tiezzi. A formal approach to autonomic systems programming: The SCEL language. *TAAS*, 9(2):7, 2014.
- [11] David Gelernter. Generative communication in linda. *ACM Trans. Program. Lang. Syst.*, 7(1):80–112, 1985.

- [12] David Gelernter and Arthur J. Bernstein. Distributed communication via global buffer. In *ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Ottawa, Canada August 18-20, 1982*, pages 10–18, 1982.
- [13] Martin Gogolla, editor. *An Extended Entity-Relationship Model: Fundamentals and Pragmatics*, chapter Formal semantics of SQL, pages 99–120. Springer Berlin Heidelberg, Berlin, Heidelberg, 1994.
- [14] E. Kuhn and T. Ludwig. Vip-mdbs: A logic multidatabase system. In *Proceedings of the First International Symposium on Databases in Parallel and Distributed Systems, DPDS '88*, pages 190–201. IEEE Computer Society Press.
- [15] Luísa Lourenço and Luís Caires. Information flow analysis for valued-indexed data security compartments. In *Trustworthy Global Computing - 8th International Symposium, TGC 2013*, pages 180–198, 2013.
- [16] Gregory Malecha, Greg Morrisett, Avraham Shinnar, and Ryan Wisnesky. Toward a verified relational database management system. In *ACM Sigplan Notices*, volume 45, pages 237–248. ACM, 2010.
- [17] M. Negri, G. Pelagatti, and L. Sbatella. Formal semantics of sql queries. *ACM Trans. Database Syst.*, 16(3):513–534, September 1991.
- [18] Chris Newcombe. Why amazon chose tla+. In Yamine Ait Ameer and Klaus-Dieter Schewe, editors, *Abstract State Machines, Alloy, B, TLA, VDM, and Z*, volume 8477 of *Lecture Notes in Computer Science*, pages 25–39. 2014.
- [19] Benjamin C. Pierce. *Types and programming languages*. MIT Press, 2002.
- [20] Gordon D. Plotkin. A structural approach to operational semantics. *J. Log. Algebr. Program.*, 60-61:17–139, 2004.
- [21] Madhan M Thirukonda and Ronaldo Menezes. On the use of Linda as a framework for distributed database systems. Technical report, 2002.
- [22] Terkel K. Tolstrup, Flemming Nielson, and René Rydhof Hansen. Locality-based security policies. In *Formal Aspects in Security and Trust, Fourth International Workshop, FAST 2006*, pages 185–201, 2006.
- [23] Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui, and John S. Fitzgerald. Formal methods: Practice and experience. *ACM Comput. Surv.*, 41(4), 2009.
- [24] Xi Wu, Ximeng Li, Alberto Lluch Lafuente, Flemming Nielson, and Hanne Riis Nielson. Klaim-db: A modeling language for distributed database applications. In *Coordination Models and Languages - 17th IFIP WG 6.1 International Conference, COORDINATION 2015*, pages 197–212, 2015.
- [25] Moshé M. Zloof. Query by example. In *Proceedings of the May 19-22, 1975, National Computer Conference and Exposition*, AFIPS '75, pages 431–438, 1975.

APPENDIX A. MULTISSET NOTATION

We use \uplus , \cap and \setminus to represent the union, intersection and subtraction, respectively, of multisets. For a multiset S , and an element s , the application $M(S, s)$ of the multiplicity function M gives the number of repetitions of s in S . Note that for $s \notin S$, $M(S, s) = 0$. Then our notions of union, intersection and subtraction are such that

$$\begin{aligned} M(S_1 \uplus S_2, s) &= M(S_1, s) + M(S_2, s) \\ M(S_1 \cap S_2, s) &= \min(M(S_1, s), M(S_2, s)) \\ M(S_1 \setminus S_2, s) &= \max(M(S_1, s) - M(S_2, s), 0) \end{aligned}$$

APPENDIX B. PROOFS

We first restate Lemma 3.14 and give its proof.

Lemma 3.14. *If $\text{no_rep}(\text{Lid}(N) \uplus \text{Lid}(N_{\text{env}}))$ and $N_{\text{env}} \vdash N \rightarrow N'$, then it holds that $\text{no_rep}(\text{Lid}(N') \uplus \text{Lid}(N_{\text{env}}))$.*

Proof. Assume that $no_rep(Lid(N) \uplus Lid(N_{env}))$ holds. We proceed with an induction on the derivation of $N_{env} \vdash N \rightarrow N'$ to establish $no_rep(Lid(N') \uplus Lid(N_{env}))$.

Case (INS), Case (DEL), Case (UPD), and Case (AGR):

we have $Lid(N) = Lid(N') = \{(l_2, I.tid)\}$ if $N' \neq ERR$, and $Lid(N') = \emptyset$ if $N' = ERR$. Hence $no_rep(Lid(N') \uplus Lid(N_{env}))$ holds.

Case (SEL): We have $Lid(N) = Lid(N') = \{(l_1, I_1.tid), \dots, (l_k, I_k.tid)\}$ if $N' \neq ERR$, and $Lid(N') = \emptyset$ if $N' = ERR$. Hence $no_rep(Lid(N') \uplus Lid(N_{env}))$ holds.

Case (DRP), Case (EVL), Case (FOR^{tt}), Case (FOR^{ff}), and Case (CALL):

We have $Lid(N') = \emptyset$. By $no_rep(Lid(N) \uplus Lid(N_{env}))$, we have $no_rep(Lid(N_{env}))$. Hence $no_rep(Lid(N') \uplus Lid(N_{env}))$ holds.

Case (CRT): If $(l_2, tid) \notin Lid(N_{env})$, then since $Lid(N')$ is the multi-set whose only element is (l_2, tid) , with multiplicity 1, we have $no_rep(Lid(N') \uplus Lid(N_{env}))$. If $(l_2, tid) \in Lid(N_{env})$, then $Lid(N') = \emptyset$. Hence $no_rep(Lid(N') \uplus Lid(N_{env}))$ also holds.

Case (SEQ^{tt}): We have $Lid(l :: P'_1; P_2 || N) = Lid(l :: P_1; P_2 || N)$.

Case (SEQ^{ff}): Analogous to Case (SEQ^{tt}).

Case (PAR): The net N is $N_1 || N_2$ and N' is $N'_1 || N_2$. We have $no_rep(Lid(N_1 || N_2) \uplus Lid(nil_N))$, which implies $no_rep(Lid(N_1) \uplus Lid(N_2))$. By the induction hypothesis for $N_2 \vdash N_1 \rightarrow N'_1$, we have $no_rep(Lid(N'_1) \uplus Lid(N_2))$. Therefore we also have $no_rep(Lid(N'_1 || N_2) \uplus Lid(nil_N))$.

Case (RES): Trivial by using the induction hypothesis.

Case (EQUIV): If neither of the equivalences, $N_1 \equiv N_2$ and $N_3 \equiv N_4$, is concerned with α -renaming, then it is obvious that $Lid(N_1) = Lid(N_2)$ and $Lid(N_3) = Lid(N_4)$. Using the induction hypothesis it is straightforward to derive $no_rep(Lid(N_4) \uplus Lid(N_{env}))$.

Suppose α -renaming is involved in $N_1 \equiv N_2$. By the global assumption that bound localities should be distinctly named, no locality name can be bound in both N_2 and N_{env} . By $no_rep(Lid(N_1) \uplus Lid(N_{env}))$, we also have $no_rep(Lid(N_2) \uplus Lid(N_{env}))$. By the induction hypothesis we have $no_rep(Lid(N_3) \uplus Lid(N_{env}))$. By similar reasoning with $N_3 \equiv N_4$ we can deduce $no_rep(Lid(N_4) \uplus Lid(N_{env}))$.

This completes the proof. \square

We proceed with proving Theorem 4.2.

Lemma B.1. *If $N_1 \equiv_\alpha N_2$, then $\Gamma, \nabla \vdash N_1 \Leftrightarrow \Gamma, \nabla \vdash N_2$.*

Lemma B.2. *If $N_1 \equiv N_2$, then $\Gamma, \nabla \vdash N_1 \Leftrightarrow \Gamma, \nabla \vdash N_2$.*

Lemma B.3. *Suppose at most one binding is present in Γ for each variable in \mathbf{D}_Γ , and Γ_1 is a re-ordering of the bindings of Γ . The the following statements hold.*

- (1) *If $\Gamma \vdash e \triangleright \tau$, then we have $\Gamma_1 \vdash e \triangleright \tau$.*
- (2) *If $\Gamma \vdash \psi \triangleright Bool$, then we have $\Gamma_1 \vdash \psi \triangleright Bool$.*
- (3) *If $\Gamma \vdash t \triangleright \tau$, then we have $\Gamma_1 \vdash t \triangleright \tau$.*
- (4) *If $\Gamma, \nabla \vdash TB \triangleright \tau$, then we have $\Gamma_1, \nabla \vdash TB \triangleright \tau$.*
- (5) *If $\Gamma, \nabla \vdash a \triangleright \Gamma'$, and $bv(a) \cap \mathbf{D}_\Gamma = \emptyset$, then $\Gamma_1, \nabla \vdash a \triangleright \Gamma'$, and at most one binding is present for each variable in Γ' , and if $\Gamma, \nabla \vdash P$, and $bv(P) \cap \mathbf{D}_\Gamma = \emptyset$, then $\Gamma_1, \nabla \vdash P$.*

Sketch of Proof. The proofs of (1) – (4) are straightforward. For the proof of (5), we proceed with an induction on the structure of actions and processes. We only present a few representative cases.

Case $\text{select}(\widetilde{TB}, T, \psi, t, !tbv)$: Since $\Gamma, \nabla \vdash \text{select}(\widetilde{TB}, T, \psi, t, !tbv) \triangleright \Gamma'$ we have $\Gamma, \nabla \vdash TB_1 \triangleright \tau_1, \dots, \Gamma, \nabla \vdash TB_n \triangleright \tau_n$ for some τ_1, \dots, τ_n such that $\tau_1 \times \dots \times \tau_n \vdash T \triangleright \Gamma''$, for some Γ'' such that $\Gamma, \Gamma' \vdash \psi \triangleright \text{Bool}$ and $\Gamma, \Gamma' \vdash t \triangleright \tau'$ for some τ' such that $\Gamma' = [tbv : \tau']$. Hence at most one binding exists in Γ' for the only variable tbv in $\mathbf{D}_{\Gamma'}$.

Using (4) we have $\Gamma_1, \nabla \vdash TB_1 \triangleright \tau_1, \dots, \Gamma_1, \nabla \vdash TB_n \triangleright \tau_n$. By our assumption that templates are linear, at most one binding exists in Γ'' for each variable in $\mathbf{D}_{\Gamma''}$. By the conditions of (5), we have $bv(T) \cap \mathbf{D}_{\Gamma} = \emptyset$. Hence it holds that $\mathbf{D}_{\Gamma''} \cap \mathbf{D}_{\Gamma} = \emptyset$. Thus at most one binding exists in Γ, Γ'' for each variable in its domain, and the same holds for Γ_1, Γ'' , which is also a reordering of the bindings of Γ, Γ'' . By (2) and (3) we have $\Gamma_1, \Gamma'' \vdash \psi \triangleright \text{Bool}$ and $\Gamma_1, \Gamma'' \vdash t \triangleright \tau'$. We can now establish $\Gamma_1, \nabla \vdash \text{select}(\widetilde{TB}, T, \psi, t, !tbv) \triangleright \Gamma'$.

Case $\text{eval}(P')@l$: By $\Gamma, \nabla \vdash \text{eval}(P')@l \triangleright []$ we have $\Gamma, \nabla \vdash P'$ and $\Gamma \vdash l \triangleright \text{Loc}$. We have $bv(\text{eval}(P')@l) = bv(P')$. Hence $bv(P') \cap \mathbf{D}_{\Gamma} = \emptyset$. By the induction hypothesis we have $\Gamma_1, \nabla \vdash P'$. By (1) we have $\Gamma_1 \vdash l \triangleright \text{Loc}$. We can then establish $\Gamma_1, \nabla \vdash \text{eval}(P')@l$.

Case $a'.P'$: By $\Gamma, \nabla \vdash a'.P'$ we have $\Gamma, \nabla \vdash a' \triangleright \Gamma''$ for some Γ'' such that $(\Gamma, \Gamma''), \nabla \vdash P'$ holds. By $bv(a'.P') \cap \mathbf{D}_{\Gamma} = \emptyset$ we have

$$bv(a') \cap \mathbf{D}_{\Gamma} = \emptyset \tag{B.1}$$

$$bv(P') \cap \mathbf{D}_{\Gamma} = \emptyset \tag{B.2}$$

Hence by the induction hypothesis we have $\Gamma_1, \nabla \vdash a' \triangleright \Gamma''$, and at most one binding exists in Γ'' for each variable in its domain. Since $bv(a') \cap \mathbf{D}_{\Gamma} = \emptyset$, it is obvious that $\mathbf{D}_{\Gamma''} \cap \mathbf{D}_{\Gamma} = \emptyset$, and $\mathbf{D}_{\Gamma''} \cap \mathbf{D}_{\Gamma_1} = \emptyset$. It is not difficult to see that at most one binding exists in Γ_1, Γ'' for each variable in $\mathbf{D}_{\Gamma_1, \Gamma''}$. It is also obvious that Γ_1, Γ'' is a reordering of the bindings of Γ, Γ'' . Since variables are renamed apart in processes, we have $bv(P') \cap bv(a') = \emptyset$. Hence $bv(P') \cap \mathbf{D}_{\Gamma''} = \emptyset$. By (B.1) we have $bv(P') \cap \mathbf{D}_{\Gamma, \Gamma''} = \emptyset$; hence $bv(P') \cap \mathbf{D}_{\Gamma_1, \Gamma''} = \emptyset$. By the induction hypothesis we have $(\Gamma_1, \Gamma''), \nabla \vdash P'$. We can now establish $\Gamma_1, \nabla \vdash a'.P'$. \square

Lemma B.4. *The following statements hold.*

- (1) *If $\Gamma \vdash e \triangleright \tau$, then $fv(e) \subseteq \mathbf{D}_{\Gamma}$.*
- (2) *If $\Gamma, \nabla \vdash TB \triangleright \tau$, then $fv(TB) \subseteq \mathbf{D}_{\Gamma}$.*
- (3) *If $\Gamma \vdash \psi \triangleright \text{Bool}$, then $fv(\psi) \subseteq \mathbf{D}_{\Gamma}$.*
- (4) *If $\Gamma \vdash t \triangleright \tau$, then $fv(t) \subseteq \mathbf{D}_{\Gamma}$.*
- (5) *If $\Gamma \vdash a \triangleright \Gamma'$ then $fv(a) \subseteq \mathbf{D}_{\Gamma}$ and $bv(a) \supseteq \mathbf{D}_{\Gamma'}$, and if $\Gamma, \nabla \vdash P$, then $fv(P) \subseteq \mathbf{D}_{\Gamma}$.*

The proof of Lemma B.4 is straightforward.

Lemma B.5. *The type environment can be shrunk as long as its domain still contains all the free variables of the objects to be typed. Formally, supposing $\Gamma' \subseteq \Gamma$, and at most one binding is present in Γ for each variable in \mathbf{D}_{Γ} , we have the following results:*

- (1) *If $\Gamma \vdash e \triangleright \tau$ and $fv(e) \subseteq \mathbf{D}_{\Gamma'}$, then $\Gamma' \vdash e \triangleright \tau$.*
- (2) *If $\Gamma, \nabla \vdash TB \triangleright \tau$, and $fv(TB) \subseteq \mathbf{D}_{\Gamma'}$, then $\Gamma', \nabla \vdash TB \triangleright \tau$.*
- (3) *If $\Gamma \vdash \psi \triangleright \text{Bool}$ and $fv(\psi) \subseteq \mathbf{D}_{\Gamma'}$, then $\Gamma' \vdash \psi \triangleright \text{Bool}$.*
- (4) *If $\Gamma \vdash t \triangleright \tau$ and $fv(t) \subseteq \mathbf{D}_{\Gamma'}$, then $\Gamma' \vdash t \triangleright \tau$.*

Proof.

- (1) The proof is by induction on the structure of e , we only present the representative cases.

Case x : By $\Gamma \vdash x \triangleright \tau$ we have $\Gamma(x) = \tau$ and $\tau \neq Loc$. Since $\Gamma' \subseteq \Gamma$, only one binding exists for x in Γ , and $fv(x) \subseteq \mathbf{D}_{\Gamma'}$, it holds that $\Gamma'(x) = \tau$. Hence $\Gamma' \vdash x \triangleright \tau$ can be established.

Case $\{e_1, \dots, e_n\}$: By $\Gamma \vdash \{e_1, \dots, e_n\} \triangleright \tau$, there exists some τ_0 such that $\tau = \tau_0 \text{ mset}$, and $\Gamma \vdash e_1 \triangleright \tau_0, \dots, \Gamma \vdash e_n \triangleright \tau_0$. By $fv(\{e_1, \dots, e_n\}) \subseteq \mathbf{D}_{\Gamma'}$ we have $fv(e_1) \subseteq \mathbf{D}_{\Gamma'}, \dots, fv(e_n) \subseteq \mathbf{D}_{\Gamma'}$. By the induction hypothesis we have $\Gamma' \vdash e_1 \triangleright \tau_0, \dots, \Gamma' \vdash e_n \triangleright \tau_0$. We thus have $\Gamma' \vdash \{e_1, \dots, e_n\} \triangleright \tau$.

- (2) The proof is by a case analysis on TB .

Case $tid@l$: By $\Gamma, \nabla \vdash tid@l \triangleright \tau$ we have $\Gamma \vdash l \triangleright Loc$ and $\tau = \nabla(tid)$. By $fv(TB) \subseteq \mathbf{D}_{\Gamma'}$ we have $fv(l) \subseteq \mathbf{D}_{\Gamma'}$. Using (1) we have $\Gamma' \vdash l \triangleright Loc$. Hence $\Gamma', \nabla \vdash tid@l \triangleright \tau$ can be established.

Case tbv : By $\Gamma, \nabla \vdash tbv \triangleright \tau$ we have $tbv \in \mathbf{D}_{\Gamma}$ and $\tau = \Gamma(tbv)$. Since $\Gamma' \subseteq \Gamma$, only one binding exists for tbv in Γ , and $fv(tbv) \subseteq \mathbf{D}_{\Gamma'}$, we have $tbv \in \mathbf{D}_{\Gamma'}$ and $\tau = \Gamma'(tbv)$. Hence $\Gamma', \nabla \vdash tbv \triangleright \tau$ can be established.

Case (I, R) : By $\Gamma, \nabla \vdash tbv \triangleright \tau$ we have $\forall t \in R : \Gamma \vdash t \triangleright I.sk$ and $I.tid \neq \perp \Rightarrow \nabla(I.tid) = I.sk$. For each tuple t in the data set R , we have $fv(t) = \emptyset$ and it is obvious that the typing of t is irrelevant to the particular typing environment in use. Hence we have $\forall t \in R : \Gamma' \vdash t \triangleright I.sk$. It can then be established that $\Gamma', \nabla \vdash tbv \triangleright \tau$.

- (3) The proof is straightforward by induction on the structure of ψ , using (1).

- (4) The proof is trivial using (1). □

Lemma B.6. *The following statements hold.*

- (1) *If $\square \vdash e \triangleright \tau$, then $\mathcal{E}[e] \Vdash \tau$, and vice versa.*
- (2) *If $\square \vdash t \triangleright \tau$, then $\mathcal{E}[t] \Vdash \tau$, and vice versa.*
- (3) *If $\square \vdash \psi \triangleright Bool$, then $\mathcal{E}[\psi] \neq err$, and vice versa.*

Proof. The proofs are by structural induction. We present only representative cases.

- (1a) We prove if $\square \vdash e \triangleright \tau$, then $\mathcal{E}[e] \Vdash \tau$, with an induction on the structure of e .

Case num : We have $\mathcal{E}[num] = num$. By $\square \vdash num \triangleright \tau$ we have $\tau = Int$. It indeed holds that $num \Vdash Int$.

Case x : The typing cannot be performed with the empty environment \emptyset ; hence the statement vacuously holds.

Case $e_1 \cdot e_2$: By $\square \vdash e_1 \cdot e_2 \triangleright \tau$ we have $\square \vdash e_1 \triangleright \tau$, $\square \vdash e_2 \triangleright \tau$ and $\tau = String$. By the induction hypothesis we have $\mathcal{E}[e_1] \Vdash String$ and $\mathcal{E}[e_2] \Vdash String$. By the definition of \Vdash , $e_1 = str_1$ and $e_2 = str_2$ for strings str_1 and str_2 . Hence $\mathcal{E}[e_1 \cdot e_2] = (str_1 \cdot str_2)$ is a string. We have $\mathcal{E}[e_1 \cdot e_2] \Vdash String$.

- (1b) We prove if $\mathcal{E}[e] \Vdash \tau$, then $\square \vdash e \triangleright \tau$, with an induction on the structure of e .

Case num : We have $\mathcal{E}[num] \Vdash Int$; hence $\tau = Int$. It indeed holds that $\square \vdash num \triangleright Int$.

Case x : We do *not* have $\mathcal{E}[[x]] \Vdash \tau$ for any τ ; hence the statement vacuously holds.

Case $e_1 \cdot e_2$: Since $\mathcal{E}[[e_1 \cdot e_2]] \Vdash \tau$, we have $\mathcal{E}[[e_1 \cdot e_2]]$ yields $\mathcal{E}[[e_1]] \cdot \mathcal{E}[[e_2]]$ rather than *err*, and $\mathcal{E}[[e_1]] = str_1$ and $\mathcal{E}[[e_2]] = str_2$ for some str_1 and str_2 . Hence $\mathcal{E}[[e_1]] \Vdash String$ and $\mathcal{E}[[e_2]] \Vdash String$. By the induction hypothesis we have $\square \vdash e_1 \triangleright String$ and $\square \vdash e_2 \triangleright String$. Hence it can be established that $\square \vdash e_1 \cdot e_2 \triangleright String$.

(2a) We prove if $\square \vdash t \triangleright \tau$, then $\mathcal{E}[[t]] \Vdash \tau$.

The typing judgment for t must be $\square \vdash e_1, \dots, e_n \triangleright \tau$. We thus have $\tau = \tau_1 \times \dots \times \tau_n$ for some τ_1, τ_2, \dots , and τ_n such that $\square \vdash e_1 \triangleright \tau_1, \dots, \square \vdash e_n \triangleright \tau_n$. By (1a) we have $\forall j \in \{1, \dots, n\} : \mathcal{E}[[e_j]] \Vdash \tau_j$. Hence $\forall j \in \{1, \dots, n\} : \mathcal{E}[[e_j]] \neq err$ and by the definitions of $\mathcal{E}[\cdot]$ and $\cdot \Vdash \cdot$, it holds that $\mathcal{E}[[e_1, \dots, e_n]] = (\mathcal{E}[[e_1]], \dots, \mathcal{E}[[e_n]]) \Vdash \tau_1 \times \dots \times \tau_n = \tau$.

(2b) We prove if $\mathcal{E}[[t]] \Vdash \tau$, then $\square \vdash t \triangleright \tau$.

Suppose $t = e_1, \dots, e_n$. Since we do *not* have $err \Vdash \tau$, it holds that $\mathcal{E}[[e_1, \dots, e_n]] \neq err$, which means $\mathcal{E}[[e_1, \dots, e_n]] = (\mathcal{E}[[e_1]], \dots, \mathcal{E}[[e_n]])$ and $\forall j \in \{1, \dots, n\} : \mathcal{E}[[e_j]] \neq err$. By $(\mathcal{E}[[e_1]], \dots, \mathcal{E}[[e_n]]) \Vdash \tau$ we have $\tau = \tau_1 \times \dots \times \tau_n$ for some τ_1, τ_2, \dots , and τ_n such that $\forall j \in \{1, \dots, n\} : \mathcal{E}[[e_j]] \Vdash \tau_j$. By (1b) we have $\forall j \in \{1, \dots, n\} : \square \vdash e_j \triangleright \tau_j$. Hence it holds that $\square \vdash e_1, \dots, e_n \triangleright \tau$.

(3a),(3b) The proof ideas are analogous. □

Pedantically, the condition that “at most one binding exists in a typing environment for each variable in its domain” would be needed in the statements of a few further lemmas whose proofs make use of Lemma B.3 and/or Lemma B.5. However, to avoid the tediousness we dispense hereafter with explicating such assumptions on (all of) our typing environments. Since bound variables are renamed apart, the extensions of typing environments made in the type system all preserve this condition.

Lemma B.7 (Substitution for Tables, Expressions, Predicates and Tuples). *The following statements hold.*

- (1) If $(\Gamma, tbv : \tau_0), \nabla \vdash TB \triangleright \tau'$, and $\square, \nabla \vdash (I, R) : \tau_0$, then $\Gamma, \nabla \vdash TB[(I, R)/tbv] \triangleright \tau'$ holds.
- (2) If $(\Gamma, u : Loc), \nabla \vdash TB \triangleright \tau'$, and l is a locality, then $\Gamma, \nabla \vdash TB[l/u] \triangleright \tau'$ holds.
- (3) If $\Gamma, x : \tau_0 \vdash e \triangleright \tau$ and val is such that $\square \vdash val : \tau_0$, then $\Gamma \vdash e[val/x] \triangleright \tau$ holds.
- (4) If $\Gamma, u : Loc \vdash e \triangleright \tau$, and l is a locality, then $\Gamma \vdash e[l/u] \triangleright \tau$ holds.
- (5) If $\Gamma, x : \tau_0 \vdash \psi \triangleright Bool$, and $\square \vdash val : \tau_0$, then $\Gamma \vdash \psi[val/x] \triangleright Bool$ holds.
- (6) If $\Gamma, u : Loc \vdash \psi \triangleright Bool$, and l is a locality, then $\Gamma \vdash \psi[l/u] \triangleright Bool$ holds.
- (7) If $\Gamma, x : \tau_0 \vdash t \triangleright \tau$, and $\square \vdash val : \tau_0$, then $\Gamma \vdash t[val/x] \triangleright \tau$ holds.
- (8) If $\Gamma, u : Loc \vdash t \triangleright \tau$, and l is a locality, then $\Gamma \vdash t[l/u] \triangleright \tau$ holds.

Proof.

(1) We proceed by a case analysis on TB .

Case $tid@l$: By $(\Gamma, tbv : \tau_0), \nabla \vdash tid@l \triangleright \tau'$ we have $\Gamma, tbv : \tau_0 \vdash l \triangleright Loc$, $tid \in \mathbf{D}_\nabla$ and $\tau' = \nabla(tid)$. By Lemma B.4, we have $fv(l) \subseteq \mathbf{D}_{\Gamma, tbv : \tau_0}$, which implies $fv(l) \subseteq \mathbf{D}_\Gamma$. Hence by Lemma B.5, we have $\Gamma \vdash l \triangleright Loc$. We can now establish $\Gamma, \nabla \vdash (tid@l)[(I, R)/tbv] \triangleright \tau'$ since $(tid@l)[(I, R)/tbv] = tid@l$.

Case tbv' : Suppose $tbv' \neq tbv$. We have $tbv'[(I, R)/tbv] = tbv'$. By $(\Gamma, tbv : \tau_0), \nabla \vdash tbv' \triangleright \tau'$ we have $tbv' \in \mathbf{D}_\Gamma$ and $\tau' = \Gamma(tbv')$. By Lemma B.5 we have $\Gamma, \nabla \vdash tbv' \triangleright \tau'$.

Suppose $tbv' = tbv$. By $\Gamma, tbv : \tau_0 \vdash tbv \triangleright \tau'$ we have $\tau' = \tau_0$. It is obvious that well-typedness is preserved under weakening of the typing environment. Hence by $\square, \nabla \vdash (I, R) \triangleright \tau_0$, we have $\Gamma, \nabla \vdash (I, R) \triangleright \tau_0$, or $\Gamma, \nabla \vdash (I, R) \triangleright \tau'$.

Case (I', R') : By $\Gamma, tbv : \tau_0 \vdash (I', R') \triangleright \tau'$ we have $\forall t \in R' : \Gamma, tbv : \tau_0 \vdash t \triangleright I'.sk$, $I'.tid \neq \perp \Rightarrow \nabla(I'.tid) = I'.sk$, and $I'.sk = \tau'$. Since for each t in R' , it holds that $fv(t) = \emptyset$, we have $\forall t \in R' : \Gamma \vdash t \triangleright I'.sk$ by Lemma B.5. We can thus establish $\Gamma, \nabla \vdash (I', R')[I, R/tbv] \triangleright \tau'$.

(2) We proceed by a case analysis on TB .

Case $tid@l$: Suppose l is some constant locality l' . We have $(tid@l')[l/u] = tid@l'$. By Lemma B.5 and $fv(tid@l') = \emptyset$, we can easily derive $\Gamma, \nabla \vdash tid@l' \triangleright \tau'$.

Suppose l is some locality variable u' such that $u' \neq u$. By $(\Gamma, u : Loc), \nabla \vdash tid@u' \triangleright \tau'$ we have $\Gamma, u : Loc \vdash u' \triangleright Loc$, $tid \in \mathbf{D}_\nabla$ and $\tau' = \nabla(tid)$. It is not difficult to see that $u' \in \mathbf{D}_\Gamma$ using Lemma B.4. By Lemma B.5 we have $\Gamma \vdash u' \triangleright Loc$. We can now establish $\Gamma, \nabla \vdash (tid@u')[l/u] \triangleright \tau'$ since $(tid@u')[l/u] = tid@u'$.

Suppose l is u . By $(\Gamma, u : Loc), \nabla \vdash tid@u \triangleright \tau'$ we have $tid \in \mathbf{D}_\nabla$ and $\tau' = \nabla(tid)$. For the locality l we have $\Gamma \vdash l \triangleright Loc$. Hence we can also establish $\Gamma, \nabla \vdash tid@l \triangleright \tau'$.

Case tbv : Analogous to the corresponding case tbv' under (1) where tbv' is not changed by the substitution.

Case (I, R) : Analogous to the corresponding case (I', R') under (1).

(3) The proof is by induction on the structure of e . We omit the trivial cases.

Case x' : Suppose $x' \neq x$. By $\Gamma, x : u_0 \vdash x' \triangleright \tau$ we have $x' \in \mathbf{D}_\Gamma$. By Lemma B.5 we have $\Gamma \vdash x' \triangleright \tau$.

Suppose $x' = x$. By $\Gamma, x : \tau_0 \vdash x \triangleright \tau$ we have $\tau_0 = \tau$ and $\tau \neq Loc$. By $\square \vdash val \triangleright \tau_0$ we have $\Gamma \vdash val \triangleright \tau_0$ (val is some num , some str or some tid). Hence it holds that $\Gamma \vdash x[val/x] \triangleright \tau$.

Case $e_1 \cdot e_2$: By $\Gamma, x : \tau_0 \vdash e_1 \cdot e_2 \triangleright \tau$ we have $\Gamma, x : \tau_0 \vdash e_1 \triangleright \tau$, $\Gamma, x : \tau_0 \vdash e_2 \triangleright \tau$, and $\tau = String$. By the induction hypothesis we have $\Gamma \vdash e_1[val/x] \triangleright String$ and $\Gamma \vdash e_2[val/x] \triangleright String$. We have $e_1[val/x] \cdot e_2[val/x] = (e_1 \cdot e_2)[val/x]$; hence it can be established that $\Gamma \vdash (e_1 \cdot e_2)[val/x] \triangleright \tau$.

Case $e_1 aop e_2$: Analogous to **Case $e_1 \cdot e_2$** .

Case $\{e_1, \dots, e_n\}$: Analogous to **Case $e_1 \cdot e_2$** .

(4) The proof is by induction on the structure of e . The main differences compared with the proof of (3) lie with the treatment of data variables and locality variables. Of the non-inductive cases, we only exhibit the one for locality variables.

Case u' : Suppose $u' \neq u$. The reasoning is analogous to the sub-case $x' \neq x$ under **Case x'** of (3).

Suppose $u' = u$. By $\Gamma, u : Loc \vdash u \triangleright \tau$ we have $\tau = Loc$. Since l is a locality, it can be established that $\Gamma \vdash u[l/u] \triangleright \tau$.

The inductive cases are analogous to those of (3).

(5) The proof is by induction on the structure of ψ . We only present the representative cases below.

Case $e_1 \text{ cop } e_2$: Because of $\Gamma, x : \tau_0 \vdash e_1 \text{ cop } e_2 \triangleright \text{Bool}$ we have $\Gamma, x : \tau_0 \vdash e_1 \triangleright \tau_d$ and $\Gamma, x : \tau_0 \vdash e_2 \triangleright \tau_d$, for some $\tau_d \in \{\text{Int}, \text{String}, \text{Id}, \text{Loc}\}$. By (3), we get $\Gamma \vdash e_1[\text{val}/x] \triangleright \tau_d$ and $\Gamma \vdash e_2[\text{val}/x] \triangleright \tau_d$. Since $(e_1 \text{ cop } e_2)[\text{val}/x] = (e_1[\text{val}/x] \text{ cop } e_2[\text{val}/x])$ holds, we have $\Gamma \vdash (e_1 \text{ cop } e_2)[\text{val}/x] \triangleright \text{Bool}$.

Case $\psi_1 \wedge \psi_2$: Because of $\Gamma, x : \tau_0 \vdash \psi_1 \wedge \psi_2 \triangleright \text{Bool}$ we have $\Gamma, x : \tau_0 \vdash \psi_1 \triangleright \text{Bool}$ and $\Gamma, x : \tau_0 \vdash \psi_2 \triangleright \text{Bool}$. By the induction hypothesis we have $\Gamma \vdash \psi_1[\text{val}/x] \triangleright \text{Bool}$ and $\Gamma \vdash \psi_2[\text{val}/x] \triangleright \text{Bool}$. Since $(\psi_1 \wedge \psi_2)[\text{val}/x] = (\psi_1[\text{val}/x] \wedge \psi_2[\text{val}/x])$ holds, it can be established that $\Gamma \vdash (\psi_1 \wedge \psi_2)[\text{val}/x] \triangleright \text{Bool}$.

(6) Analogous to (5).

(7) The proof is straightforward using (3).

(8) The proof is straightforward using (4). \square

Lemma B.8 (Substitution for Actions and Processes). *The following statements hold.*

- (1) *If $(\Gamma, x : \tau_0), \nabla \vdash a \triangleright \Gamma'$, $bv(a) \cap \mathbf{D}_{\Gamma, x : \tau_0} = \emptyset$, and val is such that $\square \vdash \text{val} : \tau_0$, then $\Gamma, \nabla \vdash a[\text{val}/x] \triangleright \Gamma'$ holds;
if $(\Gamma, x : \tau_0), \nabla \vdash P$, $bv(P) \cap \mathbf{D}_{\Gamma, x : \tau_0} = \emptyset$, and val is such that $\square \vdash \text{val} : \tau_0$, then $\Gamma, \nabla \vdash P[\text{val}/x]$ holds.*
- (2) *If $(\Gamma, tbv : \tau_0), \nabla \vdash a \triangleright \Gamma'$, $bv(a) \cap \mathbf{D}_{\Gamma, tbv : \tau_0} = \emptyset$, and $\square, \nabla \vdash (I, R) : \tau_0$, then $\Gamma, \nabla \vdash a[(I, R)/tbv] \triangleright \Gamma'$ holds;
if $(\Gamma, tbv : \tau_0), \nabla \vdash P$, $bv(P) \cap \mathbf{D}_{\Gamma, tbv : \tau_0} = \emptyset$, and $\square, \nabla \vdash (I, R) : \tau_0$, then $\Gamma, \nabla \vdash P[(I, R)/tbv]$ holds.*
- (3) *If $(\Gamma, u : \text{Loc}), \nabla \vdash a \triangleright \Gamma'$, $bv(a) \cap \mathbf{D}_{\Gamma, u : \text{Loc}} = \emptyset$, and l is a locality, then $\Gamma, \nabla \vdash a[l/u] \triangleright \Gamma'$ holds;
if $(\Gamma, u : \text{Loc}), \nabla \vdash P$, $bv(P) \cap \mathbf{D}_{\Gamma, u : \text{Loc}} = \emptyset$, and l is a locality, then $\Gamma, \nabla \vdash P[l/u]$ holds.*

Proof.

(1) We proceed with an induction on the structure of actions and processes.

Case $\text{insert}(t, \text{tid}@\ell)$: Due to $(\Gamma, x : \tau_0), \nabla \vdash \text{insert}(t, \text{tid}@\ell) \triangleright \Gamma'$ we obtain $\Gamma, x : \tau_0 \vdash t \triangleright \nabla(\text{tid})$, $\Gamma, x : \tau_0 \vdash \ell \triangleright \text{Loc}$, and $\Gamma' = \square$. By Lemma B.7, we get $\Gamma \vdash t[\text{val}/x] \triangleright \nabla(\text{tid})$. By Lemma B.4 and Lemma B.5 we have $\Gamma \vdash \ell \triangleright \text{Loc}$. Since $\text{insert}(t, \text{tid}@\ell)[\text{val}/x]$ gives $\text{insert}(t[\text{val}/x], \text{tid}@\ell)$, we can indeed establish

$$\Gamma, \nabla \vdash \text{insert}(t, \text{tid}@\ell)[\text{val}/x] \triangleright \Gamma'.$$

Case $\text{delete}(\text{tid}@\ell, T, \psi)$: Because of $(\Gamma, x : \tau_0), \nabla \vdash \text{delete}(\text{tid}@\ell, T, \psi) \triangleright \Gamma'$ there exists some typing environment Γ'' such that $\nabla(\text{tid}) \vdash T \triangleright \Gamma''$, $(\Gamma, x : \tau_0, \Gamma''), \nabla \vdash \psi \triangleright \text{Bool}$, $\Gamma, x : \tau_0 \vdash \ell \triangleright \text{Loc}$, and $\Gamma' = \square$. Since $bv(\text{delete}(\text{tid}@\ell, T, \psi)) \cap \mathbf{D}_{\Gamma, x : \tau_0} = \emptyset$, we have $\mathbf{D}_{\Gamma''} \cap \mathbf{D}_{\Gamma, x : \tau_0} = \emptyset$. By Lemma B.3, we have $\Gamma, \Gamma'', x : \tau_0 \vdash \psi \triangleright \text{Bool}$. By Lemma B.7, we have $\Gamma, \Gamma'' \vdash \psi[\text{val}/x] \triangleright \text{Bool}$. By Lemma B.4 and Lemma B.5 we have $\Gamma \vdash \ell \triangleright \text{Loc}$. Combining this with the typing for T we can establish

$$\Gamma, \nabla \vdash \text{delete}(\text{tid}@\ell, T, \psi)[\text{val}/x] \triangleright \Gamma',$$

using the fact that the substitution gives $\text{delete}(\text{tid}@\ell, T, \psi[\text{val}/x])$.

Case $\text{select}(\widetilde{TB}, T, \psi, t, !tbv)$: By $(\Gamma, x : \tau_0), \nabla \vdash \text{select}(\widetilde{TB}, T, \psi, t, !tbv) \triangleright \Gamma'$ we have $(\Gamma, x : \tau_0), \nabla \vdash TB_j \triangleright \tau_j$ for all $j \in \{1, \dots, n\}$, $\text{flatten}_s(\tau_1 \times \dots \times \tau_n) \vdash T \triangleright \Gamma''$ for some typing environment Γ'' such that $(\Gamma, x : \tau_0), \Gamma'' \vdash \psi \triangleright \text{Bool}$, and $(\Gamma, x : \tau_0), \Gamma'' \vdash t \triangleright \tau'$, for some τ' such that $\Gamma' = [tbv : \tau']$. For each $j \in \{1, \dots, n\}$ we have $TB_j[val/x] = TB_j$. By Lemma B.4, we have $fv(TB_j) \subseteq \mathbf{D}_\Gamma \cup \{x\}$ for each j . Since it is obvious that the only kind of variables each TB_j can contain are table variables, we have $fv(TB_j) \subseteq \mathbf{D}_\Gamma$. By Lemma B.5, we have

$$\forall j \in \{1, \dots, n\} : \Gamma, \nabla \vdash TB_j \triangleright \tau_j \quad (\text{B.3})$$

By Lemma B.3 and Lemma B.7 (analogously to the previous case), we have

$$\Gamma, \Gamma'' \vdash \psi[val/x] \triangleright \text{Bool} \quad (\text{B.4})$$

Again by Lemma B.3 and Lemma B.7, we have

$$\Gamma, \Gamma'' \vdash t[val/x] \triangleright \tau' \quad (\text{B.5})$$

Combining (B.3), (B.4), (B.5), and the original typing for T , we can establish $\Gamma, \nabla \vdash \text{select}(\widetilde{TB}, T, \psi[val/x], t[val/x], !tbv) \triangleright \Gamma'$, which gives the desired result since we have $\text{select}(\widetilde{TB}, T, \psi, t, !tbv)[val/x] = \text{select}(\widetilde{TB}, T, \psi[val/x], t[val/x], !tbv)$.

Case $\text{update}(tid@l, T, \psi, t)$: Analogous to the previous cases.

Case $\text{aggr}(tid@l, T, \psi, f, T')$: Analogous to the previous cases.

Case $\text{create}(tid@l, sk)$: Trivial.

Case $\text{drop}(tid@l)$: Trivial.

Case $\text{eval}(P')@l$: Because of $(\Gamma, x : \tau_0), \nabla \vdash \text{eval}(P')@l \triangleright \Gamma'$ we have $(\Gamma, x : \tau_0), \nabla \vdash P'$, $\Gamma, x : \tau_0 \vdash l \triangleright \text{Loc}$, and $\Gamma' = []$. Since $bv(\text{eval}(P')@l) \cap \mathbf{D}_{\Gamma, x:\tau_0} = \emptyset$, we have $bv(P') \cap \mathbf{D}_{\Gamma, x:\tau_0} = \emptyset$. By the induction hypothesis we have $\Gamma, \nabla \vdash P'[val/x]$, and by Lemmas B.4 and B.5 we get $\Gamma \vdash l \triangleright \text{Loc}$. Hence we can establish $\Gamma, \nabla \vdash \text{eval}(P'[val/x])@l \triangleright \Gamma'$, which gives the desired result since $(\text{eval}(P')@l)[val/x] = \text{eval}(P'[val/x])@l$.

Case nilp : Trivial.

Case $a'.P'$: By $(\Gamma, x : \tau_0), \nabla \vdash a'.P'$ we have $(\Gamma, x : \tau_0), \nabla \vdash a' \triangleright \Gamma''$ for some Γ'' such that $(\Gamma, x : \tau_0, \Gamma''), \nabla \vdash P'$. By the induction hypothesis we have $\Gamma, \nabla \vdash a'[val/x] \triangleright \Gamma''$. By Lemma B.3 we have $\Gamma, \Gamma'', x : \tau_0 \vdash P'$. By $bv(a'.P') \cap \mathbf{D}_{\Gamma, x:\tau_0} = \emptyset$ we have $bv(P') \cap \mathbf{D}_{\Gamma, x:\tau_0} = \emptyset$. Since bound variables are renamed apart we have $bv(P') \cap bv(a') = \emptyset$, and it follows that $bv(P') \cap \mathbf{D}_{\Gamma, \Gamma'', x:\tau_0} = \emptyset$. By the induction hypothesis we then have $(\Gamma, \Gamma''), \nabla \vdash P'[val/x]$. Hence we can establish $\Gamma, \nabla \vdash (a'.P')[val/x]$.

Case $P_1; P_2$: Analogous.

Case $A(\tilde{e})$: By $(\Gamma, x : \tau_0), \nabla \vdash A(\tilde{e})$, we have $\Gamma, x : \tau_0 \vdash e_1 \triangleright \tau_1, \dots, \Gamma, x : \tau_0 \vdash e_n \triangleright \tau_n$, and $(var_1 : \tau_1, \dots, var_n : \tau_n), \nabla \vdash P'$, where P' is the procedure body. By Lemma B.7, we have $\Gamma \vdash e_1[val/x] \triangleright \tau_1, \dots, \Gamma \vdash e_n[val/x] \triangleright \tau_n$. Hence we can establish $\Gamma, \nabla \vdash A(\tilde{e})[val/x]$ (Note that the substitution is not performed to P').

Case $\text{foreach}(TB, T, \psi, \leq) : P'$: By $(\Gamma, x : \tau_0), \nabla \vdash \text{foreach}(TB, T, \psi, \leq) : P'$ we have $(\Gamma, x : \tau_0), \nabla \vdash TB \triangleright \tau$ for some τ such that $\tau \vdash T \triangleright \Gamma''$ for some Γ'' such that $\Gamma, x : \tau_0, \Gamma'' \vdash \psi \triangleright \text{Bool}$, and $(\Gamma, x : \tau_0, \Gamma''), \nabla \vdash P'$. By Lemma B.4 we have $fv(TB) \subseteq \mathbf{D}_{\Gamma, x:\tau_0}$. Since the only kind of variables that TB can contain are table variables, we

have $fv(TB) \subseteq \mathbf{D}_\Gamma$. By Lemma B.5, and $TB[val/x] = TB$, we have

$$\Gamma, \nabla \vdash TB[val/x] \triangleright \tau \quad (\text{B.6})$$

Since $bv(\text{foreach}(TB, T, \psi, \leq) : P') \cap \mathbf{D}_{\Gamma, x: \tau_0} = \emptyset$, we have $bv(T) \cap \mathbf{D}_{\Gamma, x: \tau_0} = \emptyset$, and $\mathbf{D}_{\Gamma''} \cap \mathbf{D}_{\Gamma, x: \tau_0} = \emptyset$. By Lemma B.3, we have $\Gamma, \Gamma'', x : \tau_0 \vdash \psi \triangleright \text{Bool}$. By Lemma B.7, we have

$$\Gamma, \Gamma'' \vdash \psi[val/x] \triangleright \text{Bool} \quad (\text{B.7})$$

Again by Lemma B.3, we have $(\Gamma, \Gamma'', x : \tau_0), \nabla \vdash P'$. It is not difficult to see that $bv(P') \cap \mathbf{D}_{\Gamma, \Gamma'', x: \tau_0} = \emptyset$. By the induction hypothesis we have

$$(\Gamma, \Gamma''), \nabla \vdash P'[val/x] \quad (\text{B.8})$$

Combining (B.6), (B.7), (B.8), and $\tau \vdash T \triangleright \Gamma''$, we can establish

$$\Gamma, \nabla \vdash (\text{foreach}(TB, T, \psi, \leq) : P')[val/x].$$

(2) We proceed with an induction on the structure of actions and processes. The reasoning is largely similar to that of (1); hence we only present a few representative cases where the differences from (1) can be demonstrated.

Case $\text{insert}(t, tid@l)$: By $(\Gamma, tbv : \tau_0), \nabla \vdash \text{insert}(t, tid@l) \triangleright \Gamma'$ we have $\Gamma, tbv : \tau_0 \vdash t \triangleright \nabla(tid)$, $\Gamma, tbv : \tau_0 \vdash l \triangleright \text{Loc}$, and $\Gamma' = []$. Tuples cannot contain table variables; hence $t[(I, R)/tbv] = t$. By Lemma B.4 we have $fv(t) \subseteq \mathbf{D}_{\Gamma, tbv: \tau_0}$. Hence we also have $fv(t) \subseteq \mathbf{D}_\Gamma$. By Lemma B.5 we have $\Gamma \vdash t \triangleright \nabla(tid)$. Analogously to the same case under (1), we can deduce $\Gamma \vdash l \triangleright \text{Loc}$. We can now establish

$$\Gamma \vdash (\text{insert}(t, tid@l))[(I, R)/tbv] \triangleright \Gamma'.$$

Case $\text{select}(\widetilde{TB}, T, \psi, t, !tbv')$: By $(\Gamma, tbv : \tau_0), \nabla \vdash \text{select}(\widetilde{TB}, T, \psi, t, !tbv') \triangleright \Gamma'$ we have $(\Gamma, tbv : \tau_0), \nabla \vdash TB_1 \triangleright \tau_1, \dots, (\Gamma, tbv : \tau_0), \nabla \vdash TB_n \triangleright \tau_n$, for some τ_1, \dots, τ_n , such that $\text{flatten}_s(\tau_1 \times \dots \times \tau_n) \vdash T \triangleright \Gamma''$, for some Γ'' such that $\Gamma, tbv : \tau_0, \Gamma'' \vdash \psi \triangleright \text{Bool}$, and $\Gamma, tbv : \tau_0, \Gamma'' \vdash t \triangleright \tau'$ for some τ' such that $\Gamma' = [tbv' : \tau']$.

By Lemma B.7 we have $\Gamma, \nabla \vdash TB_1[(I, R)/tbv] \triangleright \tau_1, \dots, \Gamma, \nabla \vdash TB_n[(I, R)/tbv] \triangleright \tau_n$. We have $\psi[(I, R)/tbv] = \psi$ and $t[(I, R)/tbv] = t$, and using Lemma B.4 and Lemma B.5 we have $\Gamma, \Gamma'' \vdash \psi \triangleright \text{Bool}$ and $\Gamma, \Gamma'' \vdash t \triangleright \tau'$. It is essentially only the tables TB_1, \dots, TB_n that are potentially affected by the substitution, and we can now establish

$$\Gamma, \nabla \vdash (\text{select}(\widetilde{TB}, T, \psi, t, !tbv'))[(I, R)/tbv] \triangleright \Gamma'.$$

(3) The proof is still by induction on the structure of actions and processes. We omit the details since they are analogous to those of (1) and (2) above. \square

Lemma B.9. *If $\tau \vdash T \triangleright \Gamma$, then $T \Vdash \tau$.*

Proof.

We first show that if $\tau \vdash W \triangleright \Gamma$, then $W \Vdash \tau$. There are two cases:

$W = !x$: We have $\tau_m \vdash !x \triangleright [x : \tau_m]$ where $\tau = \tau_m$ and $\tau_m \neq \text{Loc}$. The type τ_m can thus be one of *Int*, *String*, *Id*, *Int mset*, *String mset*, *Id mset*, *Loc mset*. By the definition of $\cdot \Vdash \cdot$, we can always establish $!x \Vdash \tau_m$, which is $!x \Vdash \tau$.

$W = !u$: We have $\text{Loc} \vdash !u \triangleright [u : \text{Loc}]$ and $\tau = \text{Loc}$. By the definition of $\cdot \Vdash \cdot$, we can directly establish $!u \Vdash \tau$.

We now show that if $\tau \vdash T \triangleright \Gamma$, then $T \Vdash \tau$. Suppose T is W_1, \dots, W_n . By $\tau \vdash T \triangleright \Gamma$ we have $\tau_1 \vdash W_1 \triangleright \Gamma_1, \dots, \tau_n \vdash W_n \triangleright \Gamma_n$ for some τ_1, \dots, τ_n and $\Gamma_1, \dots, \Gamma_n$ such that $\tau = \tau_1 \times \dots \times \tau_n$ and $\Gamma = \Gamma_1, \dots, \Gamma_n$. Hence we have $W_1 \Vdash \tau_1, \dots, W_n \Vdash \tau_n$. By the definition of $\cdot \Vdash \cdot$ we have $W_1, \dots, W_n \Vdash \tau_1 \times \dots \times \tau_n$, which is $T \Vdash \tau$. \square

Lemma B.10. *If t does not contain operators (\cdot, aop), $\square \vdash t \triangleright \tau$ and $\tau \vdash T \triangleright \Gamma$, then $t/T \neq err$ holds. Additionally, the following statements hold.*

- (1) *If $\Gamma \vdash \psi \triangleright Bool$, and $bv(T) \supseteq fv(\psi)$, then we have $\mathcal{E}[\psi(t/T)] \neq err$.*
- (2) *If $\Gamma \vdash t' \triangleright \tau'$, and $bv(T) \supseteq fv(t')$, then we have $\mathcal{E}[t'(t/T)] \Vdash \tau'$.*

Proof.

- We first show that if t does not contain operators, $\square \vdash t \triangleright \tau$ and $\tau \vdash W \triangleright \Gamma$, then $t/W \neq err$ holds. There are two cases for W :

$W = !x$: It is not difficult to deduce from $\tau \vdash !x \triangleright \Gamma$ that τ is a multiest type τ_m that cannot be *Loc*. By $\square \vdash t \triangleright \tau_m$, the fact that t does not contain operators, and Lemma B.4, we know that t is a *num*, *str*, *tid* or a multiset of intergers/strings/table identifiers/localities. By the definition of pattern matching we have $t/W \neq err$.

$W = !u$: By $\tau \vdash !u \triangleright \Gamma$ we have $\tau = Loc$. By $\square \vdash t \triangleright Loc$ we have $t = l$ is a locality constant. We indeed have $l/!u \neq err$.

- We then show that if t does not contain operators, $\square \vdash t \triangleright \tau$ and $\tau \vdash T \triangleright \Gamma$, then $t/W \neq err$ holds. From $\tau \vdash T \triangleright \Gamma$ we have $W_1, \dots, W_n, \tau_1, \dots, \tau_n$, and $\Gamma_1, \dots, \Gamma_n$ such that $T = W_1, \dots, W_n, \tau_1 \vdash W_1 \triangleright \Gamma_1, \dots, \tau_n \vdash W_n \triangleright \Gamma_n, \tau = \tau_1 \times \dots \times \tau_n$ and $\Gamma = \Gamma_1, \dots, \Gamma_n$. By $\square \vdash t \triangleright \tau_1 \times \dots \times \tau_n$ we have $t = e_1, \dots, e_n$ for expressions e_1, \dots, e_n such that $\square \vdash e_1 \triangleright \tau_1, \dots, \square \vdash e_n \triangleright \tau_n$. Since t does not contain operators, based on the latter typing judgments for n expressions we can establish the typing of n *singleton tuples* e_1, \dots, e_n of the same form. Hence using previous result we have $e_1/W_1 \neq err, \dots, e_n/W_n \neq err$. We now have $(e_1, \dots, e_n)/(W_1, \dots, W_n) \neq err$, or $t/T \neq err$.
- We next show point (1) in the statement of the lemma. Continuing the argumentation above, since $\square \vdash e_1 \triangleright \tau_1$ and $\tau_1 \vdash W_1 \triangleright \Gamma_1$, where e_1 does not contain variables or operators, we have $\Gamma' \vdash \psi[e_1/W_1] \triangleright Bool$ using Lemma B.3 and Lemma B.7, where Γ' is obtained by removing from Γ the binding for the only variable in $bv(W_1)$. Similarly we have $\Gamma'' \vdash \psi[e_1/W_1, e_2/W_2] \triangleright Bool, \dots, \Gamma^{(n)} \vdash \psi[e_1/W_1, \dots, e_n/W_n] \triangleright Bool$, for some $\Gamma'', \dots, \Gamma^{(n)}$, that are obtained from Γ' by subsequently removing the bindings for the variables of $bv(W_2), \dots, bv(W_n)$. The last judgment is actually $\Gamma^{(n)} \vdash \psi(t/T) \triangleright Bool$. Since $bv(T) \supseteq fv(\psi)$, we have $fv(\psi(t/T)) = \emptyset$. By Lemma B.5 we have $\square \vdash \psi(t/T) \triangleright Bool$. By Lemma B.6 we have $\mathcal{E}[\psi(t/T)] \neq err$.
- The reasoning needed to establish point (2) in the statement of the lemma is analogous. \square

Lemma B.11. *If each component of t is a constant or a variable bound in T , and $\exists \Gamma : \tau \vdash T \triangleright \Gamma \wedge \Gamma \vdash t \triangleright \tau'$, then we have $\tau \downarrow_t^T = \tau'$.*

Proof. Suppose $T = W_1, \dots, W_k, t = e_1, \dots, e_n$. Then $\tau \downarrow_t^T$ is an n -ary product type. Assume $\tau \downarrow_t^T = \tau_1'' \times \dots \times \tau_n''$ for some $\tau_1'', \dots, \tau_n''$. By $\tau \vdash W_1, \dots, W_k \triangleright \Gamma$ we have $\tau = \tau_1 \times \dots \times \tau_k$ for some τ_1, \dots, τ_k . By $\Gamma \vdash e_1, \dots, e_n \triangleright \tau'$ we have $\tau' = \tau_1' \times \dots \times \tau_n'$ for some τ_1', \dots, τ_n' . We show that for each $i \in \{1, \dots, n\}$, $\tau_i'' = \tau_i'$, which will give $\tau \downarrow_t^T = \tau'$. Fixing i , we have two cases for e_i .

e_i is a constant value: By $\Gamma \vdash e_1, \dots, e_n \triangleright \tau'_1 \times \dots \times \tau'_n$ we have $\Gamma \vdash e_i \triangleright \tau'_i$. By Lemma B.5 we have $\Box \vdash e_i \triangleright \tau'_i$. Since e_i is a value ($\mathcal{E}[\![e_i]\!] = e_i$), by Lemma B.6 we have $e_i \Vdash \tau'_i$. On the other hand, by $\tau \downarrow_t^T = \tau''_1 \times \dots \times \tau''_n$ we have $e_i \Vdash \tau''_i$. Hence $\tau''_i = \tau'_i$ holds.

e_i is a variable bound in W_j : By $\tau_1 \times \dots \times \tau_k \vdash T \triangleright \Gamma$ we have the binding $e_i : \tau_j$ in Γ . By $\Gamma \vdash e_1, \dots, e_n \triangleright \tau'_1 \times \dots \times \tau'_n$ we have $\Gamma \vdash e_i \triangleright \tau'_i$ and thus $\tau'_i = \tau_j$. By $\tau_1 \times \dots \times \tau_k \downarrow_t^T = \tau''_1 \times \dots \times \tau''_n$ we also have $\tau''_i = \tau_j$. Hence $\tau''_i = \tau'_i$ holds.

This completes the proof. \square

Lemma B.12. *If $\Gamma, \nabla \vdash TB_1 \triangleright \tau_1, \dots, \Gamma, \nabla \vdash TB_n \triangleright \tau_n$, and $\Gamma, \nabla \vdash (I_k, R_k)$ holds for each (I_k, R_k) in the list $\langle l_i :: (I_i, R_i) \rangle_i$, then $\otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) = \text{flatten}_s(\tau_1 \times \dots \times \tau_n)$ holds as long as $\otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$ is defined.*

Proof. According to the definition of $\otimes_{\text{sk}}(\cdot, \cdot)$, if it is defined, there exist τ'_1, \dots, τ'_n such that

$$\otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) = \text{flatten}_s(\tau'_1 \times \dots \times \tau'_n),$$

where for each j , τ'_j falls into one of the following two cases, in each of which we show that $\tau'_j = \tau_j$.

- There exists some k such that $TB_j = I_k.\text{tid}@l_k$ and $\tau'_j = I_k.sk$. By $\Gamma, \nabla \vdash (I_k, R_k)$ we have $\nabla(I_k.\text{tid}) = I_k.sk$. From the conditions we also have $\Gamma, \nabla \vdash I_k.\text{tid}@l_k \triangleright \tau_j$. Hence we have $\tau_j = \nabla(I_k.\text{tid}) = I_k.sk$. Therefore we have $\tau'_j = \tau_j$.
- There exists some data set R_0 such that $TB_j = (I_0, R_0)$ and $\tau'_j = I_0.sk$. From the conditions we have $\Gamma, \nabla \vdash (I_0, R_0) \triangleright \tau_j$; hence $\tau_j = I_0.sk$. Therefore we have $\tau'_j = \tau_j$.

This completes the proof. \square

Lemma B.13. *If $\Gamma, \nabla \vdash TB_1 \triangleright \tau_1, \dots, \Gamma, \nabla \vdash TB_n \triangleright \tau_n$, and $\Gamma, \nabla \vdash (I_k, R_k)$ holds for each (I_k, R_k) in the list $\langle l_i :: (I_i, R_i) \rangle_i$, then $\forall t' \in \otimes_R(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) : \Box \vdash t' \triangleright \text{flatten}_s(\tau_1 \times \dots \times \tau_n)$ holds as long as $\otimes_R(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$ is defined.*

Proof Idea. The proof idea is analogous to that of Lemma B.12. A case analysis based on the definition of $\otimes_R(\cdot, \cdot)$ is needed. \square

Lemma B.14. *If $\Gamma, \nabla \vdash (I, R) \triangleright \tau$, and t is a tuple, then $\Gamma, \nabla \vdash (I, R \setminus \{t\}) \triangleright \tau$.*

The proof of Lemma B.14 is trivial.

Lemma B.15. *If N_{env} and N are closed, $bv(N) \cap \mathbf{D}_\Gamma = \emptyset$, $\Gamma, \nabla \vdash N_{\text{env}}$, $\Gamma, \nabla \vdash N$, and $N_{\text{env}} \vdash N \rightarrow N'$, then $\Gamma, \nabla \vdash N'$.*

Proof. We proceed with an induction on the derivation of $N_{\text{env}} \vdash N \rightarrow N'$.

Case (INS): The net N is $l_1 :: \text{insert}(t, \text{tid}@l_2).P \parallel l_2 :: (I, R)$. By

$$\Gamma, \nabla \vdash l_1 :: \text{insert}(t, \text{tid}@l_2).P \parallel l_2 :: (I, R)$$

we have $\Gamma, \nabla \vdash \text{insert}(t, \text{tid}@l_2).P$ and $\Gamma, \nabla \vdash (I, R)$; hence $\Gamma, \nabla \vdash \text{insert}(t, \text{tid}@l_2) \triangleright \Box$ and (by $\Gamma, \Box = \Gamma$)

$$\Gamma, \nabla \vdash P. \tag{B.9}$$

By the typing of the insertion action we have $\Gamma \vdash t \triangleright \nabla(\text{tid})$, and $\Gamma \vdash l_2 \triangleright \text{Loc}$. By $\Gamma, \nabla \vdash (I, R)$ we have $\nabla(I.\text{tid}) = I.sk$. Since the transition can happen we have $I.\text{tid} = \text{tid}$ by the semantic rule (INS). Hence $\Gamma \vdash t \triangleright I.sk$. Since N is closed, we have $fv(t) = \emptyset$. By

Lemma B.5 we have $\Box \vdash t \triangleright I.sk$. By Lemma B.6 we have $\mathcal{E}[\![t]\!] \Vdash I.sk$. Hence by (INS) we have

$$N' = l_1 :: P \parallel l_2 :: (I, R \uplus \{\mathcal{E}[\![t]\!]\}).$$

By $\Gamma \vdash t \triangleright I.sk$ and $\Gamma, \nabla \vdash (I, R)$ we have

$$\Gamma, \nabla \vdash (I, R \uplus \{\mathcal{E}[\![t]\!]\}) \quad (\text{B.10})$$

By (B.9) and (B.10), it is not difficult to derive $\Gamma, \nabla \vdash N'$.

Case (DEL): The net N is $l_1 :: \text{delete}(tid@l_2, T, \psi).P \parallel l_2 :: (I, R)$. By

$$\Gamma, \nabla \vdash l_1 :: \text{delete}(tid@l_2, T, \psi).P \parallel l_2 :: (I, R)$$

we have $\Gamma, \nabla \vdash \text{delete}(tid@l_2, T, \psi).P$ and $\Gamma, \nabla \vdash (I, R)$.

Hence $\Gamma, \nabla \vdash \text{delete}(tid@l_2, T, \psi) \triangleright \Box$, and

$$\Gamma, \nabla \vdash P. \quad (\text{B.11})$$

By the typing of the deletion action we have $\nabla(tid) \vdash T \triangleright \Gamma''$ for some Γ'' such that $\Gamma, \Gamma'' \vdash \psi \triangleright Bool$, and $\Gamma \vdash l_2 \triangleright Loc$. Since the transition can take place we have $I.tid = tid$. By $\Gamma, \nabla \vdash (I, R)$ we have $\nabla(I.tid) = I.sk$. Hence we have $I.sk \vdash T \triangleright \Gamma''$. By Lemma B.9 we have $T \Vdash I.sk$. By $\Gamma, \nabla \vdash (I, R)$ we have $\forall t \in R : \Gamma \vdash t \triangleright I.sk$. Pick an arbitrary t from the data set R , since $fv(t) = \emptyset$, by Lemma B.5 we have $\Box \vdash t \triangleright I.sk$. By Lemma B.10 we have $t/T \neq err$. By the closedness of N we have $fv(\psi) \subseteq bv(T) = \mathbf{D}_{\Gamma''}$. Hence by Lemma B.5 we have $\Gamma'' \vdash \psi \triangleright Bool$. By Lemma B.10 we have $\mathcal{E}[\![\psi(t/T)]\!] \neq err$. We are now able to assert that

$$N' = l_1 :: P \parallel l_2 :: (I, \{t \in R \mid \mathcal{E}[\![\psi(t/T)]\!] \neq tt\}).$$

It is not difficult to derive

$$\Gamma, \nabla \vdash (I, \{t \in R \mid \mathcal{E}[\![\psi(t/T)]\!] \neq tt\}). \quad (\text{B.12})$$

By (B.11) and (B.12) it is not difficult to derive $\Gamma, \nabla \vdash N'$.

Case (SEL): The net N is

$$l_0 :: \text{select}(\widetilde{TB}, T, \psi, t, !tbv).P \parallel l_2 :: (I_1, R_1) \parallel \dots \parallel l_k :: (I_k, R_k).$$

By the typing of N under Γ and ∇ we have $\Gamma, \nabla \vdash \text{select}(\widetilde{TB}, T, \psi, t, !tbv).P$, $\Gamma, \nabla \vdash (I_1, R_1)$, ..., and $\Gamma, \nabla \vdash (I_n, R_n)$; hence we have $\Gamma, \nabla \vdash \text{select}(\widetilde{TB}, T, \psi, t, !tbv) \triangleright \Gamma''$ for some Γ'' such that $(\Gamma, \Gamma''), \nabla \vdash P$. By the typing of the selection action we have $\Gamma, \nabla \vdash TB_1 \triangleright \tau_1$, ..., $\Gamma, \nabla \vdash TB_n \triangleright \tau_n$,

$$\text{flatten}_s(\tau_1 \times \dots \times \tau_n) \vdash T \triangleright \Gamma''', \quad (\text{B.13})$$

for some Γ''' such that $\Gamma, \Gamma''' \vdash \psi \triangleright Bool$ and $\Gamma, \Gamma''' \vdash t \triangleright \tau'$, for some τ' such that $\Gamma'' = [tbv : \tau']$.

Since the transition can take place, $\otimes_{sk}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$ and $\otimes_R(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$ are both defined. By Lemma B.12 we have $\otimes_{sk}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) = \text{flatten}_s(\tau_1 \times \dots \times \tau_n)$. By Lemma B.9 and (B.13), we have $T \Vdash \text{flatten}_s(\tau_1 \times \dots \times \tau_n)$; hence it holds that $T \Vdash \otimes_{sk}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$.

Pick an arbitrary $t' \in \otimes_R(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$. By Lemma B.13 we have

$$\Box \vdash t' \triangleright \text{flatten}_s(\tau_1 \times \dots \times \tau_n).$$

Hence by Lemma B.10 we have $t'/T \neq \text{err}$. Since N is closed we have $fv(\psi) \subseteq bv(T) = \mathbf{D}_{\Gamma''}$. By Lemma B.5, we have $\Gamma'' \vdash \psi \triangleright \text{Bool}$. By Lemma B.10, we have $\mathcal{E}[\psi(t'/T)] \neq \text{err}$. We also have $fv(t) \subseteq bv(T) = \mathbf{D}_{\Gamma''}$. By Lemma B.5, we have

$$\Gamma'' \vdash t \triangleright \tau'. \quad (\text{B.14})$$

By Lemma B.10, we have $\mathcal{E}[t(t'/T)] \Vdash \tau'$, which implies $\mathcal{E}[t(t'/T)] \neq \text{err}$. We can now assert that

$$N' = l_0 :: P[(I', R')/tbv] \parallel l_1 :: (I_1, R_1) \parallel \dots \parallel l_k :: (I_k, R_k),$$

where

$$I' = (\perp, \otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) \downarrow_t^T)$$

and

$$R' = \{\mathcal{E}[t\sigma] \mid \exists t' \in \otimes_{\text{R}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) : t'/T = \sigma \wedge \mathcal{E}[\psi\sigma] = tt\}.$$

We proceed with showing $\Gamma, \nabla \vdash P[(I', R')/tbv]$. This boils down to showing $\square, \nabla \vdash (I', R') \triangleright \tau'$ by Lemma B.8 since we already have $(\Gamma, tbv : \tau'), \nabla \vdash P, bv(P) \cap \mathbf{D}_{\Gamma, tbv:\tau'} = \emptyset$ (which is obvious since variables are renamed apart and $bv(N) \cap \mathbf{D}_{\Gamma} \neq \emptyset$ holds).

Take an arbitrary tuple from R' it must be expressible as $\mathcal{E}[t(t'/T)]$ for some $t' \in \otimes_{\text{R}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i)$ such that $\mathcal{E}[\psi(t'/T)] = tt$. By previous result we have $\mathcal{E}[t(t'/T)] \Vdash \tau'$. It is not difficult to see that $\mathcal{E}[\mathcal{E}[t(t'/T)]] = \mathcal{E}[t(t'/T)]$; hence $\mathcal{E}[\mathcal{E}[t(t'/T)]] \Vdash \tau'$. By Lemma B.6, we have $\square \vdash \mathcal{E}[t(t'/T)] \triangleright \tau'$. On the other hand, we have $I'.sk = \otimes_{\text{sk}}(\widetilde{TB}, \langle l_i :: (I_i, R_i) \rangle_i) \downarrow_t^T = \text{flatten}_s(\tau_1 \times \dots \times \tau_n) \downarrow_t^T$. By (B.13), (B.14) and Lemma B.11 we have $\text{flatten}_s(\tau_1 \times \dots \times \tau_n) \downarrow_t^T = \tau'$. Hence $I'.sk = \tau'$ and $\square \vdash \mathcal{E}[t(t'/T)] \triangleright I'.sk$. It can be derived that $\square, \nabla \vdash (I', R') \triangleright \tau'$. We can now establish

$$\Gamma, \nabla \vdash P[(I', R')/tbv] \quad (\text{B.15})$$

Using (B.15) we can now easily derive $\Gamma, \nabla \vdash N'$.

Case (UPD): The net N is $l_1 :: \text{update}(tid@l_2, T, \psi, t).P \parallel l_2 :: (I, R)$. By the typing of N under Γ and ∇ , we have $\Gamma, \nabla \vdash \text{update}(tid@l_2, T, \psi, t).P$ and $\Gamma, \nabla \vdash (I, R)$; hence we have $\Gamma, \nabla \vdash \text{update}(tid@l_2, T, \psi, t) \triangleright \square$ and

$$\Gamma, \nabla \vdash P. \quad (\text{B.16})$$

By the typing of the update action, we have $\nabla(tid) \vdash T \triangleright \Gamma''$, for some Γ'' such that $\Gamma, \Gamma'' \vdash \psi \triangleright \text{Bool}$ and $\Gamma, \Gamma'' \vdash t \triangleright \nabla(tid)$, and $\Gamma \vdash l_2 \triangleright \text{Loc}$.

By reasoning similar to the case (DEL), we have $T \Vdash I.sk$. Pick an arbitrary $t' \in R$, again by reasoning similar to the case (DEL), we can deduce $t'/T \neq \text{err}$ and $\mathcal{E}[\psi(t'/T)] \neq \text{err}$. Using $\Gamma, \nabla \vdash (I, R)$ we have $\Gamma \vdash t' \triangleright I.sk$ and $\nabla(I.tid) = I.sk$. Since the transition can take place, we have $tid = I.tid$. Hence $\Gamma \vdash t' \triangleright \nabla(tid)$, and $\square \vdash t' \triangleright \nabla(tid)$ holds by Lemma B.5 and the fact that $fv(t') = \emptyset$. Since N is closed, $fv(t) \subseteq bv(T) = \mathbf{D}_{\Gamma''}$. Hence $\Gamma'' \vdash t \triangleright \nabla(tid)$. By Lemma B.10 we have $\mathcal{E}[t(t'/T)] \Vdash \nabla(tid)$, which gives $\mathcal{E}[t(t'/T)] \Vdash I.sk$. Hence it also holds that $\mathcal{E}[t(t'/T)] \neq \text{err}$. We can now assert that

$$N' = l_1 :: P \parallel l_2 :: (I, R'_1 \uplus R'_2),$$

where $R'_1 = \{t' \in R \mid \mathcal{E}[\psi(t'/T)] \neq tt\}$ and $R'_2 = \{\mathcal{E}[t\sigma] \mid \exists t' : t' \in R \wedge t'/T = \sigma \wedge \mathcal{E}[\psi\sigma] = tt\}$.

We proceed with showing $\Gamma, \nabla \vdash (I, R'_1 \uplus R'_2)$. Pick arbitrary $t'' \in R'_1 \uplus R'_2$. We show that $\Gamma \vdash t'' \triangleright I.sk$ with a case analysis on t'' .

Case $t'' \in R'_1$: We have $t'' \in R$; hence it is obvious that $\Gamma \vdash t'' \triangleright I.sk$ holds.

Case $t'' \in R'_2$: We have $t'' = \mathcal{E}[\![t'(t'/T)]\!]$ for some $t' \in R$ such that $\mathcal{E}[\![\psi(t'/T)]\!] = tt$. By previous results we have $\mathcal{E}[\![t'(t'/T)]\!] \Vdash I.sk$. Using Lemma B.6 we have $\Box \vdash \mathcal{E}[\![t'(t'/T)]\!] \triangleright I.sk$. Weakening the typing environment we obtain $\Gamma \vdash \mathcal{E}[\![t'(t'/T)]\!] \triangleright I.sk$.

It is now straightforward to obtain

$$\Gamma, \nabla \vdash (I, R'_1 \uplus R'_2) \quad (\text{B.17})$$

Using (B.16) and (B.17) we can easily derive $\Gamma, \nabla \vdash N'$.

Case (AGR): The net N is $l_1 :: \text{aggr}(tid@l_2, T, \psi, f, T').P \parallel l_2 :: (I, R)$. By the typing of N under Γ and ∇ , we have $\Gamma, \nabla \vdash \text{aggr}(tid@l_2, T, \psi, f, T').P$ and $\Gamma, \nabla \vdash (I, R)$; hence we have $\Gamma, \nabla \vdash \text{aggr}(tid@l_2, T, \psi, f, T') \triangleright \Gamma''$ for some Γ'' such that

$$(\Gamma, \Gamma''), \nabla \vdash P. \quad (\text{B.18})$$

By the typing of the aggregation action we have $\nabla(tid) \vdash T \triangleright \Gamma''$, for some Γ'' such that $\Gamma, \Gamma'' \vdash \psi \triangleright Bool$, and $f : \nabla(tid) \text{ mset} \rightarrow \tau'$, for some τ' such that $\tau' \vdash T' \triangleright \Gamma''$, and $\Gamma \vdash l_2 \triangleright Loc$.

By reasoning similar to the cases (DEL) and (UPD), we have $T \Vdash I.sk$. Pick arbitrary $t \in R$, we also have $t/T \neq err$ and $\mathcal{E}[\![\psi(t/T)]\!] \neq err$ analogously to the two aforementioned cases. Picking arbitrary τ_1 and τ_2 such that $f : \tau_1 \text{ mset} \rightarrow \tau_2$, we have $\tau_1 = \nabla(tid)$ and $\tau_2 = \tau'$. Thus by $\Gamma, \nabla \vdash (I, R)$ we can deduce $\Gamma \vdash t \triangleright \tau_1$. Since $fv(t) = \emptyset$, by Lemma B.5 we have $\Box \vdash t \triangleright \tau_1$. By Lemma B.6 and the obvious fact $\mathcal{E}[\![t]\!] = t$, we have $t \Vdash \tau_1$. By Lemma B.9 we have $T' \Vdash \tau_2$. We can now assert

$$N' = l_1 :: P(t'/T') \parallel l_2 :: (I, R),$$

where $t' = f(\{t \in R \mid \mathcal{E}[\![\psi(t/T)]\!] = tt\})$.

Since the range of f is of type τ_2 , we have $t' \Vdash \tau_2$, or $\mathcal{E}[\![t']\!] \Vdash \tau_2$. By Lemma B.6 we have $\Box \vdash t' \triangleright \tau_2$. Since t' is a function value, its components are values. Suppose $t' = val_1, \dots, val_n$. We have some $\tau_m^1, \dots, \tau_m^n$ such that $\Box \vdash val_1 \triangleright \tau_m^1, \dots, \Box \vdash val_n \triangleright \tau_m^n$. By $\tau_2 \vdash T' \triangleright \Gamma''$, we have $T' = !var_1, \dots, !var_n$, where each var_j is some x_j or u_j , $\tau_m^1 \vdash !var_1 \triangleright [var_1 : \tau_m^1], \dots, \tau_m^n \vdash !var_n \triangleright [var_n : \tau_m^n]$, $\Gamma'' = [var_1 : \tau_m^1, \dots, var_n : \tau_m^n]$. By Lemma B.10 we have $t'/T' \neq err$; hence $t'/T' = [val_1/var_1, \dots, val_n/var_n]$. By (B.18), and Lemma B.8, we have

$$\Gamma, \nabla \vdash P(t'/T') \quad (\text{B.19})$$

Using (B.19) we can easily derive $\Gamma, \nabla \vdash N'$.

Case (CRT): Trivial.

Case (DRP): Trivial.

Case (EVL): The net N is $l_1 :: \text{eval}(P)@l_2.P' \parallel l_2 :: \text{nil}_P$. By the typing of N under Γ and ∇ , we have $\Gamma, \nabla \vdash \text{eval}(P)@l_2.P'$ and $\Gamma, \nabla \vdash \text{nil}_P$; hence we have $\Gamma, \nabla \vdash \text{eval}(P)@l_2 \triangleright \Box$ and $\Gamma, \nabla \vdash P'$. By the typing of the **eval** action we have $\Gamma, \nabla \vdash P$, and $\Gamma \vdash l_2 \triangleright Loc$. We have $N' = l_1 :: P' \parallel l_2 :: P$, and we can easily derive $\Gamma, \nabla \vdash N'$.

Case (FOR^{tt}): The net N is $l :: \text{foreach}((I, R), T, \psi, \leq) : P$. By the typing of N under Γ and ∇ we have $\Gamma, \nabla \vdash (I, R) \triangleright \tau$ for some τ such that $\tau \vdash T \triangleright \Gamma''$, for some Γ'' such that $\Gamma, \Gamma'' \vdash \psi \triangleright Bool$ and $(\Gamma, \Gamma''), \nabla \vdash P$.

On the other hand, according to the rule (FOR^{tt}) we have

$$N' = l :: P(t_0/T); (\text{foreach}((I, R \setminus \{t_0\}), T, \psi, \leq) : P),$$

where $t_0 \in \text{Minimal}(R, \leq)$.

Since $t_0 \in R$, by $\Gamma, \nabla \vdash (I, R) \triangleright \tau$ we have $\Gamma \vdash t_0 \triangleright I.sk$ and $\tau = I.sk$. By reasoning similar to that in the case (AGR), we can deduce $\Gamma, \nabla \vdash P(t_0/T)$. It is not difficult to deduce $\Gamma, \nabla \vdash (I, R \setminus \{t_0\}) \triangleright \tau$; hence $\Gamma, \nabla \vdash \text{foreach}((I, R \setminus \{t_0\}), T, \psi, \leq) : P$ can be established. Now it is straightforward to derive $\Gamma, \nabla \vdash N'$.

Case (FOR^{ff}): Similar to the case (FOR^{tt}), we have $\Gamma, \nabla \vdash (I, R) \triangleright \tau$ for some τ such that $\tau \vdash T \triangleright \Gamma''$, for some Γ'' such that $\Gamma, \Gamma'' \vdash \psi \triangleright Bool$ and $(\Gamma, \Gamma''), \nabla \vdash P$.

Pick an arbitrary $t_0 \in R$. By $\Gamma, \nabla \vdash (I, R) \triangleright \tau$ we have $\Gamma \vdash t_0 \triangleright I.sk$ and $\tau = I.sk$. By reasoning analogous to that used in the case (DEL), we can deduce that $t_0/T \neq err$ and $\mathcal{E}[\psi(t_0/T)] \neq err$. Hence $N' = l :: \text{nil}_P$ and it is trivial to establish $\Gamma, \nabla \vdash N'$.

Case (SEQ^{tt}): Straightforward.

Case (SEQ^{ff}): Straightforward.

Case (CALL): The net N is $l :: A(\tilde{e})$. By $\Gamma, \nabla \vdash l :: A(\tilde{e})$ we have $\Gamma \vdash e_1 \triangleright \tau_1, \dots, \Gamma \vdash e_n \triangleright \tau_n$, and $(var_1 : \tau_1, \dots, var_n : \tau_n), \nabla \vdash P$. Since N is closed, none of e_1, \dots , and e_n contains variables. By Lemma B.5 we have $\square \vdash e_1 \triangleright \tau_1, \dots, \square \vdash e_n \triangleright \tau_n$. By Lemma B.6, we have $\mathcal{E}[e_1] \Vdash \tau_1, \dots, \mathcal{E}[e_n] \Vdash \tau_n$, or $v_1 \Vdash \tau_1, \dots, v_n \Vdash \tau_n$. Using Lemma B.6 again, we have $\square \vdash v_1 \triangleright \tau_1, \dots, \square \vdash v_n \triangleright \tau_n$. By Lemma B.8 we have $\square, \nabla \vdash P[v_1/var_1] \dots [v_n/var_n]$. By Lemma B.4 we have $fv(P[v_1/var_1] \dots [v_n/var_n]) = \emptyset$. Weakening the typing environment we get $\Gamma, \nabla \vdash P[v_1/var_1] \dots [v_n/var_n]$. Hence it is straightforward to establish $\Gamma, \nabla \vdash l :: P[v_1/var_1] \dots [v_n/var_n]$.

Case (PAR): The net N is $N_1 || N_2$. By $\Gamma, \nabla \vdash N_1 || N_2$, we have $\Gamma, \nabla \vdash N_1$ and $\Gamma, \nabla \vdash N_2$. By the rule (PAR) we have $N_2 \vdash N_1 \rightarrow N'_1$. By the induction hypothesis we have $\Gamma, \nabla \vdash N'_1$. For $N' = N'_1 || N_2$ it is straightforward to establish $\Gamma, \nabla \vdash N'$.

Case (RES): Trivial.

Case (EQUIV): Straightforward with the help of Lemma B.2.

This completes the proof. \square

We restate Theorem 4.2 and point out that it is a consequence of Lemma B.15.

Theorem 4.2 (Subject Reduction). *For a closed net N , if $bv(N) \cap \mathbf{D}_\Gamma = \emptyset$, $\Gamma, \nabla \vdash N$, and $\vdash N \rightarrow N'$, then $\Gamma, \nabla \vdash N'$.*

Proof. Theorem 4.2 can be directly obtained from Lemma B.15 by instantiating the latter with $N_{\text{env}} = \text{nil}_N$. \square

Lemma B.16. *Suppose the well-typedness of each procedure body has already been decided with a typing environment that associates its formal parameters to their declared types. With a given ∇ , and the assumptions that it takes $O(1)$ time to*

- *produce the types of all constant expressions,*
- *decide the equality/inequality of table identifiers and types,*
- *construct a singleton typing environment,*
- *construct the extension (Γ_1, Γ_2) of a typing environment Γ_1 , and*
- *look up environments Γ (resp. ∇) for variables (resp. table identifiers),*

and that it takes $O(n)$ time to

- *form an n -ary product type, and*
- *perform the operation $\text{flatten}_s(\tau_1 \times \dots \times \tau_n)$,*

the following results hold

- (1) $\Gamma \vdash e \triangleright \tau$ can be decided in time $O(sz_asc(e))$
- (2) $\Gamma \vdash \psi \triangleright Bool$ can be decided in time $O(sz_asc(\psi))$
- (3) $\Gamma \vdash t \triangleright \tau$ can be decided in time $O(sz_asc(t))$
- (4) $\tau \vdash T \triangleright \Gamma'$ can be decided in time $O(sz_asc(T))$
- (5) $\Gamma, \nabla \vdash TB \triangleright \tau$ can be decided in time $O(sz_asc(TB))$
- (6) $\Gamma, \nabla \vdash a \triangleright \Gamma'$ can be decided in time $O(sz_asc(a))$, and
 $\Gamma, \nabla \vdash P$ can be decided in time $O(sz_asc(P))$
- (7) $\Gamma, \nabla \vdash C$ can be decided in time $O(sz_asc(C))$
- (8) $\Gamma, \nabla \vdash N$ can be decided in time $O(sz_asc(N))$

Sketch of Proof. A general observation is that the type system does not make use of subtyping and for each syntactical entity, there can only be one typing rule applicable. On this basis, the proofs of (1)–(8) are by straightforward induction on the structure of the corresponding syntactical entities. \square

Theorem 4.6 (Efficiency of Type Checking). *With a given ∇ , the time complexity of type checking net N is linear in $size(N)$, provided that it takes $O(1)$ time to*

- *determine the types of all constant expressions,*
- *decide the equality/inequality of table identifiers and types,*
- *construct a singleton typing environment,*
- *construct the extension (Γ_1, Γ_2) of a typing environment Γ_1 with Γ_2 , and*
- *look up environments Γ (resp. ∇) for variables (resp. table identifiers),*

and that it takes $O(n)$ time to

- *form an n -ary product type, and*
- *perform the operation $flatten_s(\tau_1 \times \dots \times \tau_n)$.*

Proof. We type check N in two steps:

- (1) we type check all the bodies of the procedures defined, with typing environments that associate their formal parameters to their declared types;
- (2) if one of these bodies fails to type check, we abort the task, declaring that N does not type check; otherwise we type check N .

In both steps, when a procedure call is encountered (in a procedure body or N), we simply ignore the procedure body.

It takes $O(\sum_{A(\{var_i:\tau_i\}_i) \triangleq P} sz_asc(P))$ time to perform step (1). By Lemma B.16, it takes $O(sz_asc(N))$ time to perform step (2). Therefore it is obvious that the time complexity of type checking N is linear in $size(N)$, which is

$$sz_asc(N) + \sum_{A(\{var_i:\tau_i\}_i) \triangleq P} sz_asc(P). \quad \square$$