

An Operational Semantics for Network Datalog

Vivek Nigam¹, Limin Jia², Anduo Wang¹,
Boon Thau Loo¹, and Andre Scedrov¹

¹ University of Pennsylvania, Philadelphia, USA
{vnigam,scedrov}@math.upenn.edu, {anduo,boonloo}@seas.upenn.edu

² Carnegie-Mellon University, Pittsburg, USA
liminjia@cmu.edu

Abstract. Network Datalog (*NDlog*) is a recursive query language that extends Datalog by allowing programs to be distributed in a network. In our initial efforts to formally specify *NDlog*'s operational semantics, we have found several problems with the current evaluation algorithm used, including unsound results, untended multiple derivations of the same table entry, and divergence. In this paper, we make a first step towards correcting these problems by formally specifying a new operational semantics for *NDlog* and proving its correctness for the fragment of non-recursive programs. Our formalization uses linear logic with sub-exponentials. We also argue that if termination is guaranteed, then the results also extend to recursive programs. Finally, we identify a number of potential implementation improvements to *NDlog*.

1 Introduction

2 Declarative networking [7–10] is based on the observation that network protocols
3 deal at their core with using basic information locally available, *e.g.*, neighbor
4 tables, to compute and maintain distributed states, *e.g.*, routes. In this frame-
5 work, network protocols are specified using a declarative logic-based recursive
6 query language called *Network Datalog (NDlog)*, which can be seen as a dis-
7 tributed variant of Datalog [16]. In prior work, it has been shown that tradi-
8 tional routing protocols can be specified in a few lines of declarative code [10],
9 and complex protocols such as Chord distributed hash table [18] in orders of mag-
10 nitude less code [9] compared to traditional imperative implementations. This
11 compact and high-level specifications enable rapid prototype development, ease
12 of customization, optimizability, and the potentiality for protocol verification.
13 When executed, these declarative networks result in efficient implementations,
14 as demonstrated in open-source implementations [15, 17].

15 An inherent feature in networking is the change of local states due to usually
16 small and incremental changes in the network topology. For example, a node
17 might need to change its local routing tables whenever a preferred connection
18 becomes available or when it is no longer available. Reconstructing a node's
19 local state from scratch whenever there is a change in topology is impractical,
20 as it would incur unnecessarily high communication overhead. For instance, in
21 the path-vector protocol used in Internet routing, recomputation from-scratch

22 would require all nodes to exchange all routing information, including those that
23 have been previously propagated.

24 Therefore in declarative networking, nodes maintain their local states incre-
25 mentally as new route messages are received from their neighbors. In literature,
26 there are well known techniques for maintaining databases incrementally [6], in
27 the form of *materialized views*, based in the traditional *semi-naïve* (SN) [2] eval-
28 uation strategy for Datalog programs. In order to accommodate these techniques
29 to a distributed setting, Loo *et al.* in [7] proposed a *pipelined semi-naïve* (PSN)
30 evaluation strategy for *NDlog* programs. PSN relaxes SN by allowing a node to
31 change its local state by following a local pipeline of update messages, specifying
32 the insertions and deletions scheduled to be performed to its local state.

33 Due to the complexity of combining incremental database view maintenance
34 with data and rule distribution, until now, there is no formal specification of
35 PSN nor a correctness proof. As PSN allows each node to compute its local
36 fixed point and disregard global update ordering, PSN does not necessarily pre-
37 serve the semantics of the centralized SN algorithm. However, in a distributed
38 setting, centralized SN evaluation is not practical. Therefore, studying the cor-
39 rectness properties of a distributed SN evaluation is crucial to the correctness of
40 declarative networking.

41 In this paper, we aim to give formal treatment of the operational semantics
42 of PSN and prove its correctness. In the process, we identify several problems
43 with PSN, namely, that it can yield unsound results; it can diverge; and it can
44 compute the same derivation multiple times. In order to address these deficien-
45 cies, we present a new evaluation algorithm for *NDlog* called PSN^ν and prove
46 its correctness for the fragment of non-recursive programs. We formalize both
47 PSN^ν and SN algorithms as the search for proofs of the same linear logic [5]
48 theory extended with subexponentials [14]. Then, we show that a PSN^ν execu-
49 tion for a distributed *NDlog* program derives the same facts as an SN execution
50 for a centralized Datalog program. This property is proved by relating the linear
51 logic proofs specifying PSN^ν computation-runs with the proofs specifying SN
52 computation-runs. We also argue that the same reasoning is applicable to prov-
53 ing correctness of PSN^ν for recursive programs provided that PSN^ν terminates
54 in the presence of messages inserting and deleting the same tuple. Finally, we
55 identify several potential implementation improvements by using PSN^ν .

56 The rest of the paper is organized as follows. In Section 2, we review the
57 basics of *NDlog*, while in Section 3 we review the SN and PSN algorithms,
58 explain the problems of PSN, and informally introduce PSN^ν . Then, in Section
59 4, we sketch our encodings of SN and PSN^ν in linear logic and in Section 5 we
60 show our main correctness results. Finally in Section 6, we comment on related
61 work and conclude with final remarks in Section 7.

62 2 Network Datalog Language

63 In this section, we review the language *Network Datalog* (*NDlog*) [7], which ex-
64 tends Datalog programs, by allowing one to distribute Datalog rules in a network.

65 2.1 Background: Datalog

66 We first review some standard definitions of Datalog, following [16]. A *Datalog*
67 program consists of a (finite) set of logic rules and a query. A rule has the form

68 $\forall \mathbf{X}.(h \mathbf{T}_h \leftarrow b_1 \mathbf{T}_1, \dots, b_n \mathbf{T}_n)$, where the commas are interpreted as conjunctions and the symbol \leftarrow as implication; $h \mathbf{T}_h$ is an atom called the head of the rule; $b_1 \mathbf{T}_1, \dots, b_n \mathbf{T}_n$ is a sequence of atoms and function relations called the body; and the \mathbf{T} s are vectors of variables and ground terms. The variables in \mathbf{X} are exactly those appearing in the union of the variables in \mathbf{T}_h and \mathbf{T}_i s. *Function relations* are simple operations such as boolean, or arithmetic (e.g. $X_1 < X_2$), or list manipulations operations (e.g. $app\ L_1\ L_2\ L_3$). Semantically the order of the elements in the body does not matter, but it does have an impact on how programs are evaluated (usually from left to right). The query is a ground atom. We say that a predicate p depends on q if there is a rule where p appears in its head and q in its body. The *dependency graph* of a program is the transitive closure of the dependency relation using its rules. We say that a program is *(non)recursive* if there are (no) cycles in its dependency graph. As a technical convenience, we assume that if predicates have different arities, then they have different names³. We classify the predicates that do not depend on any other predicates as base predicates, and the remaining predicates as derived predicates. Consider the following non-recursive Datalog program where p, s , and t are a derived predicates and u, q , and r are base predicates: $\{p \leftarrow s, t, r; s \leftarrow q; t \leftarrow u; q \leftarrow; u \leftarrow\}$. The set of all the ground atoms that are derivable from this program, called *view*, is the multiset $\{s, t, q, u\}$.

88 Datalog's predicates (atoms) correspond to tuples in databases, and logical conjunction is equivalent to a join operation in database. For the rest of the paper, these terms are used interchangeably.

91 2.2 Network Datalog by Example

92 To illustrate *NDlog* program, we provide an example based on a simplified version of the *path-vector* protocol, a standard routing protocol used for paths between any two nodes in the network. This protocol is used as a basis for Internet routing today, where different *autonomous systems* (or *Internet Service Providers*) exchange routes using this protocol.

```
97 r1 path(@S,D,P,C) :- link(@S,D,C), P=f_init(S,D).
98 r2 path(@S,D,P,C) :- link(@S,Z,C1), path(@Z,D,P2,C2), C=C1+C2,
99 P=f_concat(S,P2), f_inPath(P2,S)=false.
```

100 The program takes as input `link(@S,D,C)` tuples, where each tuple represents an edge from the node itself (S) to one of its neighbors (D) of cost C . *NDlog* supports a *location specifier* in each predicate, expressed with “@” symbol followed by an attribute. This attribute is used to denote the source location of each corresponding tuple. For example, `link` tuples are stored based on the value of the S attribute.

106 Rules `r1-r2` recursively derive `path(@S,D,P,C)` tuples, where each tuple represents the fact that there is a path P from S to D with cost C . Rule `r1` computes one-hop reachability, given the neighbor set of S stored in `link(@S,D,C)`. Rule `r2` computes transitive reachability as follows: if there exists a link from S to Z with cost $C1$, and Z knows a path $P2$ to D with cost $C2$, then S can reach D via the path `f_concatPath(S,P2)` with cost $C1+C2$. Rules `r1-r2` utilize two list manipulation functions: `P=f_init(S,D)` initializes a path vector with two nodes S and D , while `f_concatPath(S,P2)` prepends S to path vector $P2$. To prevent computing paths

³ One can easily rewrite predicate names and distinguish them by using their arities.

114 with cycles, rule `r2` uses function `f_inPath`, where `f_inPath(P,S)` returns true if
 115 `S` is in the path vector `P`.

116 To implement the path-vector protocol in the network, each node runs the
 117 exact same copy of the above program, but only stores tuples relevant to its own
 118 state. What is interesting about this program is that predicates in the body of
 119 rule `r2` have different location specifiers indicating that they are stored on differ-
 120 ent node. To improve performance and eliminate unnecessary communication,
 121 we use a *rule localization* [7] rewrite procedure that transforms a program into
 122 an equivalent one where all elements in the body of a rule have the same loca-
 123 tion, but the head of the rule may reside at a different location than the body
 124 predicates. We call a rule non-local when the rule head and body have different
 125 location specifiers. We use the convention that a non-local rule resides in the
 126 same location as its body predicates, and that when the rule is used, the derived
 127 head predicate will be *sent* to the appropriate location as specified. For the rest
 128 of this paper, we assume that the localization rewrite has been performed.

129 3 Network Datalog Program Execution

130 The evaluation of *NDlog* programs uses *pipelined semi-naïve* (PSN) algorithm,
 131 which is based on *semi-naïve fixed point* [2] Datalog evaluation mechanism (SN).
 132 We provide a brief review of SN algorithm, before describing the PSN extension.

133 3.1 Semi-Naïve Algorithm

134 When base predicates are updated, these updates need to be propagated so that
 135 the views are consistent with the Datalog rules and current base predicate. Semi-
 136 naïve (SN) evaluation iteratively updates the view until a fixed point is reached.
 137 Tuples computed for the first time in the previous iteration are used as input in
 138 the current iteration; and new tuples that are generated for the first time in the
 139 current iteration are then used as input to the next iteration.

140 Given a set of insertions, I_k , and deletions, D_k of base predicates, the Algo-
 141 rithm 1 can be used to maintain the view of a Datalog program. First, we create
 142 for each rule $\forall \mathbf{X}.(h \mathbf{T}_h \leftarrow b_1 \mathbf{T}_1, \dots, b_n \mathbf{T}_n)$ in a Datalog program the following
 143 delta insertion and deletion rules:

$$\begin{aligned} & \{ \forall \mathbf{X}.(\text{INS}(h) \mathbf{T}_h \leftarrow b_1^\nu \mathbf{T}_1, \dots, b_{i-1}^\nu \mathbf{T}_{i-1}, \Delta b_i \mathbf{T}_i, b_{i+1} \mathbf{T}_{i+1}, \dots, b_n \mathbf{T}_n) \mid 1 \leq i \leq n \} \\ & \{ \forall \mathbf{X}.(\text{DEL}(h) \mathbf{T}_h \leftarrow b_1^\nu \mathbf{T}_1, \dots, b_{i-1}^\nu \mathbf{T}_{i-1}, \Delta b_i \mathbf{T}_i, b_{i+1} \mathbf{T}_{i+1}, \dots, b_n \mathbf{T}_n) \mid 1 \leq i \leq n \} \end{aligned}$$

144 Intuitively, given a set of insertions, I_k , and deletions, D_k , of base predicates,
 145 the Algorithm 1 uses these rules to incrementally maintain a view as follows: if
 146 we are in, say, the $i^{\text{th}} + 1$ iteration, then the contents of p corresponds to the
 147 view of p at iteration $i - 1$ and the contents of p^ν to the view at iteration i . The
 148 $i^{\text{th}} + 1$ iteration consists of executing the delta-rules for all updates in I_k and D_k ,
 149 and whenever an insertion or deletion rule is fired, we store the derived tuple in
 150 respectively I_k^ν and D_k^ν . Once all rules have been executed, we update the view
 151 accordingly and proceed to a new iteration, but now using the updates stored
 152 in I_k^ν and D_k^ν , which correspond to the updates derived in iteration $i^{\text{th}} + 1$. This
 153 is done by the last lines of the code which use *set-operations*.

154 Algorithm 1 maintains correctly the view of a Datalog program [6] when-
 155 ever there is one and only one derivation for any tuple. This limitation is due
 156 to the use of set semantics. Other more complicated algorithms are available,

Algorithm 1 SN-algorithm.

```
while  $\exists I_k.size > 0$  or  $\exists D_k.size > 0$  do
  while  $\exists I_k.size > 0$  or  $\exists D_k.size > 0$  do
     $\Delta t_k \leftarrow I_k.remove$  (resp.  $\Delta t_k \leftarrow D_k.remove$ )
     $I_k^{aux}.insert(\Delta t_k)$  (resp.  $D_k^{aux}.insert(\Delta t_k)$ )
    execute all insertions (resp. deletion) delta-rules for  $t_k$ :
       $\Delta p_k^{i+1} \leftarrow p_1^v, \dots, p_{i-1}^v, \Delta t_k, p_{k+1}, \dots, p_n$ 
    for all derived tuples  $p \in \Delta p_k^{i+1}$  do
       $I_k^v.insert(p)$  (resp.  $D_k^v.insert(p)$ )
    end for
  end while
  for all predicates  $p_j$  do
     $p_j \leftarrow (p_j \cup I_j^{aux}) \setminus D_j^{aux}$ ;  $p_j^v \leftarrow (p_j \cup I_j^v) \setminus D_j^v$ ;  $I_j \leftarrow I_j^v.flush$ ;  $D_j \leftarrow D_j^v.flush$ ;
     $D_j^{aux} \leftarrow \emptyset$ ;  $I_j^{aux} \leftarrow \emptyset$ ;  $\Delta p_j^{i+1} \leftarrow \emptyset$ 
  end for
end while
```

```
@1: {} []                {p} [INS(p)]                {p} [INS(p)]                {p} []
@2: {r, s, t} [INS(r)]   {r, s, t} []                {r} [DEL(s), DEL(t)]       {r} []
@3: {} [DEL(q)]         -- INS(r)--> {} [DEL(q)] -- DEL(q), DEL(u)--> {} [] -----* {} []
@4: {} [DEL(u)]         {} [DEL(u)]                {} []                       {} []
```

Fig. 1. PSN computation-run resulting in an incorrect final state. The i^{th} row depicts the evolution of the view, in curly-brackets, and the queue, in brackets, of node i . The updates in the arrows are the ones dequeued by PSN and used to update the view of the nodes. We also elide the @ in the predicates and updates.

157 but formalizing them seems to be a non-trivial task. Moreover, Algorithm 1
158 captures most of the programs used until now in declarative networking. For
159 instance, we can use it to maintain the datalog program corresponding to the
160 path vector program described above since each `path` tuple is supported by just
161 one derivation.

162 3.2 Existing Pipelined Semi-naïve Evaluation

163 In order to maintain incrementally the states of nodes in a distributed setting,
164 Loo *et al.* in [7, 8] proposed PSN. In PSN, each node has a queue of *messages*
165 scheduling insertions and deletions of tuples to the node's local state. A node
166 proceeds in a similar fashion as in Algorithm 1; it dequeues one update; then
167 executes its corresponding insertion or deletion delta-rules; and then for each
168 derived tuple, it sends a message which is to be stored at the end of the queue
169 of the node specified by derived tuple's location specifier (@). However, when a
170 message reaches a node, it is not only stored at the end of the node's queue, but
171 it is also immediately used to update the node's local state, that is, the tuple in
172 the message is immediately inserted into or deleted from the node's view.

173 We now demonstrate that updating a node's view by using messages before
174 they are dequeued can yield unsound results. Consider the following *NDlog*
175 program whose view is `{s@2, t@2, q@3, u@4}`:

```
176 p@1 :- s@2 t@2, r@2      s@2 :- q@3      t@2 :- u@4      q@3 :-      u@4 :-
```

177 Moreover, consider the PSN computation-run depicted in Figure 1 which
178 uses the messages inserting the tuple `r@2` and deleting the tuples `q@3` and `u@4`.

179 Notice that in the first state these updates have already been used to update the
 180 view of the nodes. In the final transitions, none of the updates deleting s and t
 181 trigger the deletion of p because the bodies of the respective deletion rules are
 182 not satisfied since t and u are no longer in node 2's view. Hence, the predicate p
 183 is entailed after PSN terminates although it is not supported by any derivation.

184 The second problem that we identify is that differently from SN, PSN does
 185 not avoid redundant computations. This is because in PSN a delta rule is fired by
 186 using the contents currently stored in a node's view, and not distinguishing, as in
 187 SN, its two previous states, which in SN is accomplished by using the predicates
 188 p and p' . For example, the *NDlog* rule $p@1 :- t@1$, $t@1$ would be rewritten into
 189 the following two insertion rules, where we elide the $@$ symbols: $\text{INS}(p) :- \Delta t$,
 190 t and $\text{INS}(p) :- t, \Delta t$. Thus if we dequeue an update inserting the tuple t ,
 191 both rules are fired, and two instances inserting p are added to node 1's queue.

192 Finally, the third problem that we identify is divergence. Consider the simple
 193 *NDlog* program composed of two rules: $p@1 :- a@1$ and $p@1 :- p@1$; and that the
 194 node's 1 queue is $[\text{INS}(a), \text{DEL}(a)]$. The insertion (resp. deletion) of a will cause
 195 an insertion (resp. deletion) of p to be added at the end of the queue. Because of
 196 the second rule, the insertion and deletion of p will propagate indefinitely many
 197 insertions and deletions of p and therefore causing PSN to diverge.

198 In the informal description of PSN, presented in [7, 8], many assumptions
 199 were used, such as that messages are not lost; a *Bursty Model*, that is, the
 200 network eventually *quiesces* (does not change) for a time long enough to all the
 201 system to reach a fixed point; that message channels are assumed to be FIFO,
 202 hence no reordering of messages is allowed; and that timestamps are attached
 203 to tuples in order to evaluate delta rules. Even under these strong assumptions,
 204 the problems in PSN mentioned above persist. What is more troublesome is that
 205 this design is reflected in the current implementation of *NDlog* and therefore, all
 206 *NDlog* programs exhibit those flaws.

207 In the next section, we propose a new evaluation algorithm, called PSN' ,
 208 which not only corrects these problems, but also does not require the last two
 209 assumptions (FIFO channels and use of timestamps). The removal of these two
 210 assumptions not only simplifies the implementation, it also potentially leads to
 211 improved performance, since the implementation no longer requires receiver-
 212 based network buffers necessary to guarantee in-order delivery of messages.

213 3.3 New Pipelined Semi-naïve Evaluation

214 At a high-level, PSN' works as follows: Instead of using queues to store unpro-
 215 cessed updates, we use a single *bag*, called *upd*, that specifies the asynchronous
 216 behavior in the distributed setting by abstracting the order in which updates
 217 are used. Thus in this abstraction, we do not need to take into account the $@$
 218 specifiers since all messages go to *upd*. We process *NDlog* rules into delta-rules
 219 exactly as in the SN algorithm, so that the multiple derivation problem does
 220 not occur. Then, one PSN' -iteration is completed by executing in a sequence
 221 the following three basic commands, with the invariant that before and after a
 222 PSN' -iteration, the contents in p and in p' are the same:

223 **pick** – One picks (non-deterministically) any update, u , from the bag *upd*, ex-
 224 cept if the u is a deletion of an atom that is not (yet) in the view. Then, if u is
 225 an insertion of predicate p , we add it to the contents of p' , otherwise if it is a

226 deletion of the same predicate, we delete it from p^ν ;
 227 **fire** – After picking an update, one executes all the delta-rules corresponding to
 228 u . If a rule is fired, then we insert the derived tuple into the bag upd .
 229 **update** – Once all delta-rules are executed, we update the view according to u :
 230 if u is an insertion or deletion of predicate p , we insert it into or delete it from
 231 the contents of p .

232 The execution of an SN-iteration can also be specified with the use of the
 233 same three basic commands above. However, instead of applying just one se-
 234 quence of the three commands, the $i^{th} + 1$ SN-iteration is composed of three
 235 phases: first, all elements in upd are picked using the *pick* command. The result
 236 is that the contents in the p^ν s are updated with the updates derived in the pre-
 237 vious iteration. Hence, the contents of the p^ν s correspond exactly to the view at
 238 iteration i , while the contents in p correspond exactly to the view at iteration
 239 $i - 1$, as in Algorithm 1. Then one executes the delta-rules for all updates picked
 240 in the previous phase, deriving and storing new updates in the bag upd . After
 241 this phase, upd contains the updates derived at iteration $i + 1$. Finally, in the
 242 third phase, one executes eagerly the *update* command which then updates the
 243 contents in p to match the contents in p^ν .

244 Because both algorithms can be explained by using the same basic commands
 245 and the same delta-rules, we are able to prove correctness of PSN^ν by showing
 246 that for any computation-run of PSN^ν , which formally corresponds to a linear
 247 logic proof, there is a computation-run of SN, which corresponds to another
 248 linear logic proof of the same sequent, and vice-versa.

249 4 Encoding PSN^ν and SN in Linear Logic with 250 Subexponentials

251 We choose to use linear logic to specify the operational semantics of PSN^ν or
 252 of SN instead of a transition system, because of the following two reasons. First,
 253 linear logic is a precise and well established language, used already for both
 254 reasoning and specifying semantics of programming languages. Second, linear
 255 logic provides us with a finer detail on how data is manipulated, thus opening
 256 the possibility to use our encoding to prove the correctness not only of PSN^ν ,
 257 but also of how it is implemented.

258 Although the details of the proof system for linear logic with subexponentials
 259 are beyond the scope of this paper, in the next sections, we sketch its role for
 260 the specification of both algorithms PSN^ν and SN. The details of the encoding
 261 can be found in [13].

262 4.1 Linear Logic and Subexponentials

263 We review some of linear logic's basic proof theory. *Literals* are either atoms or
 264 their negations. The connectives \otimes and \wp and the units 1 and \perp are *multiplica-*
 265 *tive*; the connectives $\&$ and \oplus and the units \top and 0 are *additive*; \forall and \exists are
 266 (first-order) quantifiers; and $!$ and $?$ are the *exponentials*. We assume that all
 267 formulas are in *negation normal form*, that is, negation has atomic scope.

268 Due to the exponentials, one can distinguish in linear logic two kinds of
 269 formulas: the linear ones whose main connective is not a $?$ and the unbounded
 270 ones whose main connective is a $?$. The linear formulas can be seen as resources
 271 that can only be used once, while the unbounded formulas as unlimited resources

272 which can be used as many times necessary. This distinction is usually reflected
 273 in syntax by using two different contexts in the sequent, one containing only
 274 unbounded formulas and another only linear formulas [1]. Such distinction allows
 275 one to incorporate structural rules, *i.e.*, weakening and contraction, into the
 276 introduction rules of connectives.

277 However, the exponentials are not canonical [3]. In fact, we can assume the
 278 existence of a proof system containing as many exponential-like operators, ($!^l, ?^l$)
 279 called subexponentials [14], as one needs: they may or may not allow contraction
 280 and weakening, and are organized in a pre-order (\preceq) specifying the entailment
 281 relation between operators. In these proof systems the contexts for the subex-
 282 ponentials are denoted by the function \mathcal{K} which maps the set of *subexponential*
 283 *indexes* to multisets of formulas. If l is a subexponential index, we denote by $\mathcal{K}[l]$
 284 the multiset of formulas associated to l by \mathcal{K} . Notice that a context $\mathcal{K}[l]$ behaves
 285 either like the linear logic’s unbounded context or its linear context depending if
 286 the index l allows structural rules or not. The preorder \preceq is used to specify the
 287 introduction rule of subexponential bangs. As in its corresponding linear logic
 288 rule, to introduce a $!^l$ one needs to check if some type of formulas are not present,
 289 namely, that there are no formulas in the linear context nor in the contexts of
 290 the indexes k such that $l \not\preceq k$.

291 Following [14], we use subexponential indexes to encode data structures, such
 292 as views, in the context of a sequent. Given a set of ground atoms \mathcal{D} , representing
 293 a view, for each predicate p , we store its view with respect to \mathcal{D} in the contexts of
 294 the subexponentials p and p^ν using the functions: $\mathcal{K}_{\mathcal{D}}[p] = \{p[\mathbf{t}] \mid p\mathbf{t} \in \mathcal{D}\}$ and
 295 $\mathcal{K}_{\mathcal{D}}[p^\nu] = \{p^\nu[\mathbf{t}] \mid p\mathbf{t} \in \mathcal{D}\}$, where $[\mathbf{t}]$ is a list of terms. We encode in a similar
 296 fashion updates using the index *upd*, the query using the function *query*, and
 297 the encoding of program delta-rules using the index *rules*. In order to keep track
 298 of which updates have been used to fire rules from those that have not, we use
 299 the indexes *picked*, where we store updates that were picked from the *upd* bag,
 300 and *exec*, where we store updates that have been used to fire delta-rules.

301 To check if the contexts of the indexes in the set \mathcal{I} are all empty, we follow
 302 [14] and create a new index \hat{l} such that $\hat{l} \preceq k$ for all indexes, except those in \mathcal{I} .
 303 Therefore one can only introduce the subexponential bang of \hat{l} if the contexts
 304 for the indexes in \mathcal{I} are all empty.

305 4.2 Focusing and algorithmic specifications

306 Focused proof systems, first introduced by Andreoli for linear logic [1], provide
 307 normal-form proofs for proof search. Inference rules that are not necessarily
 308 invertible are classified as positive, and the remaining rules as negative. Using
 309 this classification, focused proof systems reduce proof search space by allowing
 310 one to combine a sequence of introduction rules of the same polarity into larger
 311 derivations, which can be seen as “macro-rules”. The backchaining rule in logic
 312 programming can be seen as such macro-rule.

313 In [14], Nigam and Miller propose the focused system for linear logic with
 314 subexponentials called SELLF and show how to specify imperative-like pro-
 315 grams. Consider for example the linear logic definitions depicted in Figure 2. In
 316 a focused system, these definitions are enforced to behave exactly as one would
 317 intuitively imagine. The instructions **load** and **unload** insert and delete an el-
 318 ement from a context, while **end** is just used to mark the end of a program.

$$\begin{array}{lcl}
\mathbf{load} \langle t_1, \dots, t_n \rangle l \text{ prog} & \triangleq & ?^l (l t_1 \dots t_n) \wp \text{ prog} \\
\mathbf{unload} l \langle v_1, \dots, v_n \rangle b\text{prog} & \triangleq & (l v_1 \dots v_n)^\perp \otimes (b\text{prog } v_1 \dots v_n) \\
\mathbf{loop} l \text{ kprog } \text{prog} & \triangleq & \exists v_1 \dots v_n [(l v_1 \dots v_n)^\perp \otimes \\
& & (k\text{prog } v_1 \dots v_n) (\mathbf{loop} l \text{ kprog } \text{prog})] \oplus !^i(\text{prog}) \\
\mathbf{end} & \triangleq & \perp
\end{array}$$

Fig. 2. Linear logic definitions for the basic instructions.

$$\begin{array}{lcl}
\mathit{pick} & \triangleq & \exists PLU[\mathbf{unload} \text{ upd } \langle P, L, U \rangle; \mathbf{load} \langle P, L, U \rangle \mathit{picked} \\
& & [(U = \text{INS}) \otimes \mathbf{load} \langle L \rangle P^\nu \mathbf{end}] \oplus [(U = \text{DEL}) \otimes \mathbf{unload} \langle L \rangle P^\nu \mathbf{end}]] \\
\mathit{fire} & \triangleq & \exists PLUR[\mathbf{unload} \mathit{picked} \langle P, L, U \rangle; \mathbf{unload} \text{ rules } \langle P, R, U \rangle; \\
& & \mathbf{load} \langle P, L, U \rangle \text{ exec}; \mathbf{load} \langle L \rangle \Delta P; \text{execRules } R(\mathbf{unload} \Delta P \langle L \rangle \mathbf{end})] \\
\mathit{update} & \triangleq & \exists PLU[\mathbf{unload} \text{ exec } \langle P, L, U \rangle \\
& & [(U = \text{INS}) \otimes \mathbf{load} \langle L \rangle P \mathbf{end}] \oplus [(U = \text{DEL}) \otimes \mathbf{unload} P \langle L \rangle \mathbf{end}]] \\
\mathit{query} & \triangleq & !^{\text{test}} \exists SL[\mathbf{unload} \text{ queryLoc } \langle S, L \rangle (\mathbf{unload} \langle L \rangle S \top)]
\end{array}$$

Fig. 3. Linear logic definitions specifying the basic commands. We elide from specifications the λ symbols and denote formulas of the form $A (B C)$ as $(A; B C)$.

319 In $\mathbf{loop} l \text{ kprog } \text{prog}$, we use a continuation passing style specification. It deletes
320 an atom from the context of l and focuses on the logic formula obtained from
321 applying the terms $v_1 \dots v_n$ and the continuation $(\mathbf{loop} l \text{ kprog } \text{prog})$ to $k\text{prog}$.
322 The loop ends when the context of l is empty, specified by the use of the $!^i$, and
323 then continues by introducing the logic formula prog .

324 The definition $\mathit{move} S R K \triangleq \mathbf{loop} S \lambda T \lambda \text{cont}_l (\mathbf{load} \langle T \rangle R \text{cont}_l) K$ illus-
325 trates the use of these definitions. It moves all the elements from the context S
326 to the context R , and then proceeds with the logic formula K .

327 4.3 Basic Commands

328 The linear logic definition for the basic commands described informally in Section
329 3 are depicted Figure 3. The basic command fire is the most elaborate. It starts
330 by unloading an update, $\langle p, l, u \rangle$, that is in picked , where p is predicate name,
331 l a list of terms denoting its arguments, and u is either INS or DEL denoting
332 the type of update; then retrieving the corresponding insertion or deletion delta
333 rules r , for the predicate p ; loading and unloading l into Δp , in order to execute
334 its delta rules; and finally loading the tuple $\langle p, l, u \rangle$ in the context exec , denoting
335 that the delta rules for this update have been executed. The execution of a rule
336 is done by execRules , whose definition can be found in the technical report [13].
337 Intuitively, one traverses the encoding of the body of a rule building in the process
338 a substitution that satisfies all body elements. If a predicate is encountered, one
339 checks among all elements in its view for the ones that can be used to fire the
340 rule; otherwise if a function relation is encountered, one checks if the partial
341 substitution built satisfies the relation. Once a rule is fired, we insert the derived
342 update in upd . Notice that query is the only command that can finish a proof
343 due to the presence of \top which is reached only after verifying that the query is
344 in the view. The $!^{\text{test}}$ specifies that query can only be used when the contexts for

345 *upd*, *picked*, and *exec* are all empty and therefore there are more updates being
 346 processed.

347 We insert these basic commands in a sequent by using the function $\mathcal{K}_{BC}[\infty] =$
 348 $\{!^{-\infty} \textit{pick}, !^{-\infty} \textit{fire}, !^{-\infty} \textit{update}, !^{-\infty} \textit{query}\}$, where ∞ ($-\infty$) is the maximal (min-
 349 imal) index, that is, $l \preceq \infty$ ($-\infty \preceq l$) for all index l . Since the maximal index
 350 allows both contraction and weakening, the basic commands can be used as many
 351 times as needed. The purpose of the minimal index is novel. Due to the focusing
 352 discipline, the execution of a basic command is atomic, that is, one can only use
 353 a basic command when there is no other basic command being introduced.

354 Given a set of ground atoms \mathcal{D} , a Datalog program \mathcal{P} , a multiset of updates
 355 \mathcal{U} , and a ground atom s , the sequent $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ is such that its linear context
 356 is empty and its subexponential context is $\mathcal{K}_{\mathcal{D}} \otimes \mathcal{K}_{\mathcal{P}} \otimes \mathcal{K}_{\mathcal{U}} \otimes \mathcal{K}_s \otimes \mathcal{K}_{BC}$, where
 357 \mathcal{K}_{BC} is the encoding of basic commands, \mathcal{K}_s is the encoding of the query for s ,
 358 $\mathcal{K}_{\mathcal{U}}$ is the encoding of updates, $\mathcal{K}_{\mathcal{P}}$ the encoding of delta-rules, $\mathcal{K}_{\mathcal{D}}$ the encoding
 359 of the view, and $\mathcal{K}_1 \otimes \mathcal{K}_2[l] = \mathcal{K}_1[l] \cup \mathcal{K}_2[l]$ for any l .

360 5 Correctness

361 The following definitions specify the proofs that correspond to computation runs
 362 of PSN^ν and of SN, called respectively PSN^ν and SN-proofs. The correctness
 363 proof goes by showing that if one proof exists then the other must also exist; or
 364 in other words, any query that is entailed by using PSN^ν is also entailed by SN
 365 and vice-versa.

366 **Definition 1.** *An execution of a basic command BC is any focused derivation*
 367 *that introduces a sequent focused on the formula $!^{-\infty} BC$ and whose rules intro-*
 368 *duce only descendants of $!^{-\infty} BC$. We say that the execution of pick (resp. fire*
 369 *and update) uses u if u is the element unloaded from *upd (resp. picked and*
 370 *exec).**

371 **Definition 2.** *A derivation is a complete iteration if it can be partitioned into a*
 372 *sequence of executions of pick, followed by a sequence of executions of fire, and*
 373 *finally a sequence of executions of update, such that the multiset of tuples, \mathcal{T} ,*
 374 *used by the sequence of pick executions is the same as used by the sequence of fire*
 375 *and update executions. A complete iteration is an SN-iteration if \mathcal{T} contains all*
 376 *tuples at the end-sequent that are in $\mathcal{K}[\textit{upd}]$. A complete iteration is a PSN^ν -*
 377 *iteration if \mathcal{T} contains only one element.*

378 **Definition 3.** *Let \mathcal{D} be a set of ground atoms, \mathcal{P} be a Datalog program, \mathcal{U} a*
 379 *multiset of updates, and s be a ground atom. We call any focused proof, Ξ ,*
 380 *of the sequent $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ as a PSN^ν -proof (respectively SN-proof) if it can*
 381 *be partitioned into a sequence of PSN^ν -iterations (respectively SN-iterations)*
 382 *followed by an execution of query.*

383 **Theorem 1.** *Let \mathcal{D} be a set of ground atoms, \mathcal{P} be a non-recursive Datalog*
 384 *program, \mathcal{U} be a multiset of updates, and s be a ground atom. There is a PSN^ν -*
 385 *proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ iff there is an SN-proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$.*

386 **Corollary 1.** *For non-recursive programs, a query is entailed by using PSN^ν*
 387 *iff it is entailed by using SN.*

388 We prove the theorem above by showing that: 1) we can permute the ex-
389 ecutions of two PSN^ν -iterations; 2) we can merge a complete-iteration and
390 a PSN^ν -iteration into a larger complete-iteration; and 3) conversely we can
391 split a larger complete-iteration into a smaller complete-iteration and a PSN^ν -
392 iteration. These operations are formalized by the Lemmas 2 and 3 shown in
393 the Appendix. Given a PSN^ν -proof, we construct an SN-proof by induction as
394 follows: we use the first operation to permute downwards the PSN^ν -iteration
395 that picks any element in the end-sequent's *upd*'s context, then repeat it with
396 its subproof. The resulting proof has all PSN^ν -iterations in the same order as in
397 an SN-Proof. We merge them into SN-iterations by applying the second opera-
398 tion repeatedly. For the converse direction, given an SN-proof, we can repeatedly
399 apply the third operation to split SN-iterations and obtain a PSN^ν -proof.

400 While performing these operations, however, it can happen that new rules
401 are fired. In particular, when we permute a PSN^ν -iteration that uses a deletion
402 update over a PSN^ν -iteration that uses an insertion update. The updates gener-
403 ated in these cases are necessarily conflicting, that is, are pairs of insertions and
404 deletions of the same tuple. In the general case, we cannot guarantee that PSN^ν
405 terminates when processing such conflicting updates, but we can guarantee its
406 termination if the program is non-recursive since these programs do not contain
407 dependency cycles and therefore the propagation of updates must end. This is
408 formalized by Lemma 1 in the Appendix.

409 However, if we can guarantee such termination for PSN^ν , then the proof
410 works exactly in the same way. Let us return to our path-vector example, shown
411 in Section 2, which is a recursive program. Because of the use of the function
412 `f_inPath`, one does not compute paths that contain cycles. This restriction alone
413 is enough to guarantee termination of PSN^ν : the number of `path`-updates prop-
414 agated by conflicting updates inserting and a deleting the same `link` tuple is
415 finite. Therefore we can use the same reasoning above to show that PSN^ν is
416 correct for this program.

417 In literature, there are algorithms that can be used to determine termina-
418 tion of Datalog programs [11]. It seems possible to adapt them to a distributed
419 setting, but this is left out of the scope of this paper. We are also currently
420 investigating larger classes of programs for which PSN^ν terminates.

421 6 Related Work

422 Navarro *et al.* propose in [12] an operational semantics for a variation of the
423 *NDlog* language that also includes rules with events. However, their semantics
424 also computes unsound results and therefore it is not suitable as an operational
425 semantics for *NDlog*. For instance, besides the problems we identify for PSN, one
426 is also allowed in their work to pick an update that deletes an element without
427 checking if this element is present in the view, which also yields unsound results.

428 7 Conclusions

429 In this paper, we have developed a new PSN algorithm, PSN^ν , which is key to
430 specifying the operational semantics of *NDlog* programs. We have proven that
431 PSN^ν is correct with regard to the centralized SN by using a novel approach:
432 we encode both the SN and PSN^ν in linear logic with subexponentials. The

433 correctness result is proven by showing that a proof that encodes a SN eval-
434 uation can be transformed to one that encodes a PSN^ν evaluation and vice
435 versa. Focused proofs in linear logic give well-defined operational semantics for
436 PSN^ν . Furthermore, PSN^ν lifts restrictions such as FIFO channels from $NDlog$
437 implementations and leads to significant performance improvements of protocol
438 execution.

439 This work is part of a bigger effort to formally analyze network protocol im-
440 plementations [4, 19]. The results in this paper lay a solid foundation toward clos-
441 ing the gap between verification and implementation. An important part of our
442 future work is to formalize low-level $NDlog$ implementations so that verification
443 results on high-level specifications can be applied to low-level implementations.

444 References

- 445 1. J. Andreoli. Logic programming with focusing proofs in linear logic. *Journal of*
446 *Logic and Computation*, 2(3):297–347, 1992.
- 447 2. I. Balbin and K. Ramamohanarao. A Generalization of the Differential Approach
448 to Recursive Query Evaluation. *Journal of Logic Prog*, 4(3):259–262, 1987.
- 449 3. V. Danos, J. Joinet, and H. Schellinx. The structure of exponentials: Uncovering
450 the dynamics of linear logic proofs. In *Kurt Gödel Colloquium*, p. 159–171, 1993.
- 451 4. Formally Verifiable Networking. <http://netdb.cis.upenn.edu/fvn/>
- 452 5. J.Y. Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- 453 6. A. Gupta, I. Mumick, and V. S. Subrahmanian. Maintaining views incrementally.
454 In *SIGMOD*, pages 157–166, 1993.
- 455 7. B. T. Loo, T. Condie, M. Garofalakis, D. E. Gay, J. M. Hellerstein, P. Maniatis,
456 R. Ramakrishnan, T. Roscoe, and I. Stoica. Declarative Networking: Language,
457 Execution and Optimization. In *SIGMOD*, 2006.
- 458 8. B. T. Loo, Tyson Condie, Minos Garofalakis, David E. Gay, Joseph M. Heller-
459 stein, Petros Maniatis, Raghu Ramakrishnan, Timothy Roscoe, and Ion Stoica.
460 Declarative Networking. In *Communications of the ACM (CACM)*, 2009.
- 461 9. B. T. Loo, T. Condie, J.M. Hellerstein, P. Maniatis, T. Roscoe, and I. Stoica.
462 Implementing Declarative Overlays. In *SOSP*, 2005.
- 463 10. B. T. Loo, J. M. Hellerstein, I. Stoica, and R. Ramakrishnan. Declarative Routing:
464 Extensible Routing with Declarative Queries. In *SIGCOMM*, 2005.
- 465 11. I. S. Mumick and O. Shmueli. Finiteness properties of database queries. In *Aus-*
466 *tralian Database Conference*, pages 274–288, 1993.
- 467 12. J. A. Navarro and A. Rybalchenko. Operational semantics for declarative network-
468 ing. In *PADL*, pages 76–90, 2009.
- 469 13. V. Nigam, L. Jia, A. Wang, B. T. Loo, and A. Scedrov. An operational semantics
470 for network datalog. Extended version available at
471 <http://netdb.cis.upenn.edu/fvn/ndlogsemantics.pdf>
- 472 14. V. Nigam and D. Miller. Algorithmic specifications in linear logic with subexpo-
473 nentials. In *PPDP*, pages 129–140, 2009.
- 474 15. P2: Declarative Networking System. <http://p2.cs.berkeley.edu>
- 475 16. R. Ramakrishnan and J. D. Ullman. A Survey of Research on Deductive Database
476 Systems. *Journal of Logic Programming*, 23(2):125–149, 1993.
- 477 17. RapidNet: A Declarative Toolkit for Rapid Network Simulation and Experimentation.
478 <http://netdb.cis.upenn.edu/rapidnet/>
- 479 18. I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A
480 Scalable P2P Lookup Service for Internet Applications. In *SIGCOMM*, 2001.
- 481 19. A. Wang, L. Jia, C. Liu, B. T. Loo, O. Sokolsky, and P. Basu. Formally Verifiable
482 Networking. In *SIGCOMM HotNets-VIII*, 2009.

483 8 Appendix

484 **Lemma 1.** *Let \mathcal{D} be a set of ground atoms, \mathcal{P} be a non-recursive Datalog pro-*
 485 *gram, s be a ground atom, and \mathcal{U} be a multiset of updates, such that $\langle p, L, \text{INS} \rangle,$*
 486 *$\langle p, L, \text{DEL} \rangle \in \mathcal{U}$. Let $\mathcal{U}' = \mathcal{U} \setminus \{\langle p, L, \text{INS} \rangle, \langle p, L, \text{DEL} \rangle\}$ be a multiset of updates.*
 487 *Then the sequent $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ has a PSN^ν -proof iff the sequent $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}', s)$*
 488 *has a PSN^ν -proof.*

489 *Proof.* (\Rightarrow) The updates $\langle p, L, \text{INS} \rangle, \langle p, L, \text{DEL} \rangle \in \mathcal{U}$ do not really affect the ex-
 490 ecution of *query*, since for all insertions propagated by the update $\langle p, L, \text{INS} \rangle$
 491 there are the same deletions propagated by the update $\langle p, L, \text{DEL} \rangle$. We can con-
 492 struct the a proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}', s)$ by trimming the pieces of derivations in the
 493 proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ that depend on these updates. We do so by induction
 494 on the number of PSN^ν -iterations. Let Ψ be the set of updates propagated by
 495 $\langle p, L, \text{INS} \rangle$ and $\langle p, L, \text{DEL} \rangle$. One determines this set by inspection on the proof
 496 of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$. Consider the following representative inductive case where the
 497 proof ends with a PSN^ν -iteration of the form:

$$\begin{array}{c}
 \Xi \\
 \hline
 \vdash \mathcal{K}'_2 : \cdot \uparrow \cdot \\
 \hline
 \vdash \mathcal{K}'_2 : \cdot \downarrow \mathbf{end} \\
 \hline
 \vdash \mathcal{K}_1 : \cdot \downarrow (\text{upd } p_1 L_1 u)^\perp \quad \vdash \mathcal{K}_2 : \cdot \downarrow \text{prog} \\
 \hline
 \vdash \mathcal{K} : \cdot \downarrow (\text{upd } p_1 L_1 u)^\perp \otimes \text{prog} \\
 \hline
 \vdash \mathcal{K} : \cdot \downarrow \mathbf{unload} \text{ upd } \langle p_1, L_1, u \rangle \text{ prog} \\
 \hline
 \vdash \mathcal{K} : \cdot \downarrow !^{-\infty} \text{pick}
 \end{array}$$

498 If the update $\langle p_1, L_1, u \rangle$ is an update propagated from $\langle p, L, \text{INS} \rangle$ or $\langle p, L, \text{DEL} \rangle$,
 499 then this derivation is completely deleted. Otherwise, we should not delete the
 500 whole derivation, but only the parts in the execution of *fire* that use tuples in
 501 the view which come from insertions propagated from $\langle p, L, \text{INS} \rangle$. These deletions
 502 are also done by induction, but this time on the number of “loops” in *fire*.

503 Here is a representative inductive case, where in the derivation below the
 504 **loops** are two consecutive occurrences of loops over p_1 :

$$\begin{array}{c}
 \Xi \\
 \hline
 \vdash \mathcal{K}'_2 : \downarrow \mathbf{loop} p_1 \text{ kprog } \text{prog}_2 \\
 \hline
 \vdash \mathcal{K}_1 : \downarrow (p_1 \mathbf{t})^\perp \quad \vdash \mathcal{K}_2 : \downarrow (\text{kprog } \mathbf{t}) (\mathbf{loop} p_1 \text{ kprog } \text{prog}) \\
 \hline
 \vdash \mathcal{K} : \downarrow (p_1 \mathbf{t})^\perp \otimes (\text{kprog } \mathbf{t}) (\mathbf{loop} p_1 \text{ kprog } \text{prog}) \\
 \hline
 \vdash \mathcal{K} : \downarrow \mathbf{loop} p_1 \text{ kprog } \text{prog}
 \end{array}$$

505 We delete this derivation only if p_1 is of the forms p or p^ν or p_{aux}^i ⁴ and the
 506 update $\langle p, [\mathbf{t}], \text{INS} \rangle$ is in Ψ . At the same time, we delete all occurrences of the
 507 atoms $(\text{upd } p l u)$, $(p l)$, $(p^\nu l)$, and $(p_{aux} l)$ such that the update $\langle p, l, u \rangle$ is in Ψ .

⁴ As you can see in the technical report, we assume that for each predicate p there are auxiliary subexponential indexes, p_{aux}^i , used to mark the tuples in p which were already traversed.

508 (\Leftarrow) Let Ξ be the given proof of the sequent $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}', s)$. Moreover, let Ξ_p be
509 the derivation composed of all PSN^ν -iterations in Ξ and Ξ_q be the derivation
510 composed of the *query* execution in Ξ . We can construct a proof of the sequent
511 $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ as follows. We add to the context *upd* of all sequents in Ξ_p that are
512 not introduced by an initial rule the updates $\langle p, L, \text{INS} \rangle$ and $\langle p, L, \text{DEL} \rangle$. Let Ξ'_p
513 be the resulting derivation. Then the end sequent of Ξ'_p is $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ and its
514 open premise is such that the context of *upd* is composed exactly of the updates
515 $\langle p, L, \text{INS} \rangle$ and $\langle p, L, \text{DEL} \rangle$. Now, since the program is non-recursive, it is case
516 that there is a finite sequence of PSN^ν -iterations that computes the updates
517 $\langle p, L, \text{INS} \rangle, \langle p, L, \text{DEL} \rangle$ and all the updates propagated by them. Let Ξ_u be the
518 derivation corresponding to such computation⁵. The context of *upd* of Ξ_u 's end
519 sequent is the multiset $\{\langle p, L, \text{INS} \rangle, \langle p, L, \text{DEL} \rangle\}$, while the same context for its
520 premise is the \emptyset . Finally, we can compose the derivations Ξ'_p, Ξ_u , and Ξ_q and
521 construct the proof for $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$.

522 **Lemma 2.** *Let \mathcal{D} be a set of ground atoms, \mathcal{P} be a non-recursive Datalog pro-*
523 *gram, \mathcal{U} be a multiset of updates, such that $u_1, u_2 \in \mathcal{U}$, and s be a ground atom.*
524 *Let Ξ be a PSN^ν -proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ which ends with two PSN^ν -iterations*
525 *that use u_1 and u_2 . Then there is a PSN^ν -proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ which ends with*
526 *two PSN^ν -iterations that use the updates u_2 and u_1 .*

527 *Proof.* We must consider four different cases, according to the updates u_1 and
528 u_2 :

- 529 • u_1 and u_2 are both insertions: $\langle p_1, L_1, \text{INS} \rangle$ and $\langle p_2, L_2, \text{INS} \rangle$. We show that the
530 multiset of firings obtained by first picking $\langle p_2, L_2, \text{INS} \rangle$ and then $\langle p_1, L_1, \text{INS} \rangle$ is
531 the same as before. Let F_1 be the multiset of firings in the first case and F_2 be
532 the set of firings in the second case. Let $s_1 \in F_1$. If s_1 be a firing obtained in the
533 first PSN^ν -iterations, then it must be the case that $s_1 \in F_2$ since the same delta
534 rule is executed. If s_1 is obtained in the second PSN^ν -iteration, then either it
535 did not use the insertion of $\langle p_1, L_1, \text{INS} \rangle$, in which case, $s_1 \in F_2$, since the same
536 delta-rule would be executed; or it did use the insertion of $\langle p_1, L_1, \text{INS} \rangle$, in which
537 case there is a rule that contains both p_1 and p_2 in the body, and therefore
538 $s_1 \in F_2$ because then its delta rule containing Δp_1 and t in its body is fired. To
539 prove that if $s_2 \in F_2$ then $s_2 \in F_1$ follows the same reasoning.
- 540 • u_1 and u_2 are both deletions: $\langle p_1, L_1, \text{DEL} \rangle$ and $\langle p_2, L_2, \text{DEL} \rangle$. The reasoning is
541 similar as in the previous case. Let F_1 be the multiset of firings in the first case
542 and F_2 be the set of firings in the second case.
- 543 • u_1 is an insertion and u_2 is a deletion: $\langle p_1, L_1, \text{INS} \rangle$ and $\langle p_2, L_2, \text{DEL} \rangle$. Again, we
544 show that the multiset of firings obtained by first picking $\langle p_2, L_2, \text{DEL} \rangle$ and then
545 $\langle p_1, L_1, \text{INS} \rangle$ is the same as before. Let F_1 be the multiset of firings in the first
546 case and F_2 be the set of firings in the second case. Let $s_1 = \langle s, L_s, \text{INS} \rangle \in F_1$ be
547 an update created in the first PSN^ν -iteration. Then either one did not use L_2

⁵ We can search for such computation by just following the algorithm specified in
linear logic. We do so by picking any INS update and then the corresponding DEL
update. Since in the execution of *fire* we traverse all possible combinations of tuples
in the view, it does not really matter in which order we unload elements. Hence, one
does not require to backtrack between focusing phases, but just to backtrack inside
focusing phases, which is controlled by the size of the “macro-rules”.

548 from p_2 , in which case, $s_1 \in F_2$, or one did use L_2 from p_2 , in which case it must
549 be that another update $s'_1 = \langle s, L_s, \text{DEL} \rangle \in F_2$ is created because a delta rule of
550 the same rule must be fired in the second PSN^ν -iteration. In this case, neither
551 s_1 nor s'_1 belong to F_2 because, by inverting the order of picks, no rule is fired.
552 However, from Lemma 1, the resulting sequent is still provable. The reasoning is
553 the same for the case when $s_1 = \langle s, L_s, \text{DEL} \rangle \in F_1$. To show the reverse direction
554 that if $s_2 \in F_2$ then $s_2 \in F_1$, the reasoning is similar to the next case.

555 • u_1 is a deletion and u_2 is an insertion: $\langle p_1, L_1, \text{DEL} \rangle$ and $\langle p_2, L_2, \text{INS} \rangle$. Once
556 more, we show that the multiset of firings obtained by first picking $\langle p_2, L_2, \text{INS} \rangle$
557 and then $\langle p_1, L_1, \text{DEL} \rangle$ is the same as before. Let F_1 be the multiset of firings in
558 the first case and F_2 be the set of firings in the second case. Let $s_1 \in F_1$, then
559 $s_1 \in F_2$ since the same delta rule must be fired when one picks u_2 before u_1 . Now,
560 consider that $s_2 = \langle s, L_s, \text{INS} \rangle \in F_2$ is created in the first PSN^ν -iteration. Then
561 it is created either not using L_2 from p_2 , in which case $s_2 \in F_1$, or by using L_2
562 from p_2 , in which case, it must be that another update $s'_2 = \langle s, L_s, \text{DEL} \rangle \in F_2$ is
563 created because a delta rule of the same rule must be fired in the second PSN^ν -
564 iteration. So $s_2, s'_2 \notin F_1$. However, again from Lemma 1, the resulting sequent
565 is still provable. The reasoning is the same for when $s_2 = \langle s, L_s, \text{DEL} \rangle \in F_2$.

566 **Lemma 3.** *Let \mathcal{D} be a set of ground atoms, \mathcal{P} be a non-recursive Datalog pro-*
567 *gram, \mathcal{U} be a multiset of updates, such that $\{u\} \cup \mathcal{T} \subseteq \mathcal{U}$, and s be a ground atom.*
568 *Then there is a proof of the sequent $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ which ends with a complete-*
569 *iteration that uses the multiset \mathcal{T} followed by a PSN^ν -iteration that uses the*
570 *update u iff there is a proof of the same sequent that ends with a complete-*
571 *iteration that uses the multiset $\mathcal{T} \cup \{u\}$.*

572 *Proof.* For each direction there are two cases according to the update u to con-
573 sider. Let F_1 be the multiset of updates created by a complete-iteration, C_1 ,
574 using \mathcal{T} followed by PSN^ν -iteration, P_1 , using u and F_2 be the multiset created
575 by a complete-iteration, C_2 , using $\mathcal{T} \cup \{u\}$.

576 • u is an insertion: $\langle p, L, \text{INS} \rangle$. Let $s_1 \in F_1$ be an update created. If s_1 is created
577 in C_1 , then $s_1 \in F_2$ since a delta rule of the same rule is fired in C_2 . If s_1 is
578 created in P_1 , then either the delta rule that is fired does not use any updates
579 in \mathcal{T} , in which case the same delta rule is also fired in C_2 , thus $s_1 \in F_2$; or
580 the delta rule use updates in \mathcal{T} , in which case there is another delta rule of the
581 same rule that is fired in C_2 , namely the one where the delta appears in the
582 right-most position (left-most position) if s_1 insertion (deletion) with respect to
583 the updates used; hence, $s_1 \in F_2$. Now, for the reverse direction, the reasoning
584 is much easier. Let $s_2 \in F_2$ be an update created, by using the update $\langle p, L, \text{INS} \rangle$
585 then a delta rule of the same rule is fired in P_1 ; hence $s_2 \in F_1$. Otherwise, the
586 same delta rule is fired in C_1 and therefore $s_2 \in F_1$.

587 • u is a deletion: $\langle p, L, \text{DEL} \rangle$. Again, let $s_1 \in F_1$ be an update created. If s_1 is
588 created in C_1 not using the tuple L from p , then the same rule is fired in C_2 ;
589 hence $s_1 \in F_2$. Otherwise, s_1 is created in C_1 using the tuple L from p , then
590 s_1 there is another delta rule of this rule in C_2 , hence $s_2 \in F_2$, namely the one
591 where the delta appears in the right-most position (resp. left-most position) if s_1
592 insertion (resp. deletion) with respect to the updates used. Now, for the reverse
593 direction, the reasoning is similar to the previous case.

594 **Theorem 1.** *Let \mathcal{D} be a set of ground atoms, \mathcal{P} be a non-recursive Datalog*
595 *program, \mathcal{U} be a multiset of updates, and s be a ground atom. There is a PSN^ν -*
596 *proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$ iff there is an SN-proof of $\mathcal{S}(\mathcal{D}, \mathcal{P}, \mathcal{U}, s)$.*

597 *Proof.* (\Leftarrow) Given a PSN^ν -proof, we construct an SN-proof by induction as
598 follows: use Lemma 2 to permute PSN^ν -iteration that picks an element $u \in \mathcal{U}$,
599 then repeat it with its subproof. The resulting proof has all PSN^ν -iteration in
600 the same order as in an SN-Proof, but they have to be merged into SN-iterations,
601 which is possible by applying repeatedly Lemma 3. This process terminates since
602 there are finitely many possible updates in a non-recursive program.

603 (\Rightarrow) Given an SN-proof, we repeatedly apply Lemma 3 to obtain a PSN^ν -proof.