

## LOCALIZED RETE FOR INCREMENTAL GRAPH QUERIES WITH NESTED GRAPH CONDITIONS

MATTHIAS BARKOWSKY  AND HOLGER GIESE 

Hasso Plattner Institute at the University of Potsdam, Prof.-Dr.-Helmert-Str. 2-3, D-14482 Potsdam, Germany  
*e-mail address:* {matthias.barkowsky,holger.giese}@hpi.de

**ABSTRACT.** The growing size of graph-based modeling artifacts in model-driven engineering calls for techniques that enable efficient execution of graph queries. Incremental approaches based on the RETE algorithm provide an adequate solution in many scenarios, but are generally designed to search for query results over the entire graph. However, in certain situations, a user may only be interested in query results for a subgraph, for instance when a developer is working on a large model of which only a part is loaded into their workspace. In this case, the global execution semantics can result in significant computational overhead.

To mitigate the outlined shortcoming, in this article we propose an extension of the RETE approach that enables local, yet fully incremental execution of graph queries, while still guaranteeing completeness of results with respect to the relevant subgraph.

We empirically evaluate the presented approach via experiments inspired by a scenario from software development and with queries and data from an independent social network benchmark. The experimental results indicate that the proposed technique can significantly improve performance regarding memory consumption and execution time in favorable cases, but may incur a noticeable overhead in unfavorable cases.

### 1. INTRODUCTION

In model-driven engineering, models constitute important development artifacts [Ken02]. With complex development projects involving large, interconnected models, performance of automated model operations becomes a primary concern.

Incremental graph query execution based on the RETE algorithm [For89] has been demonstrated to be an adequate solution in scenarios where an evolving model is repeatedly queried for the same information [SIRV18]. In this context, the RETE algorithm essentially tackles the problem of querying a usually graph-based model representation via operators from relational algebra [Cod70]. Thereby, the RETE algorithm supports not only graph queries in the form of plain graph patterns, but also covers advanced formalisms for query specification such as nested graph conditions [HP09]. On the flip side, the relational operators are not designed to exploit the locality typically found in graph-based encodings. Current incremental querying techniques consequently require processing of the entire model to guarantee complete results. However, in certain situations, global query execution is not required and may be undesirable due to performance considerations.

For instance, a developer may be working on only a part of a model loaded into their workspace, with a large portion of the full model still stored on disk. A concrete example of

this would be a developer using a graphical editor to modify an individual block diagram from a collection of interconnected block diagrams, which together effectively form one large model. As the developer modifies the model part in their workspace, they may want to continuously monitor how their changes impact the satisfaction of some consistency constraints via incremental model queries. In this scenario, existing techniques for incremental query execution require loading and querying the entire collection of models, even though the user is ultimately only interested in the often local effect of their own changes.

The global RETE execution semantics then results in at least three problems: First, the computation of query results for the full model, of which only a portion is relevant to the user, may incur a substantial overhead on initial query execution time. Second, incremental querying techniques are known to create and store a large number of intermediate results [SIRV18], many of which can in this scenario be superfluous, causing an overhead in memory consumption. Third, query execution requires loading the entire model into memory, potentially causing an overhead in loading time and increasing the overhead in memory consumption.

These problems can be mitigated to some extent by employing local search [CFSV04, GHS09, ABJ<sup>+</sup>10], which better exploits the locality of the problem and lazy loading capabilities of model persistence layers [cd0, DSB<sup>+</sup>17], instead of a RETE-based technique. However, resorting to local search can result in expensive redundant search operations that are only avoided by fully incremental solutions [SIRV18].

In this article, we instead propose to tackle the outlined shortcoming of incremental querying techniques via an extension of the RETE approach. This is supported by a relaxed notion of completeness for query results that accounts for situations where the computation for the full model is unnecessary. Essentially, this enables a distinction between the full model, for which query results do not necessarily have to be complete, and a relevant model part, for which complete results are required. The extended RETE approach then anchors query execution to the relevant model part and lazily fetches additional model elements required to compute query results that meet the relaxed completeness requirement. Our approach thereby avoids potentially expensive global query execution and allows an effective integration of incremental queries with model persistence layers.

The remainder of this article is structured as follows: Section 2 provides a summary of our notion of graphs, graph queries, and the RETE mechanism for query execution. Section 3 introduces a relaxed notion of completeness for results of model queries consisting of a plain graph pattern and subsequently presents an adaptation to the RETE querying mechanism that enables local, yet fully incremental execution of such plain model queries. Our technique for local query execution via RETE is then extended to the case of more complex model queries featuring nested graph conditions in Section 4. A prototypical implementation of our approach is evaluated regarding execution time and memory consumption in Section 5. Section 6 discusses related work, before Section 7 concludes the article and provides an outlook on future work.

This article is an extension of our paper at ICGT 2024 [BG24a]. The main increment of this article compared to the conference paper version is the approach for localizing execution of graph queries with nested graph conditions presented in Section 4, which is also covered by an extended evaluation in Section 5. Besides small improvements to text and figures and the inclusion of further examples, we have also added a theorem for recognizing locally relevant query results as another minor extension in Section 3.

## 2. PRELIMINARIES

In the following, we briefly reiterate the definitions of graphs and graph queries and summarize the RETE approach to incremental graph querying.

**2.1. Graphs and Graph Queries.** As defined in [EEPT06], a *graph* is a tuple  $G = (V^G, E^G, s^G, t^G)$ , with  $V^G$  the set of vertices,  $E^G$  the set of edges, and  $s^G : E^G \rightarrow V^G$  and  $t^G : E^G \rightarrow V^G$  functions mapping edges to their source respectively target vertices.

A mapping from a graph  $Q$  into another graph  $H$  is defined via a *graph morphism*  $m : Q \rightarrow H$ , which is characterized by two functions  $m^V : V^Q \rightarrow V^H$  and  $m^E : E^Q \rightarrow E^H$  such that  $s^H \circ m^E = m^V \circ s^Q$  and  $t^H \circ m^E = m^V \circ t^Q$ . A graph morphism is called a monomorphism if  $m^V$  and  $m^E$  are injective and an isomorphism if  $m^V$  and  $m^E$  are also surjective. A *partial graph morphism* consists of two partial functions  $m^V : V^Q \rightarrow V^H$  and  $m^E : E^Q \rightarrow E^H$  with  $s^H \circ m^E \leq m^V \circ s^Q$  and  $t^H \circ m^E \leq m^V \circ t^Q$  [Rib99]. In this context, for two partial functions  $f$  and  $g$ ,  $f \leq g$  denotes inclusion of the domain of  $f$  in the domain of  $g$  and equality of mappings in  $f$  and  $g$  for all domain elements for which  $f$  is defined.

A *graph modification* from a graph  $G$  into a modified graph  $G'$  can be described via an intermediate graph  $K$  and a span of graph monomorphisms  $f : K \rightarrow G$  and  $g : K \rightarrow G'$  [TELW14], which can also be denoted  $G \xleftarrow{f} K \xrightarrow{g} G'$ . Intuitively, elements in  $G \setminus f(K)$  are deleted and elements in  $G' \setminus g(K)$  are created by the modification.

A graph  $G$  can be typed over a type graph  $TG$  via a graph morphism  $\text{type}^G : G \rightarrow TG$ , forming the *typed graph*  $G^T = (G, \text{type}^G)$ . A *typed graph morphism* from  $G^T = (G, \text{type}^G)$  into another typed graph  $H^T = (H, \text{type}^H)$  with type graph  $TG$  is a graph morphism  $m^T : G \rightarrow H$  such that  $\text{type}^H \circ m^T = \text{type}^G$ .

A *model* is then simply given by a typed graph. In the remainder of the article, we therefore use the terms graph and model interchangably and implicitly assume typing when talking about graphs.

In practice, graphs often contain more edges than nodes, a characteristic that may influence the performance of certain algorithms. To capture this property for each individual edge type in a typed graph, we introduce the notion of *edge-dominated* graphs:

**Definition 2.1.** (*Edge-Dominated Typed Graphs*) We say that a typed graph  $G$  is *edge-dominated* if  $\forall e_{TG} \in E^{TG} : |\{e \in E^G \mid \text{type}(e) = e_{TG}\}| \geq \max(|\{v \in V^G \mid \text{type}(v) = s^{TG}(e_{TG})\}|, |\{v \in V^G \mid \text{type}(v) = t^{TG}(e_{TG})\}||)$ .

The above definition will be relevant in the context of an analytical evaluation of our localized querying technique, which technically requires graphs to be edge-dominated in order to guarantee desirable characteristics regarding computational complexity.

A *plain graph query* as considered in this article is characterized by a query graph  $Q$  and can be executed over a host graph  $H$  by finding graph morphisms  $m : Q \rightarrow H$ . We also call these morphisms *matches* and denote the set of all matches from  $Q$  into  $H$  by  $\mathcal{M}_H^Q$ . Typically, a set of matches is considered a complete query result if it contains all matches in  $\mathcal{M}_H^Q$ :

**Definition 2.2.** (*Completeness of Plain Query Results*) We say that a set of matches  $M$  from query graph  $Q$  into host graph  $H$  is complete iff  $\mathcal{M}_H^Q \subseteq M$ .

Note that by virtue of  $M$  being defined as a set of matches from  $Q$  into  $H$ , the statement  $\mathcal{M}_H^Q \subseteq M$  in the above definition implies that  $M = \mathcal{M}_H^Q$ .

In some cases, the formalism of plain graph queries may not be sufficiently expressive to specify the desired query, for instance if matches for the query should not be extensible by a specific negative pattern. To increase expressiveness, a plain graph query  $Q$  can therefore be equipped with a *nested graph condition*  $\psi$  [HP09], yielding the *extended graph query*  $(Q, \psi)$ .

Nested graph conditions can be defined recursively in the context of a graph morphism  $m : Q \rightarrow H$ , with  $m \models \psi$  denoting that  $m$  satisfies the condition  $\psi$ :

- $\text{true}$  is a nested graph condition that is always satisfied, that is,  $m \models \text{true}$  for any graph morphism  $m$ .
- $\neg\psi$  is a nested graph condition that is satisfied iff  $m$  does not satisfy the nested graph condition  $\psi$ , that is,  $m \models \neg\psi \Leftrightarrow m \not\models \psi$ .
- $\psi_1 \wedge \psi_2$  is a nested graph condition that is satisfied iff  $m$  satisfies both nested graph conditions  $\psi_1$  and  $\psi_2$ , that is,  $m \models \psi_1 \wedge \psi_2 \Leftrightarrow (m \models \psi_1 \wedge m \models \psi_2)$ .
- $\exists(a : Q \rightarrow Q', \psi')$ , where  $a : Q \rightarrow Q'$  is a graph monomorphism, is a nested graph condition that is satisfied iff there exists a graph morphism  $m' : Q' \rightarrow H$  such that  $m = m' \circ a$  and  $m'$  satisfies the nested graph condition  $\psi'$ , that is,  $m \models \exists(a : Q \rightarrow Q', \psi') \Leftrightarrow (\exists m' \in \mathcal{M}_H^{Q'} : m' \models \psi' \wedge m = m' \circ a)$ .

As in [BG23], we also consider the case where the morphism  $a$  of a nested graph condition  $\exists(a : Q \rightarrow Q', \psi')$  can be a partial graph morphism from some subgraph  $Q_p \subseteq Q$  of  $Q$  into  $Q'$ . In this case, satisfaction of the nested graph condition is given by  $m \models \exists(a : Q \rightarrow Q', \psi') \Leftrightarrow (\exists m' \in \mathcal{M}_H^{Q'} : m' \models \psi' \wedge m|_{Q_p} = m' \circ a)$ .

Equipping the query graph  $Q$  with the nested graph condition  $\psi$  restricts the set of valid matches for the extended graph query  $(Q, \psi)$  to those matches for  $Q$  that also satisfy  $\psi$ . In contrast to plain graph queries, completeness alone is thus insufficient as a criterion for describing what constitutes a desirable query result. Specifically, when talking about extended graph queries, a query result should not only contain all admissible matches, but also no matches that violate the equipped nested graph condition. In this article, we will hence talk about *completeness* with respect to the results of plain graph queries and *correctness* in the case of extended graph queries, with the latter defined as follows:

**Definition 2.3.** (*Correctness of Extended Query Results*) We say that a set of matches  $M$  for the extended graph query  $(Q, \psi)$  into host graph  $H$  is correct iff  $M = \{m \in \mathcal{M}_H^Q \mid m \models \psi\}$ .

**2.2. Incremental Graph Queries with RETE.** The RETE algorithm [For89] forms the basis of mature incremental graph querying techniques [VBH<sup>+</sup>16]. Therefore, the query is recursively decomposed into simpler subqueries, which are arranged in a second graph called RETE net. In the following, we will refer to the vertices of RETE nets as *(RETE) nodes*.

For plain graph queries, each RETE node  $n$  computes matches for a subgraph  $query(n) \subseteq Q$  of the full query graph  $Q$ . This computation may depend on matches computed by other RETE nodes. Such dependencies are represented by edges from the dependent node to the dependency node. For each RETE net, one of its nodes is designated as the *production node*, which computes the net's overall result. A RETE net is thus given by a tuple  $(N, p)$ , where the *RETE graph*  $N$  is a graph of RETE nodes and dependency edges and  $p \in V^N$  is the production node.

We describe the *configuration* of a RETE net  $(N, p)$  during execution by a function  $\mathcal{C} : V^N \rightarrow \mathcal{P}(\mathcal{M}_\Omega)$ , with  $\mathcal{P}$  denoting the power set and  $\mathcal{M}_\Omega$  the set of all graph morphisms.  $\mathcal{C}$  then assigns each node in  $V^N$  a *current result set*.

For a starting configuration  $\mathcal{C}$  and host graph  $H$ , executing a RETE node  $n \in V^N$  yields an updated configuration  $\mathcal{C}' = \text{execute}(n, N, H, \mathcal{C})$ , with

$$\mathcal{C}'(n') = \begin{cases} \mathcal{R}(n, N, H, \mathcal{C}) & \text{if } n' = n \\ \mathcal{C}(n') & \text{otherwise,} \end{cases}$$

where  $\mathcal{R}$  is a function defining the *target result set* of a RETE node  $n$  in the RETE graph  $N$  for  $H$  and  $\mathcal{C}$  such that  $\mathcal{R}(n, N, H, \mathcal{C}) \subseteq \mathcal{M}_H^Q$ , with  $Q = \text{query}(n)$ .

We say that  $\mathcal{C}$  is *consistent* for a RETE node  $n \in V^N$  and host graph  $H$  iff  $\mathcal{C}(n) = \mathcal{R}(n, N, H, \mathcal{C})$ . We furthermore define a consistent configuration as a configuration that is consistent for all nodes of the RETE net:

**Definition 2.4** (Consistent RETE Net Configurations). We say that a configuration  $\mathcal{C} : V^N \rightarrow \mathcal{P}(\mathcal{M}_\Omega)$  for a RETE net  $(N, p)$  is consistent for a host graph  $H$  iff  $\forall n \in V^N : \mathcal{C}(n) = \mathcal{R}(n, N, H, \mathcal{C})$ .

If  $H$  is clear from the context, we simply say that  $\mathcal{C}$  is consistent.

Given a host graph  $H$  and a starting configuration  $\mathcal{C}_0$ , a RETE net  $(N, p)$  is executed via the execution of a sequence of nodes  $O = n_1, n_2, \dots, n_x$  with  $n_i \in V^N$ . This yields the trace  $\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_x$ , with  $\mathcal{C}_y = \text{execute}(n_y, N, H, \mathcal{C}_{y-1})$ , and the result configuration  $\text{execute}(O, N, H, \mathcal{C}_0) = \mathcal{C}_x$ .

$(N, p)$  can initially be executed over  $H$  via a sequence  $O$  that yields a consistent configuration  $\mathcal{C}_x = \text{execute}(O, N, H, \mathcal{C}_0)$ , where  $\mathcal{C}_0$  is an empty starting configuration with  $\forall n \in V^N : \mathcal{C}_0(n) = \emptyset$ . Incremental execution of  $(N, p)$  can be achieved by retaining a previous result configuration and using it as the starting configuration for execution over an updated host graph. This requires incremental implementations of RETE node execution procedures that update previously computed result sets for changed inputs instead of computing them from scratch.

In the most basic form, a RETE net consists of two kinds of nodes, *edge input nodes* and *join nodes*. An edge input node  $[v \rightarrow w]$  has no dependencies, is associated with the query subgraph  $\text{query}([v \rightarrow w]) = (\{v, w\}, \{e\}, \{(e, v)\}, \{(e, w)\})$ , and directly extracts the corresponding (trivial) matches from a host graph. A join node  $[\bowtie]$  has two dependencies  $n_l$  and  $n_r$  with  $Q_l = \text{query}(n_l)$  and  $Q_r = \text{query}(n_r)$  such that  $V^{Q_l \cap Q_r} \neq \emptyset$  and is associated with  $\text{query}([\bowtie]) = Q_l \cup Q_r$ .  $[\bowtie]$  then computes matches for this union subgraph by combining the matches from its dependencies along the overlap graph  $Q_\cap = Q_l \cap Q_r$ .

In the following, we assume query graphs to be weakly connected and contain at least one edge. A RETE net that computes matches for a query graph  $Q$  can then be constructed as a tree of join nodes over edge input nodes. The join nodes thus gradually compose the trivial matches at the bottom into matches for more complex query subgraphs. We call such tree-like RETE nets *well-formed*. An execution sequence that always produces a consistent configuration is given by a reverse topological sorting of the net. The root node of the tree then computes the set of all matches for  $Q$  and is designated as the net's production node.

Connected graphs without edges consist of only a single vertex, making query execution via *vertex input nodes*, which function analogously to edge input nodes, trivial. Disconnected query graphs can be handled via separate RETE nets for all query graph components and the computation of a cartesian product.

Figure 1 shows an example plain graph query from the software domain and an associated RETE net. The query searches for paths of a package, class, and field. The RETE net

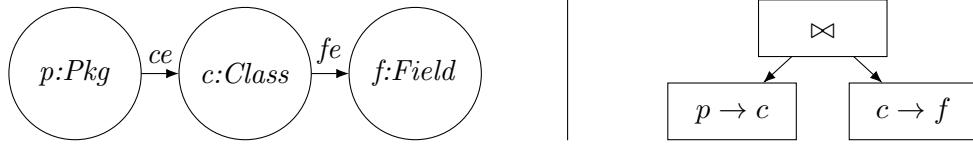


FIGURE 1. Example plain graph query (left) and corresponding RETE net (right)

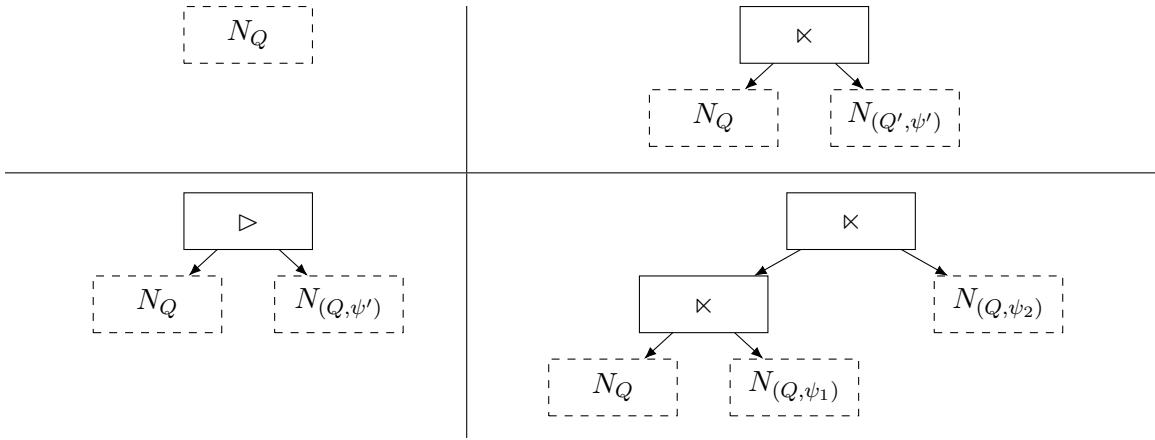


FIGURE 2. RETE net  $(N_{(Q,\psi)}, p_{(Q,\psi)})$  for an extended graph query  $(Q, \psi)$  with nested graph condition of the form  $\psi = \text{true}$  (top left),  $\psi = \exists(a : Q \rightarrow Q', \psi')$  (top right),  $\psi = \neg\psi'$  (bottom left), and  $\psi = \psi_1 \wedge \psi_2$  (bottom right)

constructs matches for the query by combining edges from a package to a class with edges from a class to a field via a natural join.

The RETE approach to graph query execution also supports extended graph queries via two additional RETE node types: *semi-join nodes* and *anti-join nodes*. Like join nodes, semi-join nodes and anti-join nodes have two dependencies  $n_l$  and  $n_r$  with  $Q_l = \text{query}(n_l)$  and  $Q_r = \text{query}(n_r)$  such that  $V^{Q_l \cap Q_r} \neq \emptyset$ . However, instead of combining matches from both dependencies, a semi-join [ $\bowtie$ ] simply computes all matches from the left dependency that have a compatible match in the right dependency. An anti-join [ $\triangleright$ ] then computes all matches from the left dependency that do *not* have a compatible match in the right dependency. Both semi-join and anti-join nodes are thus always associated with the same query graph as their left dependency, that is,  $\text{query}([\bowtie]) = Q_l$  and  $\text{query}([\triangleright]) = Q_l$ .

To support nested graph conditions of the form  $\exists(a : Q \rightarrow Q', \psi)$ , a semi-join can also be defined along a non-empty, potentially partial graph morphism [BG23]. The definition based on an overlap graph is then just a special case of a semi-join along a partial identity graph morphism.

Using semi-join and anti-join nodes and a RETE net  $(N_Q, p_Q)$  for the plain graph query  $Q$ , a RETE net  $(N_{(Q,\psi)}, p_{(Q,\psi)})$  for an extended graph query  $(Q, \psi)$  can be constructed recursively according to our description in [BG23], which is visualized in Figure 2.

Figure 3 shows an example extended graph query on the left. The extended query searches for the same pattern as the plain query in Figure 1, but also requires the class node to be connected to an interface node, without this additional node being part of the main

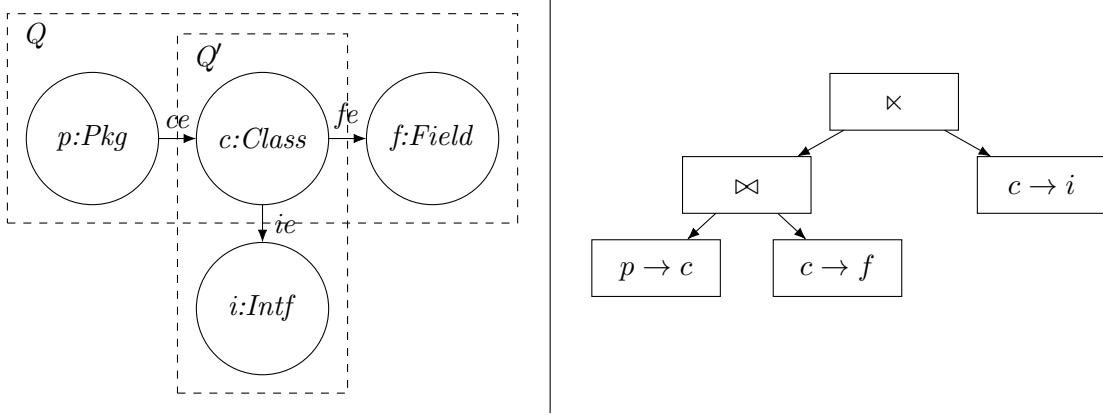


FIGURE 3. Example extended graph query  $(Q, \exists(a : Q \rightarrow Q', \text{true}))$  (left) and corresponding RETE net (right)

match. Therefore, the query is defined as  $(Q, \exists(a : Q \rightarrow Q', \text{true}))$  with a partial graph morphism  $a$  that maps only the class node in  $Q$  to the class node in  $Q'$ .

The corresponding RETE net is visualized on the right of Figure 3. For computing matches for the base pattern, the RETE net replicates the structure in Figure 1. The nested graph condition is implemented via a search for the pattern  $Q'$ , which consists of a single edge from the node  $c$  to  $i$ , and a semi-join node at the top of the net, which checks the matches for  $Q$  against the matches for  $Q'$ .

### 3. INCREMENTAL PLAIN QUERIES OVER SUBGRAPHS

As outlined in Section 1, users of model querying mechanisms may only be interested in query results related to some part of a model that is relevant to them rather than the complete model. However, simply executing a query only over this relevant subgraph and ignoring context elements is often insufficient, for instance if the full effect of modifications to the relevant subgraph should be observed, since such modifications may affect matches that involve elements outside the relevant subgraph. Instead, completeness in such scenarios essentially requires the computation of all matches that somehow *touch* the relevant subgraph.

In order to capture this need for local completeness, but avoid the requirement for global execution inherent to the characterization in Definition 2.2, we define completeness under a relevant subgraph of some host graph as follows:<sup>1</sup>

**Definition 3.1** (Completeness of Plain Query Results under Subgraphs). We say that a set of matches  $M$  from query graph  $Q$  into host graph  $H$  is complete under a subgraph  $H_p \subseteq H$  iff  $\{m \in \mathcal{M}_H^Q \mid m(Q) \cap H_p \neq \emptyset\} \subseteq M$ .

<sup>1</sup>Note that by the definition of  $H_p$  as a subgraph of  $H$  and the definition of matches as graph morphisms,  $m(Q) \cap H_p \neq \emptyset \Leftrightarrow m(V^Q) \cap H_p \neq \emptyset$ . Our slightly adapted definition in this article thus coincides with the definition in the conference paper version [BG24a].

**3.1. Marking-sensitive RETE.** Due to its reliance on edge and vertex input nodes and their global execution semantics, incremental query execution via the standard RETE approach is unable to exploit the relaxed notion of completeness from Definition 3.1 for query optimization and does not integrate well with mechanisms relying on operation locality, such as model persistence layers based on lazy loading [cd, DSB<sup>+</sup>17].

While query execution could be localized to a relevant subgraph by executing edge and vertex inputs nodes only over the subgraph, execution would then only yield matches where *all* involved elements are in the relevant subgraph. This approach would hence fail to meet the completeness criterion of Definition 3.1.

To enable incremental queries with complete results under the relevant subgraph, we instead propose to anchor RETE net execution to subgraph elements while allowing the search to retrieve elements outside the subgraph that are required to produce complete results from the full model.

This approach is based on the observation that, in order for a join node to create a match that involves elements from the relevant subgraph, at least one of the constituting matches already has to involve at least one such element. Intuitively, our localized querying technique then works as follows: At the leaves of a RETE join tree for some plain graph query, we initially compute only matches that touch the relevant subgraph via local edge navigation. We then also implement a mechanism that enables a lazy computation of certain complementary partial matches that do not touch the relevant subgraph at inner nodes of the join tree. This ensures completeness of results with respect to Definition 3.1.

A naïve implementation of this idea that retrieves complementary matches for all matches from each dependency of a join  $[\bowtie]$  until a fixpoint is reached however runs the risk of triggering the computation of superfluous matches. In this scenario, some match  $m_l$  for the dependency  $n_l$  of  $[\bowtie]$  that was only required to complement some relevant match  $m_r$  from the other dependency  $n_r$  would again trigger the computation of complementary matches for  $n_r$ , causing unnecessary computational effort.

Moreover, such an approach would create cyclic requirement relationships between the matches  $m_l$  and  $m_r$ , where the presence of  $m_l$  in the result for  $n_l$  mandates the presence of  $m_r$  in the result for  $n_r$  and vice-versa. This may prevent the unloading of matches when elements are removed from the relevant subgraph.

To mitigate these issues, we extend the standard RETE mechanism by a marking for matches in the form of a natural number. Intuitively, we will use this marking to encode up to which height in a RETE join tree derived matches should trigger the retrieval of complementary matches. An appropriate assignment of such markings then prevents the computation of the aforementioned superfluous matches as well as cyclic requirement relationships between matches, while still guaranteeing completeness according to Definition 3.1.

In our extension, an intermediate result in a *marking-sensitive RETE net* is therefore characterized by a tuple  $(m, \phi)$  of a match  $m$  and a marking  $\phi \in \bar{\mathbb{N}}$ , where we define  $\bar{\mathbb{N}} := \mathbb{N} \cup \{\infty\}$ . A configuration for a marking-sensitive RETE net  $(N^\Phi, p^\Phi)$  is then given by a function  $\mathcal{C}^\Phi : V^{N^\Phi} \rightarrow \mathcal{P}(\mathcal{M}_\Omega \times \bar{\mathbb{N}})$ .

Furthermore, in our extension, result computation distinguishes between the full host graph  $H$  and the relevant subgraph  $H_p \subseteq H$ . For marking-sensitive RETE nodes, we hence extend the function for target result sets by a parameter for  $H_p$ . The target result set of a marking-sensitive RETE node  $n^\Phi \in V^{N^\Phi}$  with  $Q = \text{query}(n^\Phi)$  for  $H, H_p$ ,

and a marking-sensitive configuration  $\mathcal{C}^\Phi$  is then given by  $\mathcal{R}^\Phi(n^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi)$ , with  $\mathcal{R}^\Phi(n^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^Q \times \bar{\mathbb{N}}$ .

Consistency of marking-sensitive configurations is then defined analogously to the standard case:

**Definition 3.2** (Consistent Marking-Sensitive RETE Net Configurations). We say that a configuration  $\mathcal{C}^\Phi : V^{N^\Phi} \rightarrow \mathcal{P}(\mathcal{M}_\Omega \times \bar{\mathbb{N}})$  for a marking-sensitive RETE net  $(N^\Phi, p^\Phi)$  is consistent for a host graph  $H$  with relevant subgraph  $H_p \subseteq H$  iff  $\forall n^\Phi \in V^{N^\Phi} : \mathcal{C}^\Phi(n^\Phi) = \mathcal{R}^\Phi(n^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi)$ .

Note that technically, the definitions of marking-sensitive configurations and target result sets of marking-sensitive RETE nodes allow the same match to be associated with multiple markings for the same RETE node. However, as shown in [BG24b], the marking of a match in the target result set of a marking-sensitive RETE node is always unique in the constructions presented in this article.

We adapt the *join node*, *union node*, and *projection node* to marking-sensitive variants that assign result matches the maximum marking of related dependency matches and otherwise work as expected from relational algebra [Cod70]. We also adapt the *vertex input node* to only consider vertices in the relevant subgraph  $H_p$  and assign matches the marking  $\infty$ . Finally, we introduce *marking assignment nodes*, which assign matches a fixed marking value, *marking filter nodes*, which filter marked matches by a minimum marking value, and *forward* and *backward navigation nodes*, which work similarly to edge input nodes but only extract edges that are adjacent to host graph vertices included in the current result set of a designated dependency. Note that an efficient implementation of the backward navigation node requires reverse navigability of host graph edges.

Formally, we define the target result sets of marking-sensitive RETE nodes as follows:

**Definition 3.3** (Target Result Sets of Marking-Sensitive RETE Nodes). Let  $H$  be a host graph with relevant subgraph  $H_p \subseteq H$ ,  $(N^\Phi, p^\Phi)$  a containing marking-sensitive RETE net, and  $\mathcal{C}^\Phi$  a configuration for  $(N^\Phi, p^\Phi)$ .

- The target result set of a *marking-sensitive join node*  $[\bowtie]^\Phi$  with marking-sensitive dependencies  $n_l^\Phi$  and  $n_r^\Phi$  with  $Q_l = \text{query}(n_l^\Phi)$  and  $Q_r = \text{query}(n_r^\Phi)$  such that  $V^{Q_l \cap Q_r} \neq \emptyset$  for  $Q_l \cap Q_r = Q_l \cap Q_r$  is given by  $\mathcal{R}^\Phi([\bowtie]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m_l \cup m_r, \max(\phi_l, \phi_r)) \mid (m_l, \phi_l) \in \mathcal{C}^\Phi(n_l^\Phi) \wedge (m_r, \phi_r) \in \mathcal{C}^\Phi(n_r^\Phi) \wedge m_l|_{Q_l \cap Q_r} = m_r|_{Q_l \cap Q_r}\}$ .
- The target result set of a *marking-sensitive union node*  $[\cup]^\Phi$  with a set of marking-sensitive dependencies  $N_\alpha^\Phi$  such that  $Q = \text{query}(n_\alpha^\Phi)$  for all dependencies  $n_\alpha^\Phi \in N_\alpha^\Phi$  is given by  $\mathcal{R}^\Phi([\cup]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m, \phi_{max}) \mid (m, \phi_{max}) \in \bigcup_{n_\alpha \in N_\alpha^\Phi} \mathcal{C}^\Phi(n_\alpha^\Phi) \wedge \phi_{max} = \max(\{\phi' \mid (m, \phi') \in \mathcal{C}^\Phi(n_\alpha^\Phi)\})\}$ .
- The target result set of a *marking-sensitive projection node*  $[\pi_Q]^\Phi$  with a single marking-sensitive dependency  $n_\alpha^\Phi$  and  $Q = \text{query}([\pi_Q]^\Phi)$  is given by  $\mathcal{R}^\Phi([\pi_Q]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m|_Q, \phi_{max}) \mid (m, \phi_{max}) \in \mathcal{C}^\Phi(n_\alpha^\Phi) \wedge \phi_{max} = \max(\{\phi' \mid (m', \phi') \in \mathcal{C}^\Phi(n_\alpha^\Phi) \wedge m'|_Q = m|_Q\})\}$ .
- The target result set of a *marking assignment node*  $[\phi := i]^\Phi$  with a single marking-sensitive dependency  $n_\alpha^\Phi$  is given by  $\mathcal{R}^\Phi([\phi := i]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m, i) \mid (m, \phi_\alpha) \in \mathcal{C}^\Phi(n_\alpha^\Phi)\}$ .
- The target result set of a *marking filter node*  $[\phi > i_{min}]^\Phi$  with marking-sensitive dependency  $n_\alpha^\Phi$  is given by  $\mathcal{R}^\Phi([\phi > i_{min}]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m, \phi_\alpha) \mid (m, \phi_\alpha) \in \mathcal{C}^\Phi(n_\alpha^\Phi) \wedge \phi_\alpha > i_{min}\}$ .

- The target result set of a *forward navigation node*  $[v \rightarrow_n w]^\Phi$  with  $Q = \text{query}([v \rightarrow_n w]^\Phi) = (\{v, w\}, \{e\}, \{(e, v)\}, \{(e, w)\})$  and marking-sensitive dependency  $n_v^\Phi$  with  $Q_v = \text{query}(n_v^\Phi) = (\{v\}, \emptyset, \emptyset, \emptyset)$  is given by  $\mathcal{R}^\Phi([v \rightarrow_n w]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m, \phi) \mid m \in \mathcal{M}_H^Q \wedge (m|_{Q_v}, \phi) \in \mathcal{C}^\Phi(n_v^\Phi)\}$ .
- The target result set of a *backward navigation node*  $[w \leftarrow_n v]^\Phi$  with  $Q = \text{query}([w \leftarrow_n v]^\Phi) = (\{v, w\}, \{e\}, \{(e, v)\}, \{(e, w)\})$  and marking-sensitive dependency  $n_w^\Phi$  with  $Q_w = \text{query}(n_w^\Phi) = (\{w\}, \emptyset, \emptyset, \emptyset)$  is given by  $\mathcal{R}^\Phi([w \leftarrow_n v]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m, \phi) \mid m \in \mathcal{M}_H^Q \wedge (m|_{Q_w}, \phi) \in \mathcal{C}^\Phi(n_w^\Phi)\}$ .
- The target result set of a *marking-sensitive vertex input node*  $[v]^\Phi$  with  $Q = \text{query}([v]^\Phi) = (\{v\}, \emptyset, \emptyset, \emptyset)$  is given by  $\mathcal{R}^\Phi([v]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m, \infty) \mid m \in \mathcal{M}_{H_p}^Q\}$ .

To obtain query results in the format of the standard RETE approach, we define the *stripped result set* of a marking-sensitive RETE node  $n^\Phi$  for a marking-sensitive configuration  $\mathcal{C}^\Phi$  as the set of matches that appear in tuples in the node's current result set in  $\mathcal{C}^\Phi$ .

**Definition 3.4** (Stripped Result Sets of Marking-Sensitive RETE Nodes). The *stripped result set* of a marking-sensitive RETE node  $n^\Phi$  for a marking-sensitive configuration  $\mathcal{C}^\Phi$  is given by  $\mathcal{R}_X^\Phi(n^\Phi, \mathcal{C}^\Phi) = \{m \mid (m, \phi) \in \mathcal{C}^\Phi(n^\Phi)\}$ .

**3.2. Localized Search with Marking-sensitive RETE.** Based on these adaptations, we introduce a recursive *localize* procedure, which takes a regular, well-formed<sup>2</sup> RETE net  $(N, p)$  as input and outputs a marking-sensitive RETE net, which performs a localized search that does not require searching the full model to produce complete results according to Definition 3.1.

If  $p = [v \rightarrow w]$  is an edge input node, the result of localization for  $(N, p)$  is given by  $\text{localize}(N, p) = (LNS(p), [\cup]^\Phi)$ . The *local navigation structure*  $LNS(p)$  consists of seven RETE nodes as shown in Figure 4 (left): (1, 2) Two marking-sensitive vertex input nodes  $[v]^\Phi$  and  $[w]^\Phi$ , (3, 4) two marking-sensitive union nodes  $[\cup]_v^\Phi$  and  $[\cup]_w^\Phi$  with  $[v]^\Phi$  respectively  $[w]^\Phi$  as a dependency, (5) a forward navigation node  $[v \rightarrow_n w]^\Phi$  with  $[\cup]_v^\Phi$  as a dependency, (6) a backward navigation node  $[w \leftarrow_n v]^\Phi$  with  $[\cup]_w^\Phi$  as a dependency, and (7) a marking-sensitive union node  $[\cup]^\Phi$  with dependencies  $[v \rightarrow_n w]^\Phi$  and  $[w \leftarrow_n v]^\Phi$ .

Importantly, the marking-sensitive vertex input nodes of the local navigation structure are executed over the relevant subgraph, whereas the forward and backward navigation nodes are executed over the full model. Intuitively, the local navigation structure thus takes the role of the edge input node, but initially only extracts edges that are adjacent to a vertex in the relevant subgraph. Connecting additional dependencies to the union nodes  $[\cup]_v^\Phi$  and  $[\cup]_w^\Phi$  in further constructions then allows the extraction of additional edges that may be required to compute complete query results.

If  $p$  is a join node, it has two dependencies  $p_l$  and  $p_r$  with  $Q_l = \text{query}(p_l)$  and  $Q_r = \text{query}(p_r)$ , which are the roots of two RETE subtrees  $N_l$  and  $N_r$ . In this case,  $(N, p)$  is localized as  $\text{localize}(N, p) = (N^\Phi, p^\Phi) = (N_{\bowtie}^\Phi \cup N_l^\Phi \cup N_r^\Phi \cup RPS_l \cup RPS_r, [\bowtie]^\Phi)$ ,

<sup>2</sup>The restriction to well-formed RETE nets prevents an optimization where redundant computation of matches for isomorphic query subgraphs is avoided via a non-tree-like structure. However, for queries as in [EALP<sup>+</sup>15], performance is often primarily determined by the computation of matches for a large query subgraph that cannot be reused.

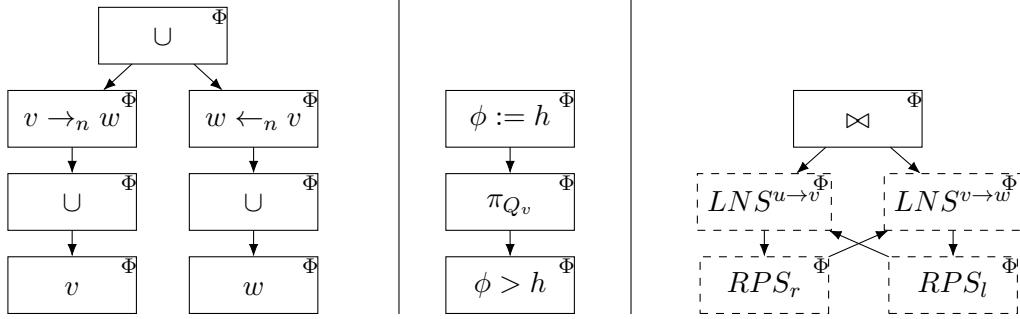


FIGURE 4. LNS (left), RPS (center), and localized RETE net (right)

where  $(N_l^\Phi, p_l^\Phi) = \text{localize}(N_l, p_l)$ ,  $(N_r^\Phi, p_r^\Phi) = \text{localize}(N_r, p_r)$ ,  $N_\bowtie^\Phi$  consists of the marking-sensitive join  $[\bowtie]^\Phi$  with dependencies  $p_l^\Phi$  and  $p_r^\Phi$ ,  $RPS_l = RPS(p_l^\Phi, N_r^\Phi)$ , and  $RPS_r = RPS(p_r^\Phi, N_l^\Phi)$ .

The *request projection structure*  $RPS_l = RPS(p_l^\Phi, N_r^\Phi)$  contains three RETE nodes as displayed in Figure 4 (center): (1) A marking filter node  $[\phi > h]^\Phi$  with  $p_l^\Phi$  as a dependency, (2) a marking-sensitive projection node  $[\pi_{Q_v}]^\Phi$  with  $[\phi > h]^\Phi$  as a dependency, and (3) a marking assignment node  $[\phi := h]^\Phi$  with  $[\pi_{Q_v}]^\Phi$  as a dependency. The value of  $h$  is given by the height of the join tree of which  $p$  is the root and  $Q_v$  is a graph consisting of a single vertex  $v \in V^{Q_l \cap Q_r}$ . The request projection structure is then connected to an arbitrary local navigation structure in  $N_r^\Phi$  that has a marking-sensitive vertex input  $[v]^\Phi$  with  $Q_v = \text{query}([v]^\Phi)$ . Therefore, it also adds a dependency from the marking-sensitive union node  $[\cup]_v^\Phi$  that already depends on  $[v]^\Phi$  to the marking assignment node  $[\phi := h]^\Phi$ . The mirrored structure  $RPS_r = RPS(p_r^\Phi, N_l^\Phi)$  is constructed analogously.

Via the request projection structures, partial matches from one join dependency can be propagated to the subnet under the other dependency. Intuitively, the inserted request projection structures thereby allow the join's dependencies to request the RETE subnet under the other dependency to fetch and process the model parts required to complement the first dependency's results. The marking of a match then signals up to which height in the join tree complementarity for that match is required. Notably, matches involving elements in the relevant subgraph are marked  $\infty$ , as complementarity for them is required at the very top to guarantee completeness of the overall result.

The result of applying *localize* to a RETE net consisting of a single join is displayed in Figure 4 (right). It consists of a marking sensitive join and two local navigation structures connected via request projection structures.

The localized RETE net for the example query in Figure 1 is shown in Figure 5. Alongside the RETE net, the figure visualizes a consistent configuration for the host graph in Figure 6. In the example host graph, the relevant subgraph consists of the node  $p_1$ , which is indicated by a bold border.

With the relevant subgraph consisting of only a single node, the only intermediate result in any of the net's vertex input nodes is the tuple  $(m_{1.1.1}, \infty)$  for the node  $[p]^\Phi$ . This intermediate result propagates up through the containing local navigation structure on the left, where it triggers the extraction of the match  $m_{1.1}$  by the forward navigation node  $[p \rightarrow_n c]^\Phi$ . Via the connected request projection structure, the restricted match  $m_{1.2.1}$  that consists only of the target node  $c_1$  of the extracted edge is then propagated into the local

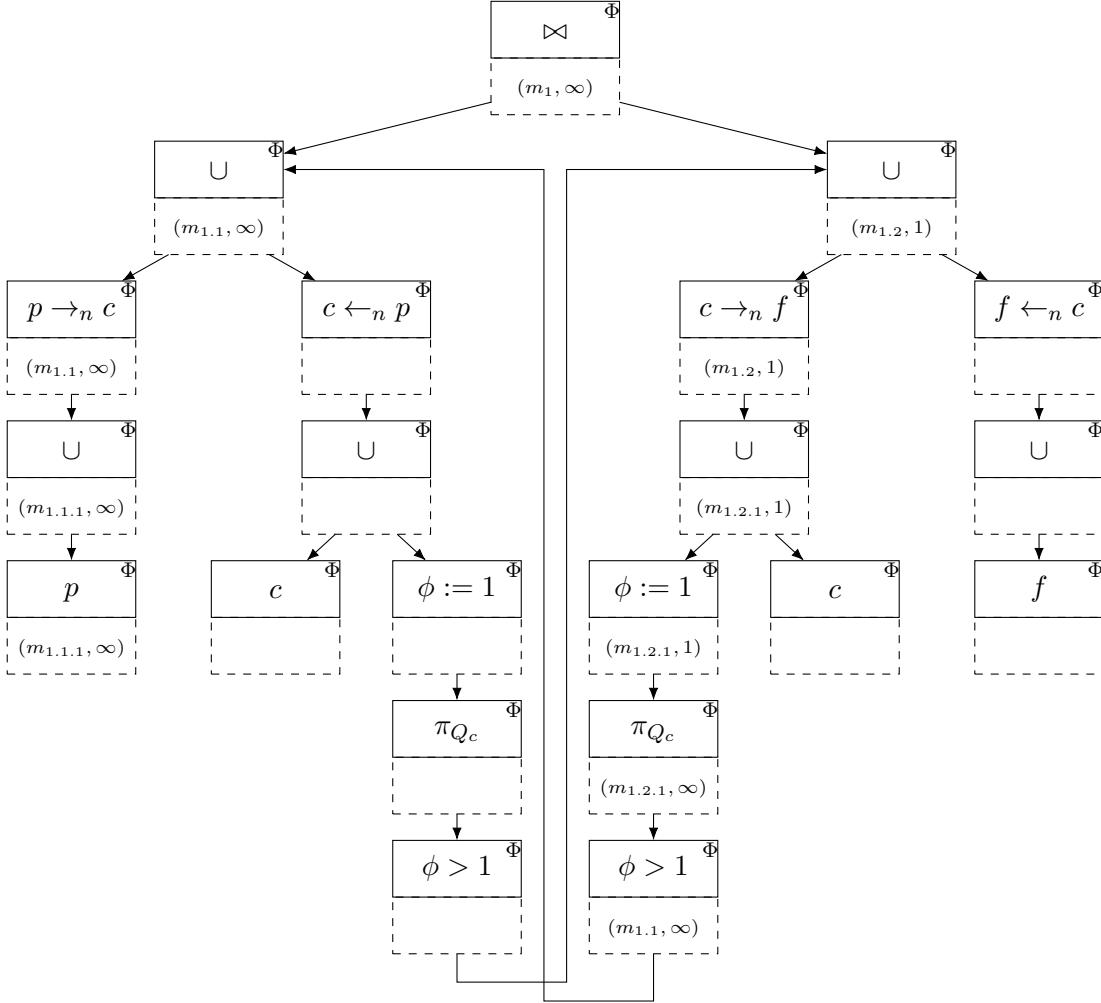


FIGURE 5. Localized RETE net for the example query in Figure 1 with a consistent configuration for the host graph in Figure 6

navigation structure on the right. As a result, the match  $m_{1.2}$  is extracted by the forward navigation node  $[c \rightarrow_n f]^\Phi$ , which is combined with the match  $m_{1.1}$  by the join node  $[\bowtie]^\Phi$  to form the complete match  $m_1$ .

Notably, the host graph contains a second match  $m_2$  for the query pattern, which however does not touch the relevant subgraph. In this example, localization allows the RETE net to completely ignore this second match and skip any related computations.

Moreover, the assignment of the marking 1 to the match  $m_{1.2.1}$  in the request projection structure on the right prevents a propagation of any of the derived intermediate results back into the left local navigation structure. Thereby, potentially problematic cyclic requirement relationships of intermediate results are avoided.

As an example, consider the case where the node  $p_1$  is removed from the relevant subgraph and this change is to be reflected in an updated configuration. Without the additional control of propagation provided by the marking filter and marking assignment nodes, the removal would not mandate an unloading of all associated intermediate results. Specifically,

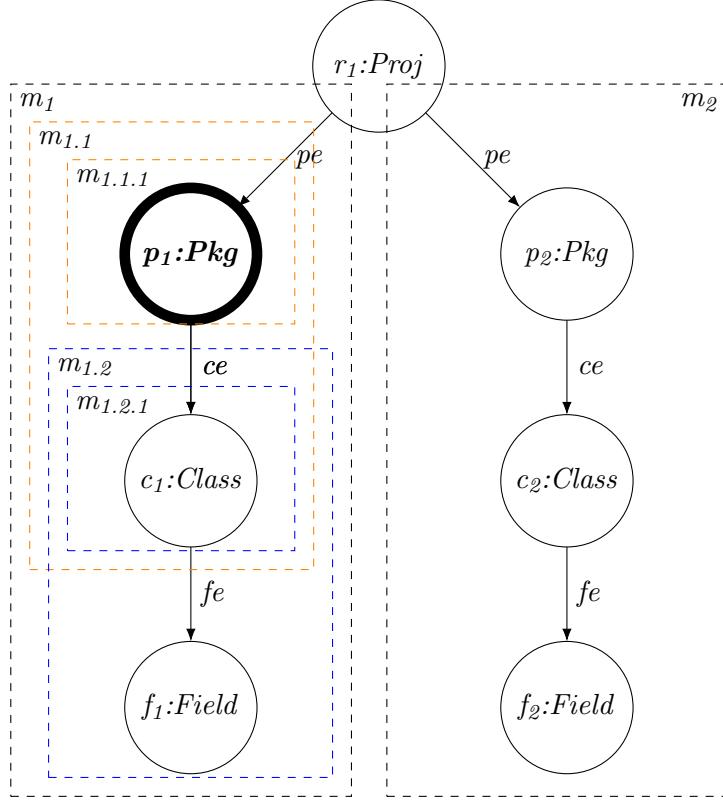


FIGURE 6. Example host graph with relevant subgraph marked in bold

the result set for the union node below the backward navigation node  $[c \leftarrow_n f]^\Phi$  would still contain the match  $m_{1.2.1}$  and the backward navigation node would hence extract the match  $m_{1.1}$  again, leading to the pollution of the RETE net's configuration with superfluous intermediate results.

Formally, a consistent configuration for a localized RETE net then indeed guarantees query results that are complete according to Definition 3.1:

**Theorem 3.5** (Consistent configurations for RETE nets localized via *localize* yield complete query results under the relevant subgraph). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net, and  $Q = \text{query}(p)$ . Furthermore, let  $\mathcal{C}^\Phi$  be a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$ . The set of matches from  $Q$  into  $H$  given by the stripped result set  $\mathcal{R}_X^\Phi(p^\Phi, \mathcal{C}^\Phi)$  is then complete under  $H_p$ .*

*Proof.* (Idea)<sup>3</sup> It can be shown via induction over the height of  $N$  that request projection structures ensure the construction of all intermediate results required to guarantee completeness of the overall result under  $H_p$ . See the proof of Theorem 5 in Appendix D of [BG24b] for a full proof.  $\square$

Interestingly, exactly those matches that touch the relevant subgraph are marked  $\infty$  in a consistent configuration for a localized RETE net, which allows an easy recognition of these matches in the result set of the net's production node:

<sup>3</sup>Detailed proofs for theorems in this article are given in Appendix A or, if indicated, in [BG24b].

**Theorem 3.6** (Matches are marked  $\infty$  in query results of RETE nets localized via *localize* iff they touch the relevant subgraph). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net, and  $Q = \text{query}(p)$ . Furthermore, let  $\mathcal{C}^\Phi$  be a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$ . It then holds that  $\forall (m, \phi) \in \mathcal{C}^\Phi(p^\Phi) : m(Q) \cap H_p \neq \emptyset \Leftrightarrow \phi = \infty$ .*

*Proof.* (Idea) Follows since marking-sensitive joins ultimately combine matches from local navigation structures in  $(N^\Phi, p^\Phi)$ , where only those matches that touch  $H_p$  are marked  $\infty$ .  $\square$

Notably, the insertion of request projection structures creates cycles in the localized RETE net, which prevents execution via a reverse topological sorting. However, the marking filter and marking assignment nodes in the request projection structures effectively prevent cyclic execution at the level of intermediate results: Because matches in the result set of a dependency of some join at height  $h$  that are only computed on request from the other dependency are marked  $h$ , these matches are filtered out in the dependent request projection structure. An execution order for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$  can thus be constructed recursively via an *order* procedure as follows:

If  $p$  is an edge input node, the RETE graph given by  $N^\Phi = \text{LNS}(p)$  is a tree that can be executed via a reverse topological sorting of the nodes in  $\text{LNS}(p)$ , that is,  $\text{order}(N^\Phi) = \text{toposort}(\text{LNS}(p))^{-1}$ .

If  $p$  is a join, according to the construction, the localized RETE graph is given by  $N^\Phi = N_{\bowtie}^\Phi \cup N_l^\Phi \cup N_r^\Phi \cup \text{RPS}_l \cup \text{RPS}_r$ . In this case, an execution order for  $(N^\Phi, p^\Phi)$  can be derived via the concatenation  $\text{order}(N^\Phi) = \text{order}(\text{RPS}_r) \circ \text{order}(N_l^\Phi) \circ \text{order}(\text{RPS}_l) \circ \text{order}(N_r^\Phi) \circ \text{order}(\text{RPS}_r) \circ \text{order}(N_l^\Phi) \circ \text{order}(N_{\bowtie}^\Phi)$ , where  $\text{order}(\text{RPS}_l) = \text{toposort}(\text{RPS}_l)^{-1}$  and  $\text{order}(\text{RPS}_r) = \text{toposort}(\text{RPS}_r)^{-1}$  and  $\text{order}(N_{\bowtie}^\Phi) = [p^\Phi]$ , that is, the sequence containing only  $p^\Phi$ .

Executing a localized RETE net  $(N^\Phi, p^\Phi)$  via  $\text{order}(N^\Phi)$  then guarantees a consistent result configuration for any starting configuration:

**Theorem 3.7** (Execution of localized RETE nets via *order* yields consistent configurations). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net, and  $\mathcal{C}_0^\Phi$  an arbitrary starting configuration. Executing  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$  via  $O = \text{order}(N^\Phi)$  then yields a consistent configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$ .*

*Proof.* (Idea) Follows because the inserted marking filter nodes prevent cyclic execution behavior at the level of intermediate results. See the proof of Theorem 6 in Appendix D of [BG24b] for a full proof.  $\square$

Combined with the result from Theorem 3.5, this means that a localized RETE net can be used to compute complete query results for a relevant subgraph in the sense of Definition 3.1, as outlined in the following corollary:

**Corollary 3.8** (Execution of RETE nets localized via *localize* yields complete query results under the relevant subgraph). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net, and  $Q = \text{query}(p)$ . Furthermore, let  $\mathcal{C}_0^\Phi$  be an arbitrary starting configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$  and  $\mathcal{C}^\Phi = \text{execute}(\text{order}(N^\Phi), N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$ . The set of matches from  $Q$  into  $H$  given by  $\mathcal{R}_X^\Phi(p^\Phi, \mathcal{C}^\Phi)$  is then complete under  $H_p$ .*

*Proof.* Follows directly from Theorem 3.5 and Theorem 3.7.  $\square$

**3.3. Performance of Localized RETE Nets.** Performance of a RETE net  $(N, p)$  with respect to both execution time and memory consumption is largely determined by the *effective size* of a consistent configuration  $\mathcal{C}$  for  $(N, p)$ , which we define as follows:

**Definition 3.9** (Effective Size of RETE Net Configurations). The *effective size* of a configuration  $\mathcal{C}$  for a RETE net  $(N, p)$  is given by  $|\mathcal{C}|_e := \sum_{n \in V^N} \sum_{m \in \mathcal{C}(n)} |m|$ , where we define the size of a match  $m : Q \rightarrow H$  as  $|m| := |m^V| + |m^E| = |V^Q| + |E^Q|$ .

It can then be shown that localization incurs only a constant factor overhead on the effective size of  $\mathcal{C}$  for any edge-dominated host graph.<sup>4</sup>

**Theorem 3.10** (RETE net localization via *localize* introduces only a constant factor overhead on effective configuration size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net with  $Q = \text{query}(p)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . It then holds that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .*

*Proof.* (Idea) Follows because for each RETE node in  $(N, p)$ ,  $(N^\Phi, p^\Phi)$  includes a constant number of nodes whose result sets contain corresponding matches and marking-sensitive result sets contain no duplicate matches.  $\square$

By Theorem 3.10, it then follows that localization of a RETE net incurs only a constant factor overhead on memory consumption even in the worst case where the relevant subgraph is equal to the full model:

**Corollary 3.11** (RETE net localization via *localize* introduces only a constant factor overhead on memory consumption). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . Assuming that storing a match  $m$  requires an amount of memory in  $O(|m|)$  and storing an element from  $\bar{\mathcal{N}}$  requires an amount of memory in  $O(1)$ , storing  $\mathcal{C}^\Phi$  requires memory in  $O(|\mathcal{C}|_e)$ .*

*Proof.* (Idea) Follows from Theorem 3.10.  $\square$

In the worst case, a localized RETE net would still be required to compute a complete result. In this scenario, the execution of the localized RETE net may essentially require superfluous recomputation of match markings, causing computational overhead. When starting with an empty configuration, the number of such recomputations per match is however limited by the size of the query graph, only resulting in a small increase in computational complexity:

**Theorem 3.12** (Execution time overhead introduced by RETE net localization via *localize* depends on query graph size compared to average match size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net for query graph  $Q$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$ , and  $\mathcal{C}_0^\Phi$  the empty configuration for  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$ . Executing*

---

<sup>4</sup>For non-edge-dominated host graphs, the number of matches for (marking-sensitive) vertex input nodes may exceed the number of matches for related edge-input nodes. If matches for marking-sensitive vertex input nodes make up the bulk of intermediate results, localization then introduces an overhead on effective configuration size that cannot be characterized by a constant factor. However, we expect this situation to be rare in practice, since it requires very sparse host graphs.

$(N^\Phi, p^\Phi)$  via  $\text{execute}(\text{order}(N^\Phi), N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  then takes  $O(T \cdot (|Q_a| + |Q|))$  steps, with  $|Q_a|$  the average size of matches in  $\mathcal{C}$  and  $T = \sum_{n \in V^N} |\mathcal{C}(n)|$ .

*Proof.* (Idea) Follows since the effort for initial construction of matches by the marking-sensitive RETE net is linear in the total size of the constructed matches and the marking of a match changes at most  $|Q|$  times. See the proof of Theorem 9 in Appendix D of [BG24b] for a full proof.  $\square$

Assuming an empty starting configuration, a regular well-formed RETE net  $(N, p)$  can be executed in  $O(|\mathcal{C}|_e)$  steps, which can also be expressed as  $O(T \cdot |Q_a|)$ . The overhead of a localized RETE net compared to the original net can thus be characterized by the factor  $\frac{|Q|}{|Q_a|}$ . Assuming that matches for the larger query subgraphs constitute the bulk of intermediate results, which seems reasonable in many scenarios,  $\frac{|Q|}{|Q_a|}$  may be approximated by a constant factor.

For non-empty starting configurations and incremental changes, no sensible guarantees can be made. On the one hand, in a localized RETE net, a host graph modification may trigger the computation of a large number of intermediate results that were previously omitted due to localization. On the other hand, in a standard RETE net, a modification may result in substantial effort for constructing superfluous intermediate results that can be avoided by localization. Depending on the exact host graph structure and starting configuration, execution may thus essentially require a full recomputation for either the localized or standard RETE net but cause almost no effort for the other variant.

#### 4. INCREMENTAL EXTENDED QUERIES OVER SUBGRAPHS

Standard RETE nets can be employed to incrementally compute results for extended graph queries via semi-join and anti-join nodes [BG23]. It is thus desirable to extend the localization technique for simple graph queries presented in Section 3 to also enable the *localized* incremental execution of extended graph queries.

However, multiple interpretations of local correctness of query results for extended graph queries are conceivable depending on the application scenario. More specifically, it is unclear whether this notion should cover only query results where the main match touches the relevant subgraph or also those where elements from the relevant subgraph play a role regarding the satisfaction of the equipped nested graph condition.

Notably, it has been shown that nested graph conditions are equivalent to first order logic [HP09, Ren04] and satisfiability for first order logic is known to be undecidable. A precise solution for the second case thus seems infeasible, as discussed in more detail in Section 4.2. In the following, we will therefore first develop an extension of our localization technique that covers the first case and then proceed by presenting a further extension that covers the second case at a practical level via overapproximation.

**4.1. Localized RETE for Extended Queries.** In the most basic case, a set of matches for an extended graph query  $(Q, \psi)$  may be considered correct under a given relevant subgraph if it contains no matches that violate  $\psi$  and it contains all matches for  $Q$  that satisfy the condition  $\psi$  and involve elements from the relevant subgraph:

**Definition 4.1** (Correctness of Extended Query Results under Subgraphs). We say that a set of matches  $M$  for the extended graph query  $(Q, \psi)$  into a host graph  $H$  is correct under a subgraph  $H_p \subseteq H$  iff  $M \subseteq \{m \in \mathcal{M}_H^Q \mid m \models \psi\}$  and  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge m(Q) \cap H_p \neq \emptyset\} \subseteq M$ .

To enable a localized execution of extended graph queries that meets this definition via the RETE approach, we first have to introduce marking-sensitive versions of the semi-join and anti-join node. These nodes work similarly to the regular versions, but also preserve the marking of matches from their left dependency:

**Definition 4.2** (Target Result Sets of Advanced Marking-Sensitive RETE Nodes). Let  $H$  be a host graph with relevant subgraph  $H_p \subseteq H$ ,  $(N^\Phi, p^\Phi)$  a containing marking-sensitive RETE net, and  $\mathcal{C}^\Phi$  a configuration for  $(N^\Phi, p^\Phi)$ .

- The target result set of a *marking-sensitive semi-join node*  $[\bowtie]^\Phi$  with marking-sensitive dependencies  $n_l^\Phi$  and  $n_r^\Phi$  with  $Q_l = \text{query}(n_l^\Phi)$  and  $Q_r = \text{query}(n_r^\Phi)$  such that  $V^{Q \cap} \neq \emptyset$  for  $Q_\cap = Q_l \cap Q_r$  is given by  $\mathcal{R}^\Phi([\bowtie]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m_l, \phi_l) \mid (m_l, \phi_l) \in \mathcal{C}^\Phi(n_l^\Phi) \wedge \exists (m_r, \phi_r) \in \mathcal{C}^\Phi(n_r^\Phi) : m_l|_{Q_\cap} = m_r|_{Q_\cap}\}$ .
- The target result set of a *marking-sensitive semi-join node*  $[\bowtie]^\Phi$  along a non-empty graph morphism  $a : Q_l \rightarrow Q_r$  with marking-sensitive dependencies  $n_l^\Phi$  and  $n_r^\Phi$  with  $Q_l = \text{query}(n_l^\Phi)$  and  $Q_r = \text{query}(n_r^\Phi)$  is given by  $\mathcal{R}^\Phi([\bowtie]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m_l, \phi_l) \mid (m_l, \phi_l) \in \mathcal{C}^\Phi(n_l^\Phi) \wedge \exists (m_r, \phi_r) \in \mathcal{C}^\Phi(n_r^\Phi) : m_l = m_r \circ a\}$ .
- The target result set of a *marking-sensitive semi-join node*  $[\bowtie]^\Phi$  along a non-empty partial graph morphism  $a : Q_l \rightarrow Q_r$  from some subgraph  $Q_p \subseteq Q_l$  into  $Q_r$  with marking-sensitive dependencies  $n_l^\Phi$  and  $n_r^\Phi$  with  $Q_l = \text{query}(n_l^\Phi)$  and  $Q_r = \text{query}(n_r^\Phi)$  is given by  $\mathcal{R}^\Phi([\bowtie]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m_l, \phi_l) \mid (m_l, \phi_l) \in \mathcal{C}^\Phi(n_l^\Phi) \wedge \exists (m_r, \phi_r) \in \mathcal{C}^\Phi(n_r^\Phi) : m_l|_{Q_p} = m_r \circ a\}$ .
- The target result set of a *marking-sensitive anti-join node*  $[\triangleright]^\Phi$  with marking-sensitive dependencies  $n_l^\Phi$  and  $n_r^\Phi$  with  $Q_l = \text{query}(n_l^\Phi)$  and  $Q_r = \text{query}(n_r^\Phi)$  such that  $V^{Q \cap} \neq \emptyset$  for  $Q_\cap = Q_l \cap Q_r$  is given by  $\mathcal{R}^\Phi([\triangleright]^\Phi, N^\Phi, H, H_p, \mathcal{C}^\Phi) = \{(m_l, \phi_l) \mid (m_l, \phi_l) \in \mathcal{C}^\Phi(n_l^\Phi) \wedge \nexists (m_r, \phi_r) \in \mathcal{C}^\Phi(n_r^\Phi) : m_l|_{Q_\cap} = m_r|_{Q_\cap}\}$ .

Based on the *localize* procedure for plain graph queries from Section 3.2 and our RETE net construction technique for extended graph queries from [BG23], we can then construct a localization procedure *localize*<sup>Ψ</sup> that can be applied to an extended graph query  $(Q, \psi)$  to create a localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$ . *localize*<sup>Ψ</sup> is a recursive procedure, which is visualized in Figure 7 and works as follows:

- For a nested graph condition of the form  $\psi = \text{true}$ , the result of *localize*<sup>Ψ</sup> is given by  $(N^\Phi, p^\Phi) = (N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ <sup>5</sup>.
- For a nested graph condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net for the query  $(N_{(Q', \psi')}^\Phi, p_{(Q', \psi')}^\Phi) = \text{localize}^\Psi(Q', \psi')$ , a request projection structure  $RPS_l^\infty = RPS^\infty(p_Q^\Phi, N_{(Q', \psi')}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\bowtie]^\Phi$  along  $a$  with left dependency  $p_Q^\Phi$  and right dependency  $p_{(Q', \psi')}^\Phi$ .

<sup>5</sup>Here, we use *localize*( $Q$ ) as a shorthand for *localize*( $N, p$ ) with an arbitrary well-formed RETE net  $(N, p)$  for  $Q$ .

- For a nested graph condition of the form  $\psi = \neg\psi'$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net for the query  $(N_{(Q,\psi')}^\Phi, p_{(Q,\psi')}^\Phi) = \text{localize}^\Psi(Q, \psi')$ , a request projection structure  $RPS_l^\infty = RPS^\infty(p_Q^\Phi, N_{(Q,\psi')}^\Phi)$ , and a marking-sensitive anti-join node  $p^\Phi = [\triangleright]^\Phi$  with left dependency  $p_Q^\Phi$  and right dependency  $p_{(Q,\psi')}^\Phi$ .
- For a nested graph condition of the form  $\psi = \psi_1 \wedge \psi_2$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the RETE net  $(N_{(Q,\psi_1)}^\Phi, p_{(Q,\psi_1)}^\Phi) = \text{localize}^\Psi(Q, \psi_1)$ , a request projection structure  $RPS_1^\infty = RPS^\infty(p_Q^\Phi, N_{(Q,\psi_1)}^\Phi)$ , a marking-sensitive semi-join node  $[\times]_1^\Phi$  with left dependency  $p_Q^\Phi$  and right dependency  $p_{(Q,\psi_1)}^\Phi$ , the RETE net  $(N_{(Q,\psi_2)}^\Phi, p_{(Q,\psi_2)}^\Phi) = \text{localize}^\Psi(Q, \psi_2)$ , a request projection structure  $RPS_2^\infty = RPS^\infty([\times]_1^\Phi, N_{(Q,\psi_2)}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\times]_2^\Phi$  with left dependency  $[\times]_1^\Phi$  and right dependency  $p_{(Q,\psi_2)}^\Phi$ .

Note that we include semi-join and anti-join nodes when computing join tree height for determining the value of marking filter nodes in request projection structures. Moreover, by  $RPS^\infty$ , we denote a request projection structure where the marking assignment node assigns a marking of  $\infty$ . While the marking assignment nodes in these request projection structures are technically redundant when considering the constructions in this article and could be removed, we include the nodes to mirror the structure of the request projection structures from Section 3.

Intuitively, by employing localized RETE nets in the computation of nested graph condition satisfaction, it is no longer necessary to compute related matches over the entire host graph. The insertion of request projection structures via  $\text{localize}^\Psi$  then ensures that matches that are required to correctly evaluate a nested graph condition are computed on demand for every relevant match for the context pattern.

The localized RETE net for the extended graph query in Figure 3 is visualized in Figure 8 along with a consistent configuration for the host graph in Figure 9. For local navigation structures and request projection structures, the displayed current result set is the current result set of the topmost union node respectively the marking assignment node of the structure.

Analogously to Figure 3, the localized RETE net essentially extends the RETE net from Figure 5 by a marking sensitive semi-join and a local navigation structure for matching the pattern  $Q'$ . In addition, it contains a request projection structure that controls the computation of matches for  $Q'$  based on matches for  $Q$  that touch the relevant subgraph.

Via this new request projection structure, the match  $m_1$  for  $Q$  triggers the extraction of the match  $m_{1.3}$  for  $Q'$  in the local navigation structure  $LNS^{c \rightarrow i}$ . This ensures that  $m_1$  is included in the result set of the semi-join node at the top of the net, which thereby provides a locally correct query result.

Formally, the following theorem states the correctness of localization via  $\text{localize}^\Psi$  with respect to Definition 4.1:

**Theorem 4.3** (Consistent configurations for RETE nets localized via  $\text{localize}^\Psi$  yield correct query results under the relevant subgraph). *Let  $H$  be a graph,  $H_p \subseteq H$ , and  $(Q, \psi)$  an extended graph query. Furthermore, let  $\mathcal{C}^\Phi$  be a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$ . The set of matches given by the stripped result set  $\mathcal{R}_X^\Phi(p^\Phi, \mathcal{C}^\Phi)$  is then correct under  $H_p$ .*

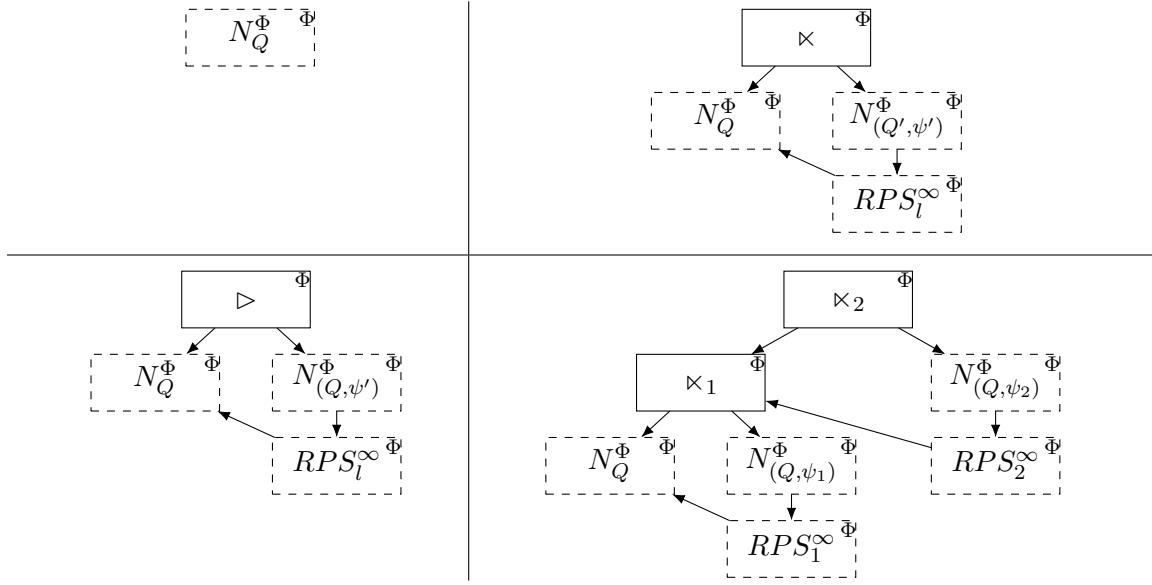


FIGURE 7. Results of applying  $localize^\Psi$  to an extended graph query  $(Q, \psi)$  with nested graph conditions of the form  $\psi = \mathbf{true}$  (top left),  $\psi = \exists(a : Q \rightarrow Q', \psi')$  (top right),  $\psi = \neg\psi'$  (bottom left), and  $\psi = \psi_1 \wedge \psi_2$  (bottom right)

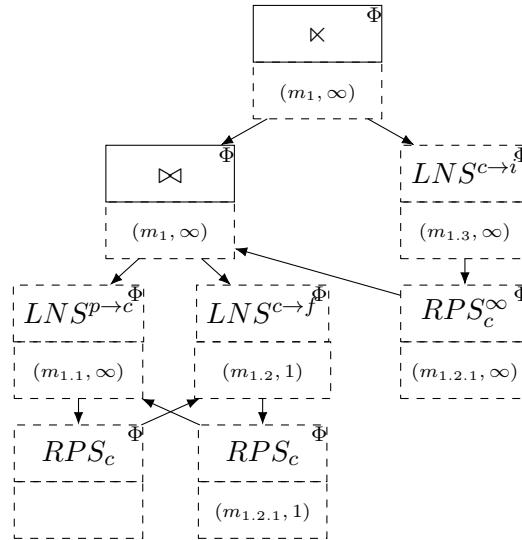


FIGURE 8. Localized RETE net for the example query in Figure 3 with a consistent configuration for the host graph in Figure 9

*Proof.* (Idea) It can be shown via induction over  $\psi$  that the inserted request projection structures trigger the computation of the required intermediate results for nested queries of  $(Q, \psi)$ .  $\square$

An execution order for a localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$  can then be computed via the recursive procedure  $\text{order}^\Psi$ . This procedure works as follows:

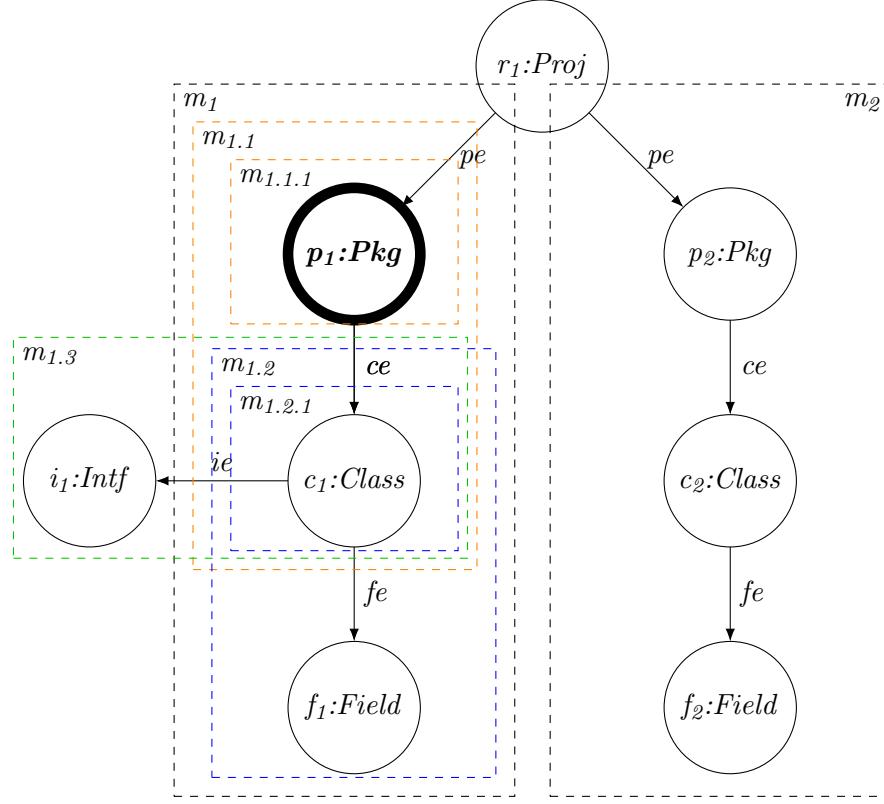


FIGURE 9. Example host graph with relevant subgraph marked in bold

- For a nested graph condition of the form  $\psi = \text{true}$ , the result of  $\text{order}^\Psi$  is simply given by  $\text{order}^\Psi(N^\Phi) = \text{order}(N^\Phi)$ .
- For a nested graph condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ , the result of  $\text{order}^\Psi$  is given by  $\text{order}^\Psi(N^\Phi) = \text{order}(N_Q^\Phi) \circ \text{toposort}(RPS_l^\infty)^{-1} \circ \text{order}^\Psi(N_{(Q', \psi')}^\Phi) \circ [\times]^\Phi$ .
- For a nested graph condition of the form  $\psi = \neg\psi'$ , the result of  $\text{order}^\Psi$  is given by  $\text{order}^\Psi(N^\Phi) = \text{order}(N_Q^\Phi) \circ \text{toposort}(RPS_l^\infty)^{-1} \circ \text{order}^\Psi(N_{(Q, \psi')}^\Phi) \circ [\triangleright]^\Phi$ .
- Finally, for a nested graph condition of the form  $\psi = \psi_1 \wedge \psi_2$ , the result of  $\text{order}^\Psi$  is given by  $\text{order}^\Psi(N^\Phi) = \text{order}(N_Q^\Phi) \circ \text{toposort}(RPS_1^\infty)^{-1} \circ \text{order}^\Psi(N_{(Q, \psi_1)}^\Phi) \circ [\times]_1^\Phi \circ \text{toposort}(RPS_2^\infty)^{-1} \circ \text{order}^\Psi(N_{(Q, \psi_2)}^\Phi) \circ [\times]_2^\Phi$ .

This execution order then guarantees a consistent resulting configuration:

**Theorem 4.4** (Execution of localized RETE nets via  $\text{order}^\Psi$  yields consistent configurations). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  a graph query, and  $\mathcal{C}_0^\Phi$  an arbitrary starting configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$ . Executing  $(N^\Phi, p^\Phi)$  via  $O = \text{order}^\Psi(N^\Phi)$  then yields a consistent configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$ .*

*Proof.* (Idea) Follows because all cyclical structures in  $N^\Phi$  are encapsulated in localized RETE subnets for plain graph queries, for which  $\text{order}$  creates an admissible ordering.  $\square$

It thus follows that a RETE net localized via  $\text{localize}^\Psi$  can be executed to produce a set of matches that is correct under a given subgraph:

**Corollary 4.5** (Execution of RETE nets localized via  $localize^\Psi$  yields correct query results under the relevant subgraph). *Let  $H$  be a graph,  $H_p \subseteq H$ , and  $(Q, \psi)$  a graph query. Furthermore, let  $\mathcal{C}_0^\Phi$  be an arbitrary configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \psi)$  and  $\mathcal{C}^\Phi = execute(order^\Psi(N^\Phi), N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$ . The set of matches from  $Q$  into  $H$  given by  $\mathcal{R}_X^\Phi(p^\Phi, \mathcal{C}^\Phi)$  is then correct under  $H_p$ .*

*Proof.* Follows directly from Theorems 4.4 and 4.3.  $\square$

As  $localize^\Psi$  essentially mirrors the construction procedure for RETE nets for extended graph queries from [BG23], each RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \psi)$  for some extended query  $(Q, \psi)$  has a corresponding regular RETE net  $(N, p)$  created via the procedure described in [BG23] that uses the same RETE net structure for  $Q$  and all patterns in  $\psi$ . As for  $localize$  and well-formed RETE nets for simple queries, we can therefore show that a consistent configuration for  $(N^\Phi, p^\Phi)$  only exhibits a constant-factor increase in configuration size compared to  $(N, p)$  in the worst case:

**Theorem 4.6** (RETE net localization via  $localize^\Psi$  introduces only a constant factor overhead on effective configuration size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \psi)$  corresponding to  $(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . It then holds that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .*

*Proof.* (Idea) Follows because each node in  $N$  has only a constant number of associated nodes with result sets containing corresponding matches in  $N^\Phi$ .  $\square$

Consequently, localization via  $localize^\Psi$  then also incurs at most a constant-factor overhead on memory consumption:

**Corollary 4.7** (RETE net localization via  $localize^\Psi$  introduces only a constant factor overhead on memory consumption). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \psi)$  corresponding to  $(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . Assuming that storing a match  $m$  requires an amount of memory in  $O(|m|)$  and storing an element from  $\bar{N}$  requires an amount of memory in  $O(1)$ , storing  $\mathcal{C}^\Phi$  requires memory in  $O(|\mathcal{C}|_e)$ .*

*Proof.* Follows directly from Theorem 4.6 and the assumptions.  $\square$

Finally, by a similar argumentation as for  $localize$ ,  $localize^\Psi$  also only leads to an increase in computational complexity by factor  $\frac{|Q|}{|Q_a|}$ , with  $|Q_a|$  the average size of matches in a consistent configuration for  $(N, p)$ :

**Theorem 4.8** (Execution time overhead introduced by RETE net localization via  $localize^\Psi$  depends on query graph size compared to average match size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}_0^\Phi$  the empty configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \psi)$  corresponding to  $(N, p)$ . Executing  $(N^\Phi, p^\Phi)$  via  $execute(order^\Psi(N^\Phi), N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  then takes  $O(T \cdot |Q_{max}|)$  steps, with  $T = \sum_{n \in V^N} |\mathcal{C}(n)|$  and  $|Q_{max}|$  the size of the largest graph among  $Q$  and graphs in  $\psi$ .*

*Proof.* (Idea) Follows from Theorem 3.12 and the fact that all semi-joins, anti-joins, and request projection structures not encapsulated in RETE subnets for plain graph queries in  $N^\Phi$  are only executed once.  $\square$

Interestingly, only a subset of the vertex input nodes in a localized RETE net for an extended graph query are required for the net to produce correct results with respect to Definition 4.1. Specifically, all vertex input nodes except those in the local navigation structures of the RETE net for the query pattern created in the topmost recursive call of  $localize^\Psi$  are ultimately superfluous. By removing these superfluous vertex input nodes, localization via  $localize^\Psi$  may then actually improve performance even in cases where the relevant subgraph coincides with the full host graph.

Essentially, the resulting RETE net no longer necessarily computes all matches for all subpatterns of a query's nested graph condition. Instead, match computation for these subpatterns is actually guided by requests for complementarity resulting from matches for the base pattern. Restrictive base patterns with few matches may then create a similar effect to localized execution, with only a subset of matches for subpatterns in nested graph conditions being computed in order to complement existing matches for the base pattern. This is similar to what can be achieved with local-search-based querying techniques but with the advantage of enabling incremental execution. As with localization, the effectiveness however depends on the exact query and host graph.

**4.2. Localized Detection of NGC Satisfaction Changes.** While RETE nets constructed via  $localize^\Psi$  can be used to locally and incrementally compute matches for an extended query that touch a relevant subgraph, the notion of completeness in Definition 4.1 no longer covers the use case of correctly tracking the full effect of local changes on a global query result. This is due to the fact that for extended queries, local changes may result in changes to the satisfaction of a nested graph condition for matches that do not touch the relevant subgraph, effectively manipulating query results that do not fall under Definition 4.1.

To track the effect of local host graph changes on the global set of results for an extended graph query, we thus also have to maintain the set of matches for which a modification of the relevant subgraph may change the nested graph condition's satisfaction. We ideally also want to precompute this set of matches before the host graph is modified to enable incremental maintenance. We thus have no prior knowledge regarding the modifications we may have to consider except for the fact that they concern the relevant subgraph.

This proves problematic when considering a case where the relevant subgraph constitutes the entire host graph as an example. In this scenario, checking whether an arbitrary modification to the relevant subgraph may change the satisfaction of a nested graph condition for some context match corresponds to checking satisfiability of the nested graph condition.

Notably however, nested graph conditions have been shown to be equivalent to first-order logic over graphs [HP09, Ren04]. The satisfiability problem for first-order logic being undecidable thus makes a direct computation of matches for which NGC satisfaction somehow depends on the relevant subgraph impossible in the general case.

In order to still enable a localized tracking of changes to the results of an extended graph query, we therefore employ the following overapproximation for the set of matches where NGC satisfaction may change as a result of local modifications:

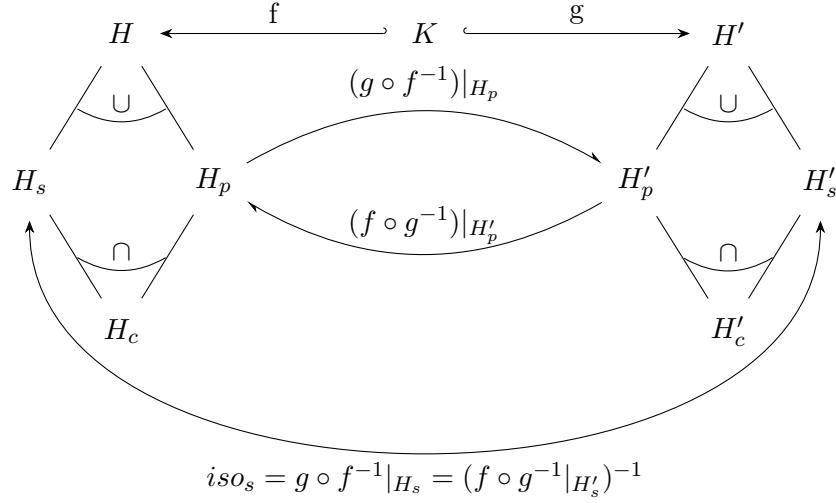


FIGURE 10. Relationship between graphs in Definition 4.10

**Definition 4.9** (Subgraph Satisfaction Dependence). We say that a match  $m : Q \rightarrow H$  for an extended graph query  $(Q, \psi)$  into a host graph  $H$  with relevant subgraph  $H_p \subseteq H$  is *subgraph satisfaction dependent* if

- $\psi = \exists(a : Q \rightarrow Q', \psi')$  and there exists a graph morphism  $m' : Q' \rightarrow H$  such that  $m = m' \circ a$  and  $m'(Q') \cap H_p \neq \emptyset$  or  $m'$  is subgraph satisfaction dependent
- $\psi = \neg\psi'$  and  $m$  is subgraph satisfaction dependent for  $(Q, \psi')$
- $\psi = \psi_1 \wedge \psi_2$  and  $m$  is subgraph satisfaction dependent for  $(Q, \psi_1)$  or  $(Q, \psi_2)$

If  $\psi = \text{true}$ ,  $m$  is never subgraph satisfaction dependent.

We furthermore define a subgraph-restricted, that is, local graph modification based on the basic notion of graph modifications introduced in Section 2.1. In a subgraph-restricted graph modification, changes are limited to relevant subgraphs of source and target graph and relevant subgraphs are mapped onto each other:

**Definition 4.10** (Subgraph-restricted Graph Modification). We say that a graph modification  $H \xleftarrow{f} K \xrightarrow{g} H'$  from a graph  $H$  with relevant subgraph  $H_p$  into a modified graph  $H'$  with relevant subgraph  $H'_p$  is *subgraph-restricted* if  $iso_s = g \circ f^{-1}|_{H_s} = (f \circ g^{-1}|_{H'_s})^{-1}$  is an isomorphism between  $H_s = \{V^H \setminus (V^{H_p} \setminus V^{H_c}), E^H \setminus E^{H_p}, s^H|_{E^H \setminus E^{H_p}}, t^H|_{E^H \setminus E^{H_p}}\}$  and  $H'_s = \{V^{H'} \setminus (V^{H'_p} \setminus V^{H'_c}), E^{H'} \setminus E^{H'_p}, s^{H'}|_{E^{H'} \setminus E^{H'_p}}, t^{H'}|_{E^{H'} \setminus E^{H'_p}}\}$ , where  $V^{H_c} = \{v_c \in V^{H_p} \mid e \in E^H \setminus E^{H_p} : s^H(e) = v_c \vee t^H(e) = v_c\}$  and with  $V^{H'_c}$  defined analogously. Furthermore, it must hold that  $(g \circ f^{-1})(H_p) \subseteq H'_p \wedge (f \circ g^{-1})(H'_p) \subseteq H_p$ .

The relationship between the graphs in the above definition is visualized in the diagram in Figure 10.  $H$  is the source graph of the modification, which is modified into  $H'$  via the intermediate graph  $K$ .  $H_p$  and  $H'_p$  are the relevant subgraphs of  $H$  respectively  $H'$ , which have to be mapped onto each other via  $f \circ g^{-1}$  and  $g \circ f^{-1}$ .  $H_s$  and  $H'_s$  constitute the remainder of  $H$  and  $H'$  and are required to be isomorphic, which means that the subgraph-restricted graph modification cannot make changes outside the relevant subgraph.

Note that  $H_s$  and  $H'_s$  may share nodes with  $H_p$  respectively  $H'_p$  if they are adjacent to edges in  $H \setminus H_p$  respectively  $H' \setminus H'_p$  to avoid dangling edges.  $H_c$  and  $H'_c$  then are the discrete graphs consisting of these shared nodes  $V^{H_c}$  respectively  $V^{H'_c}$  and no edges.

The subgraph satisfaction dependent matches for an extended graph query are then a superset of all matches for which the satisfaction of the query's nested graph condition may change as a result of a subgraph-restricted graph modification. Essentially, a change in nested graph condition satisfaction requires a match to be subgraph satisfaction dependent before or after the host graph modification:

**Theorem 4.11** (Subgraph-restricted graph modifications can only change NGC satisfaction for subgraph satisfaction dependent matches). *Let  $(Q, \psi)$  be an extended graph query and  $H \xleftarrow{f} K \xrightarrow{g} H'$  a subgraph-restricted graph modification from host graph  $H$  with relevant subgraph  $H_p$  into the modified host graph  $H'$  with relevant subgraph  $H'_p$ . It then holds for any graph morphisms  $m_K : Q \rightarrow K$  that  $(f \circ m_K \models \psi \wedge g \circ m_K \not\models \psi) \vee (f \circ m_K \not\models \psi \wedge g \circ m_K \models \psi) \Rightarrow f \circ m_K$  is subgraph satisfaction dependent or  $g \circ m_K$  is subgraph satisfaction dependent.*

*Proof.* (Idea) Can be shown via induction over  $\psi$ , exploiting the fact that host graph modifications can only directly change the satisfaction of existential subconditions.  $\square$

A RETE net for locally computing all subgraph satisfaction dependent matches for an extended graph query  $(Q, \psi)$  can then be constructed recursively via the procedure  $localize^{sat}$ , which is visualized in Figure 11 and works as follows:

- For a nested graph condition of the form  $\psi = true$ ,  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  consists of a single dummy RETE node  $p^{sat} = [\emptyset]^\Phi$ , whose result set is always empty.
- For a nested graph condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ ,  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = localize(Q)$ , the localized RETE net for the plain pattern  $Q'$ ,  $(N_{Q'}^\Phi, p_{Q'}^\Phi) = localize(Q')$ , the RETE net  $(N_{(Q', \psi')}^{sat}, p_{(Q', \psi')}^{sat}) = localize^{sat}(Q', \psi')$ , a marking-sensitive union node  $[\cup]^\Phi$  with dependencies  $p_{Q'}^\Phi$  and  $p_{(Q', \psi')}^\Phi$ , a request projection structure given by  $RPS_r^\infty = RPS^\infty([\cup]^\Phi, N_Q^\Phi)$ , and a marking-sensitive semi-join node  $p^{sat} = [\times]^\Phi$  along  $a$  with left dependency  $p_Q^\Phi$  and right dependency  $[\cup]^\Phi$ .
- For a nested graph condition of the form  $\psi = \neg\psi'$ ,  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi) = localize^{sat}(Q, \psi')$ .
- For a nested graph condition of the form  $\psi = \psi_1 \wedge \psi_2$ ,  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  consists of the RETE nets  $(N_{(Q, \psi_1)}^{sat}, p_{(Q, \psi_1)}^{sat}) = localize^{sat}(Q, \psi_1)$  and  $(N_{(Q, \psi_2)}^{sat}, p_{(Q, \psi_2)}^{sat}) = localize^{sat}(Q, \psi_2)$  and a marking-sensitive union node  $p^{sat} = [\cup]^\Phi$  with dependencies  $p_{(Q, \psi_1)}^{sat}$  and  $p_{(Q, \psi_2)}^{sat}$ .

A localized RETE net for computing subgraph-satisfaction dependent matches for the example query in Figure 3 is displayed in Figure 12 alongside a consistent configuration for the host graph in Figure 13. It consists of the localized RETE net in Figure 5 and the structure created for handling the nested graph condition according to Figure 11.

Notably, there exist no matches for the base pattern  $Q$  that touch the relevant subgraph, which only consists of the interface node  $i_1$ . Instead, the computation of matches for  $Q$  is controlled via the request projection structure that feeds matches for the right dependency of the semi-join node into the localized RETE net for  $Q$ .

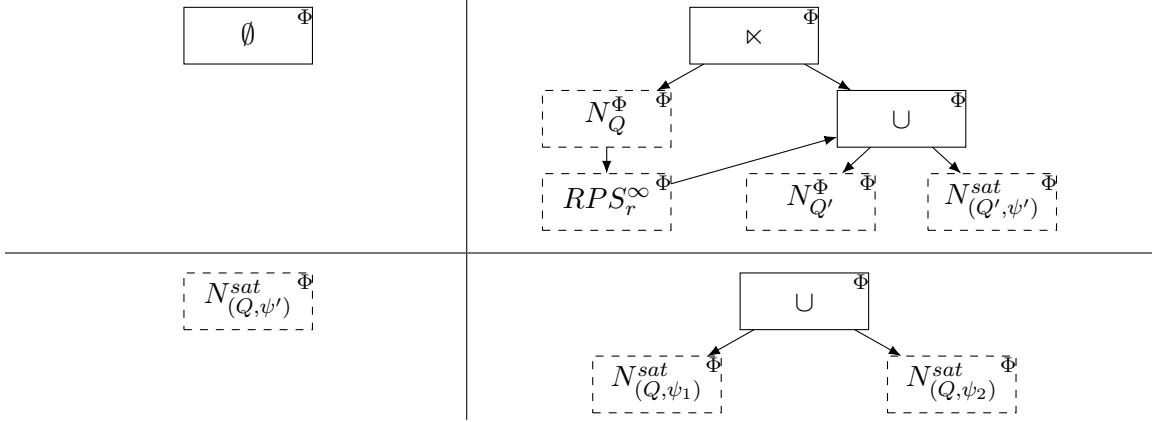


FIGURE 11. Results of applying  $localize^{sat}$  to an extended graph query  $(Q, \psi)$  with nested graph conditions of the form  $\psi = \text{true}$  (top left),  $\psi = \exists(a : Q \rightarrow Q', \psi')$  (top right),  $\psi = \neg\psi'$  (bottom left), and  $\psi = \psi_1 \wedge \psi_2$  (bottom right)

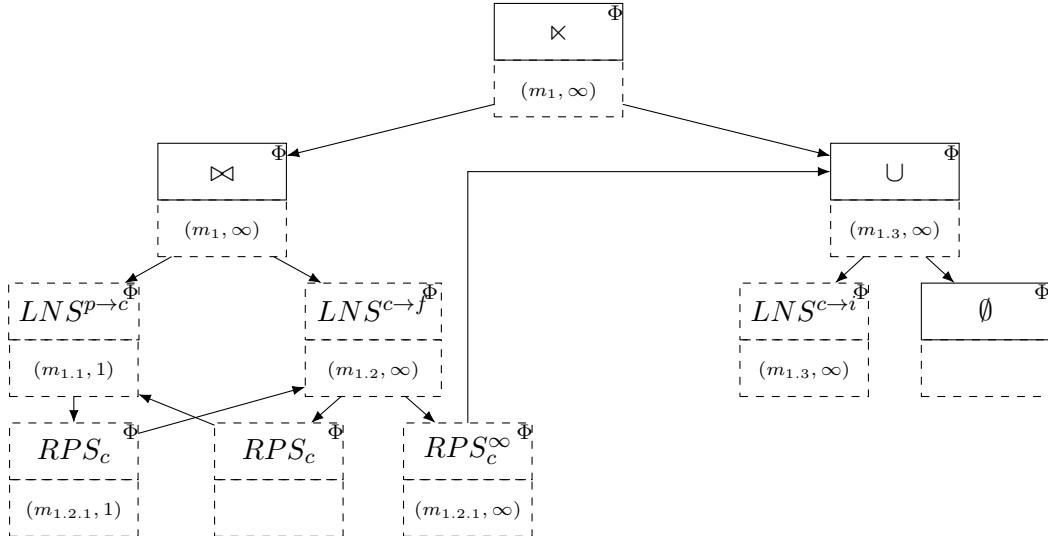


FIGURE 12. Localized RETE net for computing subgraph-satisfaction dependent matches for the example query in Figure 3 with a consistent configuration for the host graph in Figure 13

Compared to Figure 8, propagation of matches is essentially reversed: First, the local navigation structure  $LNS^{c \rightarrow i}$  extracts all matches for the pattern  $Q'$  that touch the relevant subgraph. The request projection structure connected to the union node then causes the local navigation structure  $LNS^{c \rightarrow f}$  to extract the match  $m_{1.2}$ . The leftmost request projection structure finally ensures that the required complementary match  $m_{1.1}$  is extracted by the local navigation structure  $LNS^{p \rightarrow c}$ .

A consistent configuration for the resulting RETE net then indeed guarantees a complete query result with respect to Definition 4.9:

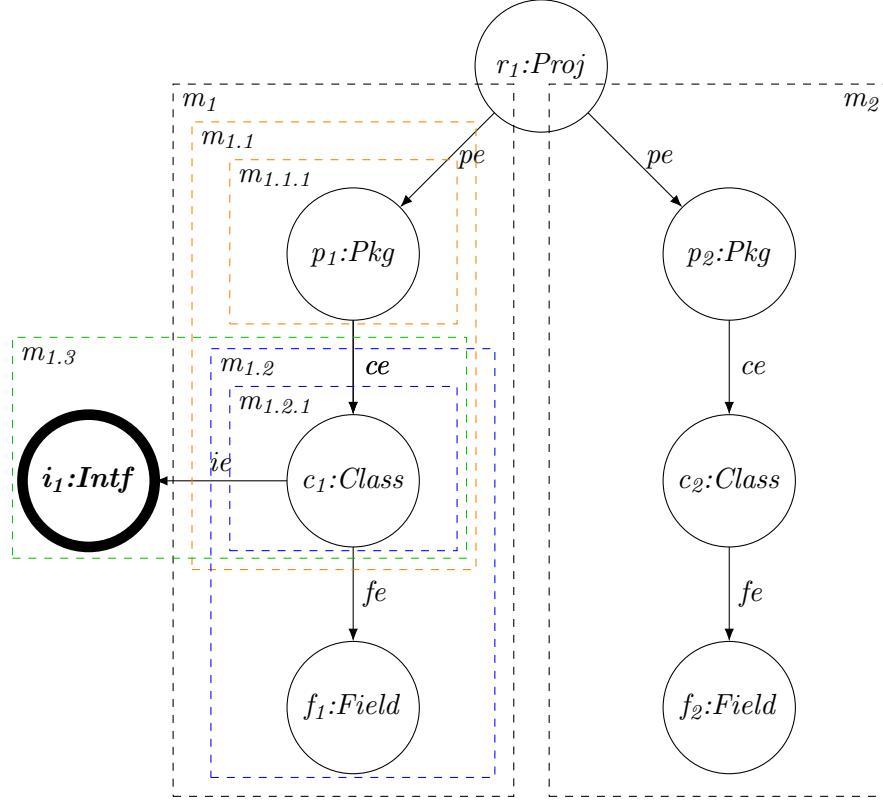


FIGURE 13. Example host graph with relevant subgraph marked in bold

**Theorem 4.12** (Consistent configurations for RETE nets localized via  $localize^{sat}$  yield all subgraph satisfaction dependent matches). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  an extended graph query, and  $\mathcal{C}^\Phi$  a consistent configuration for the RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$ . It then holds that  $\forall m \in \mathcal{M}_H^Q : m$  is subgraph satisfaction dependent  $\Rightarrow m \in \mathcal{R}_X^\Phi(p^{sat}, \mathcal{C}^\Phi)$ .*

*Proof.* (Idea) Can be shown via induction over  $\psi$ , exploiting the close alignment of  $localize^{sat}$  with Definition 4.9.  $\square$

Given an extended graph query  $(Q, \psi)$  and the RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$ , an execution order for  $(N^{sat}, p^{sat})$  can be constructed in a straightforward manner via the procedure  $order^{sat}$ :

- For a nested graph condition of the form  $\psi = \text{true}$ , the result of  $order^{sat}$  is given by  $order^{sat}(N^{sat}) = [\emptyset]^\Phi$ .
- For a nested graph condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ , the result of  $order^{sat}$  is given by  $order^{sat}(N^{sat}) = order(N_{Q'}^\Phi) \circ order^{sat}(N_{(Q', \psi')}^{sat}) \circ [\cup]^\Phi \circ toposort(RPS_r^\infty)^{-1} \circ order(N_Q^\Phi) \circ [\times]^\Phi$ .
- For a nested graph condition of the form  $\psi = \neg\psi'$ ,  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi')$  and thus the ordering via  $order^{sat}(N^{sat})$  is covered by one of the other cases.
- For a nested graph condition of the form  $\psi = \psi_1 \wedge \psi_2$ , the result of  $order^{sat}$  is given by  $order^{sat}(N^{sat}, p^{sat}) = order^{sat}(N_{(Q, \psi_1)}^{sat}) \circ order^{sat}(N_{(Q, \psi_2)}^{sat}) \circ [\cup]^\Phi$ .

This execution order then guarantees consistency of the resulting configuration:

**Theorem 4.13** (Execution of localized RETE nets via  $order^{sat}$  yields consistent configurations). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  a graph query, and  $\mathcal{C}_0^\Phi$  an arbitrary starting configuration for the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$ . Executing  $(N^{sat}, p^{sat})$  via  $O = order^{sat}(N^{sat})$  then yields a consistent configuration  $\mathcal{C}^\Phi = execute(O, N^{sat}, H, H_p, \mathcal{C}_0^\Phi)$ .*

*Proof.* (Idea) Follows because all cyclical structures in  $N^{sat}$  are encapsulated in localized RETE subnets for plain graph queries, for which  $order$  creates an admissible ordering.  $\square$

We hence again get a corollary stating the correctness of RETE net execution via  $localize^{sat}$ :

**Corollary 4.14** (Execution of RETE nets localized via  $localize^{sat}$  yields all subgraph satisfaction dependent matches). *Let  $H$  be a graph,  $H_p \subseteq H$ , and  $(Q, \psi)$  a graph query. Furthermore, let  $\mathcal{C}_0^\Phi$  be an arbitrary starting configuration for the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  and  $\mathcal{C}^\Phi = execute(order^{sat}(N^{sat}), N^{sat}, H, H_p, \mathcal{C}_0^\Phi)$ . It then holds that  $\forall m \in \mathcal{M}_H^Q : m \text{ is subgraph satisfaction dependent} \Rightarrow m \in \mathcal{R}_X^\Phi(p^{sat}, \mathcal{C}^\Phi)$ .*

*Proof.* Follows directly from Theorems 4.12 and 4.13.  $\square$

Technically, it can be shown that a RETE net for an extended graph query  $(Q, \psi)$  constructed via  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  has similar performance characteristics as a RETE net constructed via  $localize^\Psi$  compared to a regular RETE net  $(N, p)$  for  $(Q, \psi)$  constructed via the basic procedure described in [BG23]. Specifically, the size of a consistent configuration for  $(N^{sat}, p^{sat})$  is only increased by a constant factor compared to a consistent configuration for  $(N, p)$ , resulting in only a constant-factor increase in memory consumption and a minor increase in runtime complexity.<sup>6</sup>

However, it is important to note that  $(N, p)$  and  $(N^{sat}, p^{sat})$  potentially compute very different sets of matches. By Definition 4.9, RETE nets constructed via  $localize^{sat}$  may have to compute matches for  $Q$  that do not actually satisfy  $\psi$  if  $\psi$  involves negations or conjunctions, whereas such matches have to be excluded from the results of  $(N, p)$ .

Relating the performance of  $(N^{sat}, p^{sat})$  to the performance of  $(N, p)$  therefore relies on redundancy in the basic construction in [BG23]. Many optimizations to  $(N, p)$ , such as the elimination of redundancy or the downward propagation of nested graph conditions in the RETE net described in [BG23], are applicable to localized RETE nets constructed via  $localize^\Psi$  and likely yield similar benefits. In contrast, due to the different construction procedure and desired computation results, such optimizations can have a very different effect on  $(N^{sat}, p^{sat})$  or may simply not be applicable at all. Consequently, executing  $(N^{sat}, p^{sat})$  as part of any execution of  $(Q, \psi)$  may create substantial overhead on both memory consumption and execution time in unfavorable cases if further optimizations are involved.

Still, using the  $localize^{sat}$  procedure, we can now construct a RETE net that allows the localized computation of changes to the results of an extended graph query  $(Q, \psi)$  caused by a subgraph-restricted graph modification  $H \xleftarrow{f} K \xrightarrow{g} H'$ , which may be beneficial in practice.

---

<sup>6</sup>See Theorem A.1, Corollary A.2, and Theorem A.3 in Appendix A.

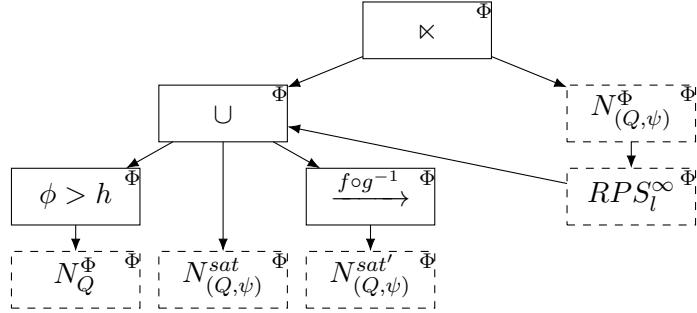


FIGURE 14. Result of applying  $localize^{\Delta}$  to an extended graph query  $(Q, \psi)$

The result  $(N^{\Delta}, p^{\Delta}) = localize^{\Delta}(Q, \psi)$  of this construction via the procedure  $localize^{\Delta}$  is visualized in Figure 14. The resulting RETE net consists of the localized version  $(N_Q^{\Phi}, p_Q^{\Phi}) = localize(N^Q, p^Q)$  of a regular RETE net  $(N^Q, p^Q)$  with height  $h$  for the plain graph query  $Q$ , a marking filter node  $[\phi > h]^{\Phi}$  with dependency  $p_Q^{\Phi}$ , the RETE net  $(N_{(Q, \psi)}^{sat}, p_{(Q, \psi)}^{sat}) = localize^{sat}(Q, \psi)$ , and a marking-sensitive union node  $[\cup]^{\Phi}$  with dependencies  $[\phi > h]^{\Phi}$  and  $p_{(Q, \psi)}^{sat}$ . In addition,  $N^{\Delta}$  comprises the RETE net  $(N_{(Q, \psi)}^{\Phi}, p_{(Q, \psi)}^{\Phi}) = localize^{\Psi}(Q, \psi)$ , a request projection structure  $RPS_l^{\infty} = RPS^{\infty}([\cup]^{\Phi}, N_{(Q, \psi)}^{\Phi})$ , and a marking-sensitive semi-join  $[\times]^{\Phi}$  with left dependency  $[\cup]^{\Phi}$  and right dependency  $p_{(Q, \psi)}^{\Phi}$ .

Additionally, another copy of the RETE net for computing subgraph satisfaction dependent matches,  $(N_{(Q, \psi)}^{sat'}, p_{(Q, \psi)}^{sat'}) = localize^{sat}(Q, \psi)$ , is created for computing subgraph satisfaction dependent matches in  $H'$ . These matches then have to be translated into matches into  $H$  along the graph morphisms  $f$  and  $g$  via a dedicated *marking-sensitive translation node*  $[\xrightarrow{f \circ g^{-1}}]^{\Phi}$  with dependency  $p_{(Q, \psi)}^{sat'}$ , which is also added to  $N^{\Delta}$ . The target result set of this marking-sensitive translation node is given by  $\mathcal{R}^{\Phi}([\xrightarrow{f \circ g^{-1}}]^{\Phi}, N^{\Delta}, H, H_p, \mathcal{C}^{\Phi}) = \{(m, \phi) | m \in \mathcal{M}_H^Q \wedge \exists (m', \phi) \in \mathcal{C}^{\Phi}(p^{sat'}) : f \circ g^{-1} \circ m' = m\}$ . The node  $[\xrightarrow{f \circ g^{-1}}]^{\Phi}$  can then be added as an additional dependency to  $[\cup]^{\Phi}$ .

The created net thereby computes the matches into  $H$  that satisfy  $\psi$  and touch the relevant subgraph or are subgraph satisfaction dependent in  $H$  or  $H'$ .

**Theorem 4.15** (For a subgraph-restricted graph modification, consistent configurations for RETE nets localized via  $localize^{\Delta}$  yield all matches that touch the relevant subgraph in the source or are subgraph satisfaction dependent in the source or target of the modification). *Let  $(Q, \psi)$  be an extended graph query and  $H \xleftarrow{f} K \xrightarrow{g} H'$  a subgraph-restricted graph modification modifying the graph  $H$  with relevant subgraph  $H_p \subseteq H$  into the graph  $H'$  with relevant subgraph  $H'_p \subseteq H'$ . Furthermore, let  $\mathcal{C}^{\Phi}$  be a configuration that is consistent for  $(N^{\Delta}, p^{\Delta}) = localize^{\Delta}(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  and consistent for  $(N_{(Q, \psi)}^{sat'}, p_{(Q, \psi)}^{sat'}) = localize^{sat}(Q, \psi)$  for host graph  $H'$  and relevant subgraph  $H'_p$ . It must then hold that  $\forall m \in \mathcal{M}_H^Q : m \models \psi \wedge (m(Q) \cap$*

$H_p \neq \emptyset$  or  $m$  is subgraph satisfaction dependent or  $\exists m' \in \mathcal{M}_{H'}^Q : f \circ g^{-1} \circ m' = m \wedge m'$  is subgraph satisfaction dependent)  $\Rightarrow m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi)$ .

*Proof.* (Idea) Follows from Theorems 3.5, 4.3 and 4.12.  $\square$

An execution order for a RETE net  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$  can simply be constructed as  $\text{order}^\Delta(N^\Delta) = \text{order}(N_Q^\Phi) \circ [\phi > h]^\Phi \circ \text{order}^{\text{sat}}(N_{(Q,\psi)}^{\text{sat}}) \circ [\xrightarrow{f \circ g^{-1}}]^\Phi \circ [\cup]^\Phi \circ \text{toposort}(RPS_l^\infty)^{-1} \circ \text{order}^\Psi(N_{(Q,\psi)}^\Phi) \circ [\bowtie]^\Phi$ :

**Theorem 4.16** (Execution of localized RETE nets via  $\text{order}^\Delta$  yields consistent configurations). *Let  $(Q, \psi)$  be an extended graph query and  $H$  a host graph with relevant subgraph  $H_p$  and  $\mathcal{C}_0^\Phi$  an arbitrary starting configuration for the marking-sensitive RETE net  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$ . Executing  $(N^\Delta, p^\Delta)$  via  $O = \text{order}^\Delta(N^\Delta)$  then yields a consistent configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Delta, H, H_p, \mathcal{C}_0^\Phi)$ .*

*Proof.* (Idea) Follows because all cyclical structures in  $N^\Delta$  are encapsulated in localized RETE subnets for plain graph queries, for which  $\text{order}$  creates an admissible ordering.  $\square$

As usual, we thus get a corollary stating the correctness of RETE net execution via  $\text{order}^\Delta$ :

**Corollary 4.17** (For a subgraph-restricted graph modification, execution of RETE nets localized via  $\text{localize}^\Delta$  yields all matches that touch the relevant subgraph in the source or are subgraph satisfaction dependent in the source or target graph of the modification). *Let  $(Q, \psi)$  be an extended graph query and  $H \xleftarrow{f} K \xrightarrow{g} H'$  a subgraph-restricted graph modification modifying the graph  $H$  with relevant subgraph  $H_p \subseteq H$  into the graph  $H'$  with relevant subgraph  $H'_p \subseteq H'$ . Furthermore, let  $\mathcal{C}_0^\Phi$  be a configuration for  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$  and  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'}) = \text{localize}^{\text{sat}}(Q, \psi)$  that is consistent for  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'})$  for host graph  $H'$  and relevant subgraph  $H'_p$ . It then holds for the configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Delta, H, H_p, \mathcal{C}_0^\Phi)$  that  $\forall m \in \mathcal{M}_H^Q : m \models \psi \wedge (m(Q) \cap H_p \neq \emptyset$  or  $m$  is subgraph satisfaction dependent or  $\exists m' \in \mathcal{M}_{H'}^Q : f \circ g^{-1} \circ m' = m \wedge m'$  is subgraph satisfaction dependent)  $\Rightarrow m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi)$ .*

*Proof.* Follows directly from Theorem 4.15 and Theorem 4.16.  $\square$

Changes to the results of an extended graph query  $(Q, \psi)$  resulting from a subgraph-restricted graph modification  $H \xleftarrow{f} K \xrightarrow{g} H'$  for relevant subgraphs  $H_p \subseteq H$  and  $H'_p \subseteq H'$  can then be computed by mirroring the construction with  $\text{localize}^\Delta$  for  $H$  and  $H'$ . This entails creating RETE nets  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$  and  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'}) = \text{localize}^{\text{sat}}(Q, \psi)$  as well as duplicates  $(N^{\Delta'}, p^{\Delta'}) = \text{localize}^\Delta(Q, \psi)$  and  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'}) = \text{localize}^{\text{sat}}(Q, \psi)$ .  $(N^\Delta, p^\Delta)$  and  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'})$  are executed over  $H$ , whereas  $(N^{\Delta'}, p^{\Delta'})$  and  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'})$  are executed over  $H'$ . After appropriate translation of matches, the set of matches that are deleted by the modification or for which the satisfaction of  $\psi$  changes from `true` to `false` can be computed as the difference between results for  $p^\Delta$  and  $p^{\Delta'}$ :

**Theorem 4.18** (Removal of matches from the results of an extended graph query caused by a subgraph-restricted graph modification can be detected using RETE nets localized via  $\text{localize}^\Delta$ ). *Let  $(Q, \psi)$  be an extended graph query and  $H \xleftarrow{f} K \xrightarrow{g} H'$  a subgraph-restricted graph modification modifying the graph  $H$  with relevant subgraph  $H_p \subseteq H$  into*

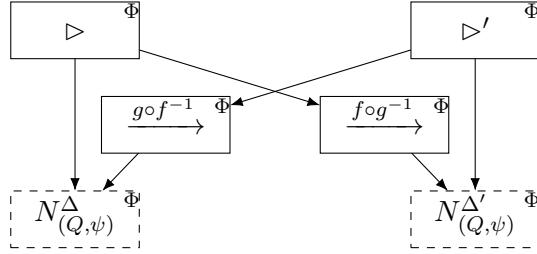


FIGURE 15. RETE construction for locally computing changes in results for an extended graph query  $(Q, \psi)$

the graph  $H'$  with relevant subgraph  $H'_p \subseteq H'$ . Furthermore, let  $\mathcal{C}^\Phi$  be a configuration that is consistent for  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$  and  $(N_{(Q, \psi)}^{\text{sat}}, p_{(Q, \psi)}^{\text{sat}}) = \text{localize}^{\text{sat}}(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  and consistent for  $(N^{\Delta'}, p^{\Delta'}) = \text{localize}^{\Delta'}(Q, \psi)$  and  $(N_{(Q, \psi)}^{\text{sat}'}, p_{(Q, \psi)}^{\text{sat}'}) = \text{localize}^{\text{sat}'}(Q, \psi)$  for host graph  $H'$  and relevant subgraph  $H'_p$ . It then holds that  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge \nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi\} = \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

*Proof.* (Idea) Inclusion in both directions can be shown by exploiting the fact that ultimately, both  $(N^\Delta, p^\Delta)$  and  $(N^{\Delta'}, p^{\Delta'})$  check the satisfaction of  $\psi$  for sets of matches that can only differ in matches touching the relevant subgraph.  $\square$

Analogously, matches created by the modification or for which the satisfaction of  $\psi$  changes from **false** to **true** can be computed as the difference between  $p^{\Delta'}$  and  $p^\Delta$ . The required translation of matches can be realized in a straightforward manner via marking-sensitive translation nodes and difference computation can be performed via marking-sensitive anti-joins, as displayed in Figure 15.

## 5. EVALUATION

We aim to investigate whether RETE net localization can improve performance of query execution in scenarios where the relevant subgraph constitutes only a fraction of the full model, considering initial query execution time, execution time for incrementally processing model updates, and memory consumption as performance indicators. We experiment<sup>7</sup> with the following querying techniques:

- **STANDARD:** Our own implementation of a regular RETE net with global execution semantics [BG23].
- **LOCALIZED:** Our own implementation of the RETE net used for STANDARD, localized according to the description in Section 3.2 and Section 4.1.

<sup>7</sup>Experiments were executed on a Linux SMP Debian 4.19.67-2 computer with Intel Xeon E5-2630 CPU (2.3 GHz clock rate) and 386 GB system memory running OpenJDK 11.0.6. Reported measurements correspond to the arithmetic mean of measurements for 10 runs. Memory measurements were obtained using the Java Runtime class. To represent graph data, all experiments use EMF [emf] object graphs that enable reverse navigation of edges. Our implementation is available under [BG25]. More details and query visualizations can be found in [BG24b] and [BG25] or in Appendix B.

- **DELTA**: Our own implementation of the RETE net used for STANDARD, localized for detecting changes to the results of extended graph queries according to the description in Section 4.2.
- **VIATRA**: The external RETE-based VIATRA tool [VBH<sup>+</sup>16].
- **SDM\***: Our own local-search-based Story Diagram Interpreter tool [GHS09].

Note that STANDARD and LOCALIZED implement the additional optimizations for elimination of redundancy and downward propagation of nested graph conditions described in [BG23]. To the extent possible, DELTA employs these optimizations as well.

For SDM\*, we only consider searching for new query results. We thus underapproximate the time and memory required for a full solution with this strategy, which would also require maintaining previously found results.

**5.1. Plain Graph Queries over Synthetic Abstract Syntax Graphs.** We first attempt a systematic evaluation via a synthetic experiment, which emulates a developer loading part of a large model into their workspace and monitoring some well-formedness constraints as they modify the loaded part, that is, relevant subgraph, without simultaneous changes to other model parts.

We therefore generate Java abstract syntax graphs with 1, 10, 100, 1000, and 10000 packages, with each package containing 10 classes with 10 fields referencing other classes in the same or a different package. As relevant subgraph, we consider a single package and its contents. We then execute a plain graph query searching for paths consisting of a package and four classes connected via fields.

After the initial query execution, we modify the relevant subgraph by creating a class in the considered package along with 10 fields referencing other existing classes in the relevant subgraph, that is, the same package. We then perform an incremental update of query results. This step of modifying the relevant subgraph and updating query results is repeated 10 times.

Due to the query being a simple path, the effort for query execution largely depends on the number of overall query results. This number is in turn directly determined by the number of classes considered during execution and their number of connections to other classes. By fixing the former for localized execution via the selection of the relevant subgraph and fixing the latter for both localized and global execution via the construction of the host graph, we thus aim to isolate the effect of localization on the performance of query execution.

Figure 16 (left) displays the execution times for the initial execution of the query. The execution time of LOCALIZED remains around 120 ms for all model sizes. The execution time for SDM\* slowly grows from around 350 ms to 1025 ms due to indexing effort that is necessary for observing model changes. In contrast, the execution time for the other RETE-based strategies clearly scales with model size, with the execution time for STANDARD growing from around 13 ms for the smallest model to more than 184 000 ms for the largest model. On the one hand, localization thus incurs a noticeable overhead in initial execution time for the smallest model, where even localized query execution is essentially global. On the other hand, it significantly improves execution time for the larger models and even achieves better scalability than the local-search-based tool in this scenario.

The average times for processing a model update are displayed in Figure 16 (center). Here, all strategies achieve execution times mostly independent of model size. While the

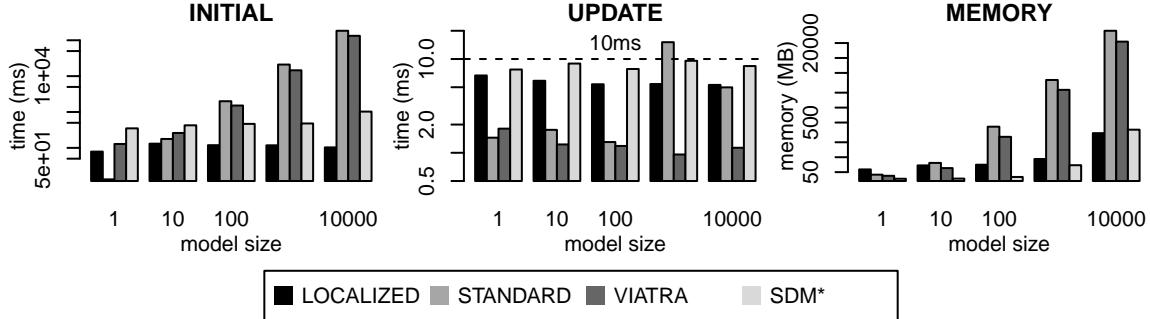


FIGURE 16. Measurements for the synthetic abstract syntax graph scenario (log scale)

measurements for STANDARD fluctuate, likely due to the slightly unpredictable behavior of hash-based indexing structures, average execution times remain low overall and below 10 ms for LOCALIZED. Still, localization incurs a noticeable overhead up to factor 6 compared to STANDARD and VIATRA. This overhead is expected, since in this scenario, all considered updates affect the relevant subgraph and thus impact the results of the localized RETE net similarly to the results of the standard RETE nets. Consequently, localization does not reduce computational effort, but causes overhead instead.

Finally, Figure 16 (right) shows the memory measurements for all strategies and models after the final update. Here, LOCALIZED again achieves a substantial improvement in scalability compared to the other RETE-based strategies, with a slightly higher memory consumption for the smallest model and an improvement by factor 120 over STANDARD for the largest model. This is a result of the localized RETE net producing the same number of intermediate results for all model sizes, with the slight growth in memory consumption likely a product of the growing size of the model itself. SDM\*, not storing any matches, performs better for all but the largest model, where memory consumption surpasses the measurement for LOCALIZED. This surprising result can probably be explained by the fact that SDM\* has to index the full model to observe modifications, causing an overhead in memory consumption.

In addition to these experiments, where the full model is always stored in main memory, we also experimented with a model initially stored on disk via the persistence layer CDO [cdo]. Measurements mostly mirror those for the main-memory-based experiment. Notably though, in conjunction with CDO the LOCALIZED strategy achieves almost ideal scalability regarding memory consumption for this scenario, with measurements around 70 MB for all model sizes.<sup>8</sup>

**5.2. Plain Graph Queries over Real Abstract Syntax Graphs.** To evaluate our approach in the context of a more realistic application scenario, we perform a similar experiment using real data and plain graph queries inspired by object-oriented design patterns [GHJV93]. In contrast to the synthetic scenario, this experiment emulates a situation where modifications may concern not only the relevant subgraph but the entire

<sup>8</sup>See [BG24b] for measurement results.

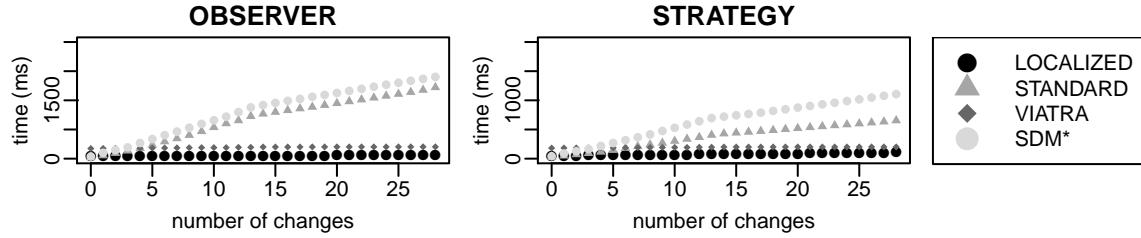


FIGURE 17. Execution times for the real abstract syntax graph scenario

model, for instance when multiple developers are simultaneously working on different model parts.

We therefore extract a history of real Java abstract syntax graphs with about 16 000 vertices and 45 000 edges from a software repository using the MoDisco tool [BCJM10, BCDM14]. After executing the queries over the initial commit, we replay the history and perform incremental updates of query results after each commit. The employed queries are visualized in Appendix B. As relevant subgraph, we again consider a single package and its contents.

Figure 17 displays the aggregate execution time for processing the commits one after another for the queries where LOCALIZED performed best and worst compared to STANDARD, with the measurement at  $x = 0$  indicating the initial execution time for the starting model. Initial execution times are similarly low due to a small starting model and in fact slightly higher for LOCALIZED. However, on aggregate LOCALIZED outperforms STANDARD with an improvement between factor 5 and 18 due to significantly lower update times, which are summarized in Figure 18 (left).

In this case, the improvement mostly stems from the more precise monitoring of the model for modifications: The RETE nets of both STANDARD and LOCALIZED remain small due to strong filtering effects in the query graphs. However, while STANDARD spends significant effort on processing model change notifications due to observing all appropriately typed model elements, this effort is substantially reduced for LOCALIZED, which only monitors elements relevant to query results required for completeness under the relevant subgraph. The execution times of SDM\* can be explained by the same effect. Interestingly, VIATRA seems to implement an improved handling of such notifications, achieving improved execution times for particularly small updates even compared to LOCALIZED, but requiring more time if an update triggers changes to the RETE net. Combined with a higher RETE net initialization time, this results in LOCALIZED also outperforming VIATRA for all considered queries.

Regarding memory consumption, all strategies perform very similarly, which is mostly a result of the size of the model itself dominating the measurement and hiding the memory impact of the rather small RETE nets.

**5.3. LDBC Social Network Benchmark.** Finally, we also perform experiments inspired by the independent LDBC Social Network Benchmark [EALP<sup>+</sup>15, AAA<sup>+</sup>24], simulating a

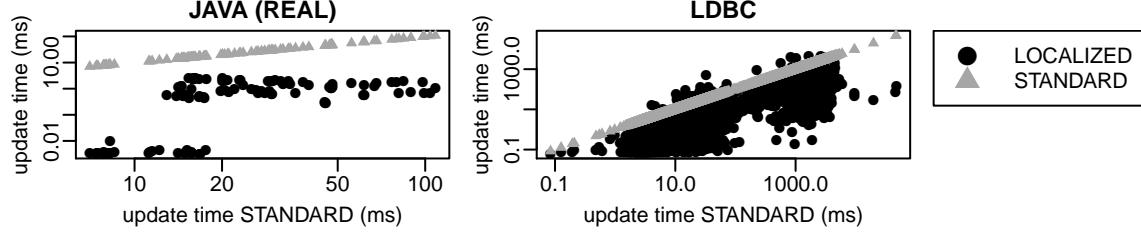


FIGURE 18. Summary of update processing times for plain graph queries in different scenarios (log scale)

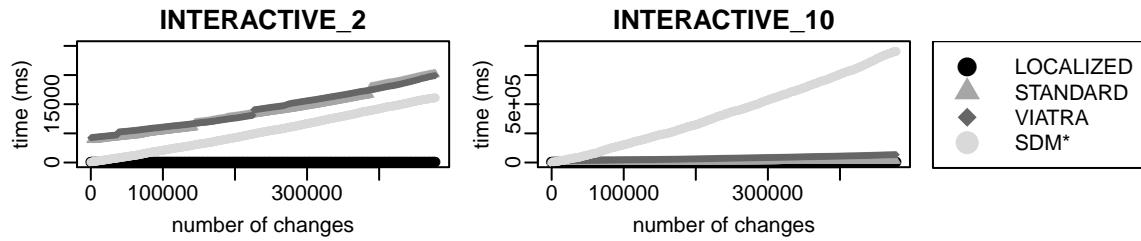


FIGURE 19. Execution times for the LDBC scenario

case where a user of a social network wants to incrementally track query results relating to them personally.<sup>9</sup>

We therefore generate a synthetic social network consisting of around 850 000 vertices, including about 1700 persons, and 5 500 000 edges using the benchmark’s data generator and the predefined scale factor 0.1. We subsequently transform this dataset into a sequence of element creations and deletions based on the timestamps included in the data. We then create a starting graph by replaying the first half of the sequence and perform an initial execution of adapted versions of twelve benchmark queries consisting of plain graph patterns, with a person with a close-to-average number of contacts in the final social network designated as relevant subgraph. After the initial query execution, we replay the remaining changes, incrementally updating the query results after each change. To evaluate our approach for localizing execution of extended graph queries via the RETE mechanism, we also perform analogous experiments with versions of two queries that are equipped with nested graph conditions according to their original formulation in the benchmark. Visualizations of the employed queries can be found in Appendix B.

**5.3.1. Plain Graph Queries.** The resulting execution times for the plain queries where LOCALIZED performed best and worst compared to STANDARD are displayed in Figure 19. A summary of all update time measurements for LOCALIZED in comparison with STANDARD is also displayed in Figure 18 (right). For all queries, LOCALIZED ultimately outperforms the other approaches by a substantial margin, as the localized RETE version forgoes the computation of a large number of irrelevant intermediate results due to the small relevant subgraph on the one hand and avoids redundant computations on the other hand.

<sup>9</sup>Our experiments are not to be confused with an official run of the LDBC Social Network Benchmark. The benchmark specification and data generator only serve as a source of plausible data and queries.

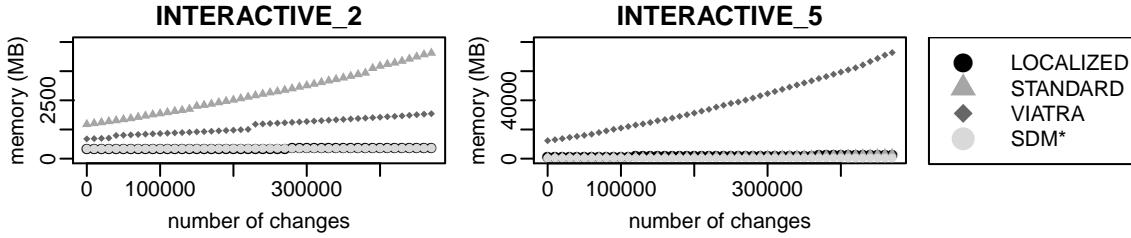


FIGURE 20. Memory measurements for the LDBC scenario

The memory measurements in Figure 20 mostly mirror execution times for RETE-based approaches, with the memory consumption for LOCALIZED always lower than for STANDARD and VIATRA except for a period at the beginning of the execution of the query INTERACTIVE\_5, where STANDARD outperforms LOCALIZED. The weaker performance for INTERACTIVE\_5 and INTERACTIVE\_10 likely stems from the fact that, disregarding edge direction, the associated query graphs contain cycles that act as strong filters for subsequent (intermediate) results. These filters then achieve a somewhat similar effect as localization. The weaker performance of VIATRA for INTERACTIVE\_5 is a product of the usage of a suboptimal RETE net. As expected, memory consumption is lowest for SDM\* for all queries.

**5.3.2. Extended Graph Queries.** The resulting execution times for the considered extended graph queries are displayed in Figure 21. A visual comparison between update times for LOCALIZED and DELTA with STANDARD is shown in Figure 22.

For both queries, LOCALIZED outperforms STANDARD and VIATRA significantly with respect to both initial execution time and update times, as it benefits from the same localized computation of matches for the base pattern as in the plain query case. Moreover, checking effort for the equipped nested graph conditions is also reduced, with checks only performed locally in the context of relevant base pattern matches rather than globally.

Compared to LOCALIZED, DELTA requires noticeably more initial execution time and time for processing updates for both queries. The deterioration in performance is likely a product of several factors. First, the required duplication of the RETE nets for localized match computation and checking of nested graph condition creates obvious redundancy. Second, DELTA introduces additional RETE nets for computing subgraph satisfaction dependent matches for the base pattern. Third, effectively enabling the execution of queries over two versions of the host graph at the same time comes with additional indexing effort required for the DELTA strategy.

The deterioration in performance is slightly more pronounced for the query INTERACTIVE\_4\_NGC. This is likely due to the fact that the typing of vertices in the nested graph conditions of INTERACTIVE\_3\_NGC prevents changes to the relevant subgraph from impacting the conditions' satisfaction, which leads to empty RETE subnets for the computation of subgraph satisfaction dependent matches.

Most observations related to execution time again also apply to memory consumption due to the close link between RETE net execution time and configuration size. The most notable exception is the lower memory consumption measured for VIATRA and the query INTERACTIVE\_4\_NGC compared to DELTA. This is likely a product of a generally more memory efficient RETE implementation in VIATRA in comparison with our own, which is

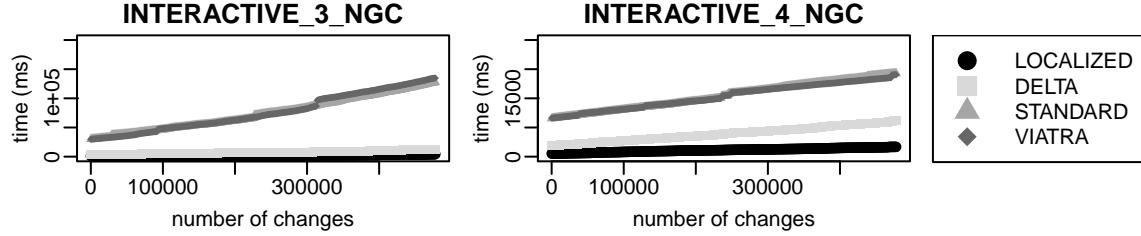


FIGURE 21. Execution times for the LDBC scenario

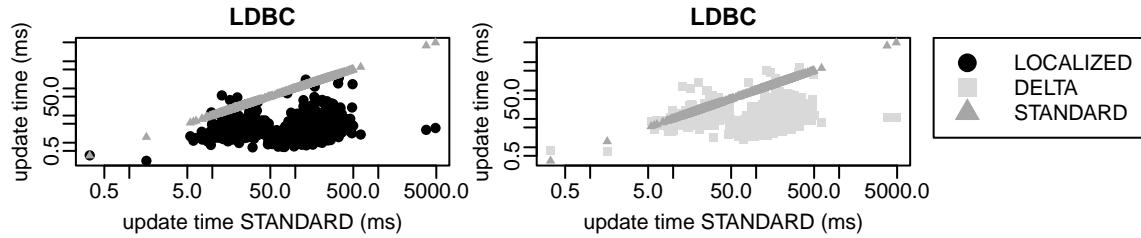


FIGURE 22. Summary of update processing times for extended graph queries in the LDBC scenario (log scale)

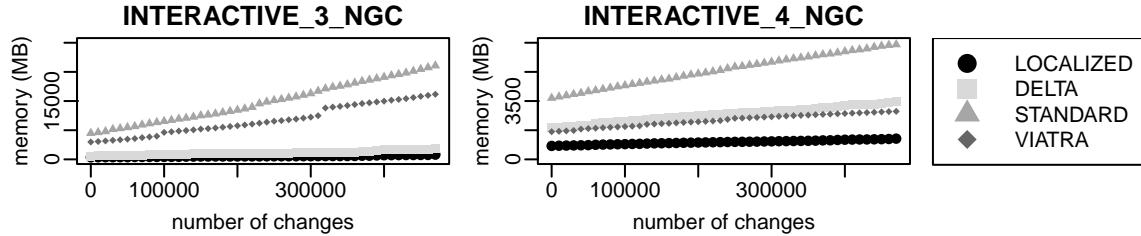


FIGURE 23. Memory measurements for the LDBC scenario

also illustrated by the difference in memory consumption between VIATRA and STANDARD despite very similar execution times.

Note that SDM\* is excluded from our experiments with extended graph queries. This is due to the fact that it is unclear how to implement change-based search of matches for extended graph queries purely with local search and without implementing some kind of auxiliary data structure for storing information about the search process.

**5.4. Discussion.** On the one hand, our experimental results demonstrate that in situations where the relevant subgraph constitutes only a fraction of the full model, RETE net localization can improve the performance of incremental query execution compared to both the standard RETE approach and a solution based on local search. In such scenarios, localization can improve scalability with respect to initial query execution time and memory consumption, as demonstrated in Sections 5.1 and 5.3 and, if changes are not restricted to the relevant subgraph, also update processing time, as shown in Sections 5.2 and 5.3.

On the other hand, as demonstrated in Section 5.1, localization incurs an overhead on update processing time if changes are only made to the relevant subgraph and on initial

execution time and memory consumption if the relevant subgraph contains most of the elements in the full model. While this overhead will essentially be limited to a constant factor in many scenarios, as analyzed in Section 3.3 and 4.1, the standard RETE approach remains preferable for query execution with global semantics or if modifications are restricted to the relevant subgraph and initial query execution time and memory consumption are irrelevant.

These observations apply to both plain and extended graph queries when it comes to the application scenario of tracking all matches for the base query pattern that include elements from the relevant subgraph. In the case of plain graph queries, the application scenario of localized tracking of changes to a global query result coincides with the first. However, for extended graph queries, localization requires noticeably more effort in the second scenario and can lead to increased theoretical complexity if further optimizations are involved, as discussed in Section 4.2. However, our evaluation results in Section 5.3 indicate that it may still prove beneficial in practice.

**5.5. Threats to Validity.** To mitigate internal threats to the validity of our results resulting from unexpected JVM behavior, we have performed 10 runs of all experiments. However, with reliable memory measurements a known pain point of Java-based experiments, the reported memory consumption values are still not necessarily accurate and can only serve as an indicator. To minimize the impact of the implementation on conceptual observations, we compare the prototypical implementation of our approach to a regular RETE implementation [BG23], which shares a large portion of the involved code, and to two existing tools [VBH<sup>+</sup>16, GHS09].

We have attempted to address external threats to validity via experiments accounting for different application domains and a combination of synthetic and real-world queries and data, including a setting from an established, independent benchmark. Still, our results cannot be generalized and do not support quantitative claims, but serve to demonstrate conceptual advantages and disadvantages of the presented approach.

## 6. RELATED WORK

With graph query execution forming the foundation of many applications, there already exists an extensive body of research regarding its optimization.

Techniques based on local search [CFSV04, GBG<sup>+</sup>06, GHS09, ABJ<sup>+</sup>10, HLL13, BCL<sup>+</sup>16, JM18] constitute one family of graph querying approaches. While they are designed to exploit locality in the host graph to improve execution time, repeated query execution leads to redundant computations that are only avoided by fully incremental techniques.

In [Egy06], Egyed proposes a scoping mechanism for local search to support incremental query execution, only recomputing query results when a graph element touched during query execution changes. While this approach offers some degree of incrementality, it is limited to queries with designated root elements that serve as an anchor in the host graph and may still result in redundant computations, since query reevaluation is only controlled at the granularity of root elements.

A second family of solutions is based on discrimination networks [HBC02, VD13, VBH<sup>+</sup>16, Bey18], the most prominent example of which are RETE nets.

VIATRA [VBH<sup>+</sup>16] is a mature tool for incremental graph query execution based on the RETE algorithm [For89], which supports advanced concepts for query specification and

optimization not considered in this article. Notably, VIATRA allows reuse of matches for isomorphic query subgraphs within a single RETE net. This is achieved via RETE structures not covered by the rather restrictive definition of well-formedness used in this article, which points to a possible direction for future work. However, while VIATRA also has a local search option for query execution, it does not integrate local search with the incremental query engine but rather offers it as an alternative.

Beyhl [Bey18] presents an incremental querying technique based on a generalized version of RETE nets, called Generalized Discrimination Networks (GDNs) [HBC02]. The main difference compared to the RETE algorithm is the lack of join nodes. Instead, more complex nodes that directly compute complex matches using local search are employed. The approach however represents more of a means of controlling the trade-off between local search and RETE rather than an integration and still requires a global computation of matches for the entire host graph.

In previous work [BBG21], we have made a first step in the direction of localizing RETE-based query execution. While this earlier technique already allowed anchoring the execution of a RETE net to certain host graph vertices, this anchoring was based on typing information and its results did not meet the definition of completeness introduced in this article.

Model repositories such as CDO [cd0] and NeoEMF [DSB<sup>+</sup>17] provide support for query execution over partial models via lazy loading. As persistence layers, these solutions however focus on implementing an interface of atomic model access operations in order to be agnostic regarding the employed query mechanism.

The Mogwai tool [DSC18] aims to improve query execution over persistence layers like CDO and NeoEMF by mapping model queries to native queries for the underlying database system instead of using the atomic model access operations provided by the layer's API, avoiding loading the entire model into main memory. The tool however does not consider incremental query execution.

Jahanbin et al. propose an approach for querying partially loaded models stored via persistence layers [JKGS22] or as XMI files [JKG23]. In contrast to the solution presented in this article, their approach still aims to always provide complete query results for the full model and is thus based on static analysis and typing information rather than dynamic exploitation of locality.

Query optimization for relational databases is a research topic that has been under intense study for decades [KBZ86, LSC01, LGM<sup>+</sup>15]. Generally, many of the techniques from this field are applicable to RETE nets, which are ultimately based on relational algebra and related to materialized views in relational databases [GM99]. However, relational databases lack the notion of locality inherent to graph-based encodings and are hence not tailored to exploit local navigation.

This shortcoming has given rise to a number of graph databases [Ang12], which employ a graph-based data representation instead of a relational encoding and form the basis of some model persistence layers like NeoEMF [DSB<sup>+</sup>17]. While these database systems are designed to accommodate local navigation, to the best of our knowledge, support for incremental query execution is still lacking.

## 7. CONCLUSION

In this article, we have presented a relaxed notion of completeness for query results that lifts the requirement of strict completeness of results for graph queries and thereby the need for necessarily global query execution. Based on this relaxed notion of completeness, we have developed an extension of the RETE approach that allows local, yet fully incremental execution of both plain and extended graph queries. An initial evaluation demonstrates that the approach can improve scalability in scenarios with small relevant subgraphs, but causes an overhead in unfavorable cases.

In future work, we want to further evaluate the performance of the presented approach. In particular, we are interested in exploring whether the RETE net localization technique for extended graph queries can improve performance in certain scenarios even if host graph and relevant subgraph coincide. We also plan to investigate whether the proposed solution can be adapted to support bulk loading of partial models in order to reduce overhead caused by lazy loading strategies employed by model persistence layers.

## ACKNOWLEDGMENTS

This work was developed in the course of the project modular and incremental Global Model Management (project number 336677879) funded by the DFG.

## REFERENCES

- [AAA<sup>+</sup>24] Renzo Angles, János Benjamin Antal, Alex Averbuch, Altan Birler, Peter Boncz, Márton Búr, Orri Erling, Andrey Gubichev, Vlad Haprian, Moritz Kaufmann, Josep Lluís Larriba Pey, Norbert Martínez, József Marton, Marcus Paradies, Minh-Duc Pham, Arnau Prat-Pérez, David Püroja, Mirko Spasić, Benjamin A. Steer, Dávid Szakállas, Gábor Szárnyas, Jack Waudby, Mingxi Wu, and Yuchen Zhang. The LDBC Social Network Benchmark, 2024. [arXiv:2001.02299](https://arxiv.org/abs/2001.02299), doi:10.48550/arXiv.2001.02299.
- [ABJ<sup>+</sup>10] Thorsten Arendt, Enrico Biermann, Stefan Jurack, Christian Krause, and Gabriele Tautzner. Henshin: advanced concepts and tools for in-place EMF model transformations. In *International Conference on Model Driven Engineering Languages and Systems*, pages 121–135. Springer, 2010. doi:10.1007/978-3-642-16145-2\_9.
- [Ang12] Renzo Angles. A comparison of current graph database models. In *International Conference on Data Engineering Workshops*, pages 171–177. IEEE, 2012. doi:10.1109/ICDEW.2012.31.
- [BBG21] Matthias Barkowsky, Thomas Brand, and Holger Giese. Improving adaptive monitoring with incremental runtime model queries. In *Symposium on Software Engineering for Adaptive and Self-Managing Systems*, pages 71–77. IEEE, 2021. doi:10.1109/SEAMS51251.2021.00019.
- [BCDM14] Hugo Bruneliere, Jordi Cabot, Grégoire Dupé, and Frédéric Madiot. Modisco: A model driven reverse engineering framework. *Information and Software Technology*, 56(8):1012–1032, 2014. doi:10.1016/j.infsof.2014.04.007.
- [BCJM10] Hugo Bruneliere, Jordi Cabot, Frédéric Jouault, and Frédéric Madiot. MoDisco: a generic and extensible framework for model driven reverse engineering. In *IEEE/ACM International Conference on Automated Software Engineering*, pages 173–174, 2010. doi:10.1145/1858996.1859032.
- [BCL<sup>+</sup>16] Fei Bi, Lijun Chang, Xuemin Lin, Lu Qin, and Wenjie Zhang. Efficient subgraph matching by postponing cartesian products. In *International Conference on Management of Data*, pages 1199–1214. ACM, 2016. doi:10.1145/2882903.2915236.
- [Bey18] Thomas Beyhl. *A framework for incremental view graph maintenance*. PhD thesis, Hasso Plattner Institute at the University of Potsdam, 2018.

- [BG23] Matthias Barkowsky and Holger Giese. Host-graph-sensitive RETE nets for incremental graph pattern matching with nested graph conditions. *Journal of Logical and Algebraic Methods in Programming*, 131, 2023. doi:10.1016/j.jlamp.2022.100841.
- [BG24a] Matthias Barkowsky and Holger Giese. Localized RETE for incremental graph queries. In *International Conference on Graph Transformation*, pages 118–137. Springer, 2024. doi:10.1007/978-3-031-64285-2\_7.
- [BG24b] Matthias Barkowsky and Holger Giese. Localized RETE for incremental graph queries. *arXiv preprint*, 2024. doi:10.48550/arXiv.2405.01145.
- [BG25] Matthias Barkowsky and Holger Giese. Localized RETE for Incremental Graph Queries with Nested Graph Conditions Evaluation Artifacts, 2025. Zenodo. doi:10.5281/zenodo.15754399.
- [cd0] Eclipse CDO Model Repository. <https://projects.eclipse.org/projects/modeling.emf.cdo>. Last accessed 13 December 2024.
- [CFSV04] Luigi P Cordella, Pasquale Foggia, Carlo Sansone, and Mario Vento. A (sub) graph isomorphism algorithm for matching large graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1367–1372, 2004. doi:10.1109/TPAMI.2004.75.
- [Cod70] Edgar F Codd. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387, 1970. doi:10.1145/362384.362685.
- [DSB<sup>+</sup>17] Gwendal Daniel, Gerson Sunyé, Amine Benelallam, Massimo Tisi, Yoann Vernageau, Abel Gómez, and Jordi Cabot. NeoEMF: A multi-database model persistence framework for very large models. *Science of Computer Programming*, 149:9–14, 2017. doi:10.1016/j.scico.2017.08.002.
- [DSC18] Gwendal Daniel, Gerson Sunyé, and Jordi Cabot. Scalable queries and model transformations with the Mogwai tool. In *Theory and Practice of Model Transformation*, pages 175–183. Springer, 2018. doi:10.1007/978-3-319-93317-7\_9.
- [EALP<sup>+</sup>15] Orri Erling, Alex Averbuch, Josep Larriba-Pey, Hassan Chafi, Andrey Gubichev, Arnau Prat, Minh-Duc Pham, and Peter Boncz. The LDBC social network benchmark: Interactive workload. In *ACM SIGMOD International Conference on Management of Data*, pages 619–630. ACM, 2015. doi:10.1145/2723372.2742786.
- [EEPT06] Hartmut Ehrig, Karsten Ehrig, Ulrike Prange, and Gabriele Taentzer. *Fundamentals of algebraic graph transformation*. Springer, 2006. doi:10.1007/3-540-31188-2.
- [Egy06] Alexander Egyed. Instant consistency checking for the UML. In *International Conference on Software Engineering*, pages 381–390, 2006. doi:10.1145/1134285.1134339.
- [emf] EMF: Eclipse Modeling Framework. <https://www.eclipse.org/modeling/emf/>. Last accessed 13 December 2024.
- [For89] Charles L Forgy. Rete: A fast algorithm for the many pattern/many object pattern match problem. In *Readings in Artificial Intelligence and Databases*, pages 547–559. Elsevier, 1989. doi:10.1016/0004-3702(82)90020-0.
- [GBG<sup>+</sup>06] Rubino Geiß, Gernot Veit Batz, Daniel Grund, Sebastian Hack, and Adam Szalkowski. GrGen: A fast SPO-based graph rewriting tool. In *International Conference on Graph Transformation*, pages 383–397. Springer, 2006. doi:10.1007/11841883\_27.
- [GHJV93] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. Design patterns: Abstraction and reuse of object-oriented design. In *European Conference on Object-Oriented Programming*, pages 406–431. Springer, 1993. doi:10.1007/3-540-47910-4\_21.
- [GHS09] Holger Giese, Stephan Hildebrandt, and Andreas Seibel. Improved flexibility and scalability by interpreting story diagrams. *Electronic Communications of the EASST*, 18, 2009. doi:10.14279/tuj.eceasst.18.268.
- [GM99] Ashish Gupta and Inderpal Singh Mumick. Maintenance of materialized views: problems, techniques, and applications. In *Materialized Views: Techniques, Implementations, and Applications*. The MIT Press, 05 1999. doi:10.7551/mitpress/4472.003.0016.
- [HBC02] Eric N Hanson, Sreenath Bodagala, and Ullas Chadaga. Trigger condition testing and view maintenance using optimized discrimination networks. *IEEE Transactions on Knowledge and Data Engineering*, 14(2):261–280, 2002. doi:10.1109/69.991716.
- [HLL13] Wook-Shin Han, Jinsoo Lee, and Jeong-Hoon Lee. Turboiso: towards ultrafast and robust subgraph isomorphism search in large graph databases. In *ACM SIGMOD International Conference on Management of Data*, pages 337–348. ACM, 2013. doi:10.1145/2463676.2465300.

- [HP09] Annegret Habel and Karl-Heinz Pennemann. Correctness of high-level transformation systems relative to nested conditions. *Mathematical Structures in Computer Science*, 19(2):245–296, 2009. [doi:10.1017/S0960129508007202](https://doi.org/10.1017/S0960129508007202).
- [JKG23] Sorour Jahanbin, Dimitris Kolovos, and Simos Gerasimou. Towards memory-efficient validation of large XMI models. In *ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*, pages 241–250. IEEE, 2023. [doi:10.1109/MODELS-C59198.2023.00053](https://doi.org/10.1109/MODELS-C59198.2023.00053).
- [JKGS22] Sorour Jahanbin, Dimitris Kolovos, Simos Gerasimou, and Gerson Sunyé. Partial loading of repository-based models through static analysis. In *ACM SIGPLAN International Conference on Software Language Engineering*, pages 266–278, 2022. [doi:10.1145/3567512.3567535](https://doi.org/10.1145/3567512.3567535).
- [JM18] Alpár Jüttner and Péter Madarasi. VF2++—An improved subgraph isomorphism algorithm. *Discrete Applied Mathematics*, 242:69–81, 2018. [doi:10.1016/j.dam.2018.02.018](https://doi.org/10.1016/j.dam.2018.02.018).
- [KBZ86] Ravi Krishnamurthy, Haran Boral, and Carlo Zaniolo. Optimization of nonrecursive queries. In *International Conference on Very Large Data Bases*, volume 86, pages 128–137, 1986.
- [Ken02] Stuart Kent. Model driven engineering. In *International Conference on Integrated Formal Methods*, pages 286–298. Springer, 2002. [doi:10.1007/3-540-47884-1\\_16](https://doi.org/10.1007/3-540-47884-1_16).
- [LGM<sup>+</sup>15] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. How good are query optimizers, really? *Proceedings of the VLDB Endowment*, 9(3):204–215, 2015. [doi:10.14778/2850583.2850594](https://doi.org/10.14778/2850583.2850594).
- [LSC01] Chiang Lee, Chi-Sheng Shih, and Yaw-Huei Chen. Optimizing large join queries using a graph-based approach. *IEEE Transactions on Knowledge and Data Engineering*, 13(2):298–315, 2001. [doi:10.1109/69.917567](https://doi.org/10.1109/69.917567).
- [Ren04] Arend Rensink. Representing first-order logic using graphs. In *International Conference on Graph Transformation*, pages 319–335. Springer, 2004. [doi:10.1007/978-3-540-30203-2\\_23](https://doi.org/10.1007/978-3-540-30203-2_23).
- [Rib99] Leila Ribeiro. Parallel composition of graph grammars. *Applied categorical structures*, 7(4):405–430, 1999. [doi:10.1023/A:1008691205954](https://doi.org/10.1023/A:1008691205954).
- [SIRV18] Gábor Szárnyas, Benedek Izsó, István Ráth, and Dániel Varró. The Train Benchmark: cross-technology performance evaluation of continuous model queries. *Software & Systems Modeling*, 17:1365–1393, 2018. [doi:10.1007/s10270-016-0571-8](https://doi.org/10.1007/s10270-016-0571-8).
- [TELW14] Gabriele Taentzer, Claudia Ermel, Philip Langer, and Manuel Wimmer. A fundamental approach to model versioning based on graph modifications: from theory to implementation. *Software & Systems Modeling*, 13(1):239–272, 2014. [doi:10.1007/s10270-012-0248-x](https://doi.org/10.1007/s10270-012-0248-x).
- [VBH<sup>+</sup>16] Dániel Varró, Gábor Bergmann, Ábel Hegedüs, Ákos Horváth, István Ráth, and Zoltán Ujhelyi. Road to a reactive and incremental model transformation platform: three generations of the VIATRA framework. *Software & Systems Modeling*, 15(3):609–629, 2016. [doi:10.1007/s10270-016-0530-4](https://doi.org/10.1007/s10270-016-0530-4).
- [VD13] Gergely Varró and Frederik Deckwerth. A Rete network construction algorithm for incremental pattern matching. In *International Conference on Theory and Practice of Model Transformations*, pages 125–140. Springer, 2013. [doi:10.1007/978-3-642-38883-5\\_13](https://doi.org/10.1007/978-3-642-38883-5_13).

## APPENDIX A. TECHNICAL DETAILS

We first introduce a few technical definitions required only for proofs in this appendix.

Given a local navigation structure  $LNS(p)$  containing the marking-sensitive vertex input nodes  $[v]^\Phi$  and  $[w]^\Phi$ , we call  $\chi(LNS(p)) = \{[\cup]_v^\Phi, [\cup]_w^\Phi\}$  the *extension points* of  $LNS(p)$ .

Let  $(N, p)$  be a well-formed RETE net with  $p$  a join node. For the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$  with  $N^\Phi = N_{\bowtie}^\Phi \cup N_l^\Phi \cup N_r^\Phi \cup RPS_l \cup RPS_r$ , the extension points of  $N^\Phi$  are given by  $\chi(N^\Phi) = \chi(N_l^\Phi) \cup \chi(N_r^\Phi)$ .

Any RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$  directly contains a RETE subnet  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$  regardless of the form of the nested graph condition  $\psi$ . The extension points of  $(N^\Phi, p^\Phi)$  are then given by  $\chi(N^\Phi) = \chi(N_Q^\Phi)$ .

We say that a marking-sensitive RETE net  $(X^\Phi, p_X^\Phi)$  is a *modular extension* of  $(N^\Phi, p^\Phi)$  if  $N^\Phi \subseteq X^\Phi$  and  $\forall e \in E^{X^\Phi} : s^{X^\Phi}(e) \in V^{N^\Phi} \wedge t^{X^\Phi}(e) \in V^{N^\Phi} \Rightarrow e \in E^{N^\Phi}$ .

### A.1. Theorems in Section 3.2.

**Theorem 3.6** (Matches are marked  $\infty$  in query results of RETE nets localized via *localize* iff they touch the relevant subgraph). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net, and  $Q = \text{query}(p)$ . Furthermore, let  $\mathcal{C}^\Phi$  be a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$ . It then holds that  $\forall (m, \phi) \in \mathcal{C}^\Phi(p^\Phi) : m(Q) \cap H_p \neq \emptyset \Leftrightarrow \phi = \infty$ .*

*Proof.* Since each extension point  $x \in \chi(N^\Phi)$  is only preceded by marking assignment nodes that assign a marking other than  $\infty$  and some marking-sensitive vertex input node  $[v]^\Phi$  that only extracts matches  $m$  with  $m(v) \in H_p$ , it must hold for the union node  $[\cup]^\Phi$  at the top of each local navigation structure that  $\forall (m, \phi) \in \mathcal{C}^\Phi([\cup]^\Phi) : m(Q) \cap H_p \neq \emptyset \Leftrightarrow \phi = \infty$ .

All joins in  $N^\Phi$  ultimately only combine matches from the top union nodes of local navigation structures, assigning the maximum marking in the process. Consequently, for any marking-sensitive join node  $[\bowtie]^\Phi \in N^\Phi$  and a tuple  $(m, \infty) \in \mathcal{C}^\Phi([\bowtie]^\Phi)$ , there must be a tuple  $(m', \infty) \in \mathcal{C}^\Phi([\cup]^\Phi)$  for some local navigation structure with top union node  $[\cup]^\Phi$  with  $Q' = \text{query}([\cup]^\Phi)$  such that  $m(Q') = m'(Q')$ . Thus, it must hold that  $m(Q) \cap H_p \neq \emptyset$ .

Conversely, by the semantics of the marking-sensitive join node, which always assigns the maximum marking of related dependent matches to the resulting matches, a marking of  $\infty$  ultimately always propagates to the top of the RETE net. By Theorem 3.5 and because the construction of a match that touches the relevant subgraph via joins must always involve at least one join with a match that is marked  $\infty$  in a local navigation structure, it thus follows that  $\forall m \in \mathcal{M}_H^Q : (m, \infty) \in \mathcal{C}^\Phi([\bowtie]^\Phi)$ .

Since  $p^\Phi$  is either a marking-sensitive join node or the top union node of a local navigation structure, it follows that  $\forall (m, \infty) \in \mathcal{C}^\Phi(p^\Phi) : \exists v \in V^Q : m(v) \in V^{H_p}$ .  $\square$

### A.2. Theorems in Section 3.3.

**Theorem 3.10** (RETE net localization via *localize* introduces only a constant factor overhead on effective configuration size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net with  $Q = \text{query}(p)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net*

$(N^\Phi, p^\Phi) = \text{localize}(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . It then holds that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .

*Proof.* Follows directly from Lemma A.4.  $\square$

### A.3. Theorems in Section 4.1.

**Theorem 4.3** (Consistent configurations for RETE nets localized via  $\text{localize}^\Psi$  yield correct query results under the relevant subgraph). *Let  $H$  be a graph,  $H_p \subseteq H$ , and  $(Q, \psi)$  an extended graph query. Furthermore, let  $\mathcal{C}^\Phi$  be a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$ . The set of matches given by the stripped result set  $\mathcal{R}_X^\Phi(p^\Phi, \mathcal{C}^\Phi)$  is then correct under  $H_p$ .*

*Proof.* Regardless of what  $\psi$  looks like, there is a matching marking-sensitive vertex input node connected to each  $x \in \chi(N^\Phi)$ . Because there is at least one extension point  $x \in \chi(N^\Phi)$  with  $\text{query}(x) = (\{v\}, \emptyset, \emptyset, \emptyset)$  for each  $v \in V^Q$ , it follows from Lemma A.6 that  $\forall m \in \mathcal{R}_X^\Phi(p^\Phi, \mathcal{C}^\Phi) : m \models \psi \wedge \forall m \in \mathcal{M}_H^Q : (m' \models \psi \wedge m'(Q) \cap H_p \neq \emptyset) \Rightarrow m' \in \mathcal{R}_X^\Phi(p^\Phi, \mathcal{C}^\Phi)$ . Thus, the theorem holds.  $\square$

**Theorem 4.4** (Execution of localized RETE nets via  $\text{order}^\Psi$  yields consistent configurations). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  a graph query, and  $\mathcal{C}_0^\Phi$  an arbitrary starting configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$ . Executing  $(N^\Phi, p^\Phi)$  via  $O = \text{order}^\Psi(N^\Phi)$  then yields a consistent configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$ .*

*Proof.* Follows directly from Lemma A.7.  $\square$

**Theorem 4.6** (RETE net localization via  $\text{localize}^\Psi$  introduces only a constant factor overhead on effective configuration size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$  corresponding to  $(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . It then holds that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .*

*Proof.* We show the correctness of the theorem via structural induction over the nested graph condition  $\psi$ .

In the base case where  $\psi = \text{true}$ , it holds that  $\text{localize}^\Psi(Q, \psi) = \text{localize}(N, p)$ . From Theorem 3.10, it then immediately follows that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .

We now proceed to showing that, under the induction hypothesis that the lemma holds for a nested graph condition of nesting depth  $d$ , the lemma holds for any nested graph condition of nesting depth  $d + 1$ .

For a nested condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ ,  $N^\Phi$  consists of the localized RETE net for  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the RETE net  $(N_{(Q', \psi')}^\Phi, p_{(Q', \psi')}^\Phi) = \text{localize}^\Psi(Q', \psi')$ , a request projection structure  $RPS_l^\infty = RPS^\infty(p_Q^\Phi, N_{(Q', \psi')}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\bowtie]^\Phi$ .

By the construction described in [BG23],  $(N, p)$  then consists of a RETE net for the plain graph query  $Q$ ,  $(N_Q, p_Q)$ , the RETE net for the extended graph query  $(Q', \psi')$ ,  $(N', p')$ , and a semi-join  $p = [\bowtie]$  with dependencies  $p_Q$  and  $p'$ .

Lemma A.4 entails that  $\sum_{n^\Phi \in V^{N_Q^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \sum_{n \in V^{N_Q} \setminus \{p_Q\}} \sum_{m \in \mathcal{C}(n)} |m| + 5 \cdot \sum_{m \in \mathcal{C}_{\{p_Q\}}} |m|$ , since  $\text{localize}(Q) = \text{localize}(N_Q, p_Q)$ .

In the worst case, the extension points of  $(N^\Phi, p^\Phi)$  already contain all possible single-vertex matches into  $H$  even without the additional input via  $RPS_l^\infty$ . Therefore, by the induction hypothesis, it must hold that  $\sum_{n^\Phi \in V^{N^\Phi(Q', \psi')}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_{V^{N'}}|_e$ .

From Lemma 12 in [BG24b], it follows that  $\sum_{n_Q^\Phi \in V^{RPS_l^\infty}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n_Q^\Phi)} |m| \leq 2 \cdot |\mathcal{C}|_{\{p_Q\}}|_e$ . Finally, from Theorem 4.3 and Lemma 12 in [BG24b], it follows that  $\sum_{(m, \phi) \in \mathcal{C}^\Phi(p_{(Q, \psi)}^\Phi)} |m| \leq |\mathcal{C}|_{\{[\times]\}}|_e$ .

Consequently, it must hold that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_{V^{N_Q} \setminus \{p_Q\}}|_e + 7 \cdot |\mathcal{C}|_{\{p_Q\}}|_e + 7 \cdot |\mathcal{C}|_{V^{N'}}|_e + |\mathcal{C}|_{\{[\times]\}}|_e$  and thus  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .

For a nested condition of the form  $\psi = \neg\psi'$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net  $(N_{(Q, \psi')}^\Phi, p_{(Q, \psi')}^\Phi) = \text{localize}^\Psi(Q, \psi')$ , a request projection structure  $RPS_l^\infty = RPS^\infty(p_Q^\Phi, N_{(Q, \psi')}^\Phi)$ , and marking-sensitive anti-join node  $p^\Phi = [\triangleright]^\Phi$ .

By the construction described in [BG23],  $(N, p)$  then consists of a RETE net for the plain graph query  $Q$ ,  $(N_Q, p_Q)$ , the RETE net for the extended graph query  $(Q, \psi')$ ,  $(N', p')$ , and an anti-join  $p = [\triangleright]$  with dependencies  $p_Q$  and  $p'$ .

By analogous argumentation as for the existential case, it must then also hold that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .

For a nested condition of the form  $\psi = \psi_1 \wedge \psi_2$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net  $(N_{(Q, \psi_1)}^\Phi, p_{(Q, \psi_1)}^\Phi) = \text{localize}^\Psi(Q, \psi_1)$ , a request projection structure  $RPS_1^\infty = RPS^\infty(p_Q^\Phi, N_{(Q, \psi_1)}^\Phi)$ , a marking-sensitive semi-join node  $[\times]_1^\Phi$ , the localized RETE net  $(N_{(Q, \psi_2)}^\Phi, p_{(Q, \psi_2)}^\Phi) = \text{localize}^\Psi(Q, \psi_2)$ , a request projection structure  $RPS_2^\infty = RPS^\infty([\times]_1^\Phi, N_{(Q, \psi_2)}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\times]_2^\Phi$ .

By the construction described in [BG23],  $(N, p)$  then consists of a RETE net for the plain graph query  $Q$ ,  $(N_Q, p_Q)$ , the RETE net for the extended graph query  $(Q, \psi_1)$ ,  $(N_1, p_1)$ , the RETE net for the extended graph query  $(Q, \psi_2)$ ,  $(N_2, p_2)$ , a semi-join  $[\times]_1$  with dependencies  $p_Q$  and  $p_1$ , and a semi-join  $p = [\times]_2$  with dependencies  $[\times]_1$  and  $p_2$ .

As in the existential case, it then also holds that  $\sum_{n^\Phi \in V^{N_Q^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \sum_{n \in V^{N_Q} \setminus \{p_Q\}} \sum_{m \in \mathcal{C}(n)} |m| + 5 \cdot \sum_{m \in \mathcal{C}(p_Q)} |m|$ . By the induction hypothesis, it holds that  $\sum_{n^\Phi \in V^{N_{(Q, \psi_1)}^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_{V^{N_1}}|_e$  and  $\sum_{n^\Phi \in V^{N_{(Q, \psi_2)}^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_{V^{N_2}}|_e$ .

From Lemma 12 in [BG24b], it follows that  $\sum_{n_Q^\Phi \in V^{RPS_1^\infty}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n_Q^\Phi)} |m| \leq 2 \cdot |\mathcal{C}|_{\{p_Q\}}|_e$ .

We furthermore know by Theorem 4.3 that  $\mathcal{C}^\Phi([\times]_1^\Phi)$  can only contain matches for  $Q$  that satisfy  $\phi_1$  and  $\mathcal{C}([\times]_1)$  contains all such matches. From Lemma 12 in [BG24b], it thus follows that  $\sum_{(m, \phi) \in \mathcal{C}^\Phi([\times]_1^\Phi)} |m| \leq |\mathcal{C}|_{\{[\times]_1\}}|_e$ . Consequently, it must hold that  $\sum_{n_Q^\Phi \in V^{RPS_2^\infty}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n_Q^\Phi)} |m| \leq 2 \cdot |\mathcal{C}|_{\{[\times]_1\}}|_e$ . By Theorem 4.3 and Lemma 12 in [BG24b], it also follows that  $\sum_{(m, \phi) \in \mathcal{C}^\Phi([\times]_2^\Phi)} |m| \leq |\mathcal{C}|_{\{[\times]_2\}}|_e$ .

We thus know that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_{V^{N_Q} \setminus \{p_Q\}}|_e + 7 \cdot |\mathcal{C}|_{\{p_Q\}}|_e + 7 \cdot |\mathcal{C}|_{V^{N_1}}|_e + 7 \cdot |\mathcal{C}|_{V^{N_2}}|_e + 3 \cdot |\mathcal{C}|_{\{[\times]_1\}}|_e + |\mathcal{C}|_{\{[\times]_2\}}|_e$  and thus  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_e$ .

From the correctness of base case and induction step then follows the correctness of the theorem.  $\square$

**Theorem 4.8** (Execution time overhead introduced by RETE net localization via  $localize^\Psi$  depends on query graph size compared to average match size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}_0^\Phi$  the empty configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \psi)$  corresponding to  $(N, p)$ . Executing  $(N^\Phi, p^\Phi)$  via  $execute(order^\Psi(N^\Phi), N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  then takes  $O(T \cdot |Q_{max}|)$  steps, with  $T = \sum_{n \in V^N} |\mathcal{C}(n)|$  and  $|Q_{max}|$  the size of the largest graph among  $Q$  and graphs in  $\psi$ .*

*Proof.* By the construction of  $order^\Psi(N^\Phi)$ , each node in  $N^\Phi$  is either executed only once and none of its dependencies are executed after it or is part of some subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi) = localize(N^{Q'}, p^{Q'})$  for some subnet  $(N^{Q'}, p^{Q'})$  of  $(N, p)$  and only executed as part of a single execution of  $order(N_{Q'}^\Phi)$ .

By Theorem 3.12, the execution of  $order(N_{Q'}^\Phi)$  for  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  in isolation is in  $O(T_{Q'} \cdot |Q'|)$ , with  $T_{Q'} = \sum_{n \in V^{N^{Q'}}} |\mathcal{C}(n)|$ . Each subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  has at most one union node in one of its local navigation structures that has the marking assignment node of a request projection structure inserted by  $localize^\Psi$  as an additional dependency. By Lemma 17 in [BG24b], the time for executing all nodes in  $N^\Phi$  that are part of such a subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  is in  $O((T_Q^\Phi + T_{RPS}^\Phi) \cdot Q_{max})$ , with  $T_Q^\Phi = \sum_{n \in N_Q^\Phi} |\mathcal{C}(n)|$  and  $T_{RPS}^\Phi = \sum_{n \in N_{RPS}^\Phi} |\mathcal{C}(n)|$ , where  $N_Q^\Phi$  denotes the set of all RETE nodes in such subnets  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  and  $N_{RPS}^\Phi$  denotes the set of all RETE nodes in request projection structures inserted by  $localize^\Psi$ .

All remaining nodes in  $(N^\Phi, p^\Phi)$  are then marking-sensitive semi-joins, anti-joins or projections, marking assignments, or marking filters that are only executed once. Moreover, each node in  $(N^\Phi, p^\Phi)$  can be the dependency of at most one of these semi-joins or anti-joins and at most one of these marking filters. By Lemmata 16, 18, and 19 in [BG24b] and Lemmata A.8 and A.9, the execution of these nodes is thus in  $O((T_\Psi^\Phi + T_{RPS}^\Phi + T_Q^\Phi) \cdot Q_{max})$ , with  $T_\Psi^\Phi = \sum_{n \in N_\Psi^\Phi} |\mathcal{C}(n)|$ , where  $N_\Psi^\Phi$  denotes the set of all semi-join and anti-join nodes inserted by  $localize^\Psi$ .

Since  $T_\Psi^\Phi + T_{RPS}^\Phi + T_Q^\Phi \leq 7 \cdot T$  by Lemma A.10, it then follows that executing  $(N^\Phi, p^\Phi)$  via  $execute(order^\Psi(N^\Phi), N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  takes  $O(T \cdot |Q_{max}|)$  steps.  $\square$

#### A.4. Theorems in Section 4.2.

**Theorem 4.11** (Subgraph-restricted graph modifications can only change NGC satisfaction for subgraph satisfaction dependent matches). *Let  $(Q, \psi)$  be an extended graph query and  $H \xleftarrow{f} K \xrightarrow{g} H'$  a subgraph-restricted graph modification from host graph  $H$  with relevant subgraph  $H_p$  into the modified host graph  $H'$  with relevant subgraph  $H'_p$ . It then holds for any graph morphisms  $m_K : Q \rightarrow K$  that  $(f \circ m_K \models \psi \wedge g \circ m_K \not\models \psi) \vee (f \circ m_K \not\models \psi \wedge g \circ m_K \models \psi) \Rightarrow f \circ m_K$  is subgraph satisfaction dependent or  $g \circ m_K$  is subgraph satisfaction dependent.*

*Proof.* We show the correctness of the theorem via structural induction over the nested graph condition  $\psi$ .

In the base case of  $\psi = \text{true}$ , the theorem is trivially satisfied, since any match always satisfies  $\text{true}$ .

We now proceed to showing that, under the induction hypothesis that the lemma holds for a nested graph condition of nesting depth  $d$ , the lemma holds for any nested graph

condition of nesting depth  $d + 1$ . Let therefore  $m_K : Q \rightarrow K$  be a graph morphism such that  $(f \circ m_K \models \psi \wedge g \circ m_K \not\models \psi) \vee (f \circ m_K \not\models \psi \wedge g \circ m_K \models \psi)$ .

For a nested condition of the form  $\psi = \neg\psi'$ , it follows that  $(f \circ m_K \not\models \psi' \wedge g \circ m_K \models \psi') \vee (f \circ m_K \models \psi' \wedge g \circ m_K \not\models \psi')$ . From the induction hypothesis and the definition of subgraph satisfaction dependence then follows the correctness of the theorem.

For a nested condition of the form  $\psi = \psi_1 \wedge \psi_2$ , it follows that  $(f \circ m_K \models \psi_1 \wedge g \circ m_K \not\models \psi_1) \vee (f \circ m_K \not\models \psi_1 \wedge g \circ m_K \models \psi_1)$  or  $(f \circ m_K \models \psi_2 \wedge g \circ m_K \not\models \psi_2) \vee (f \circ m_K \not\models \psi_2 \wedge g \circ m_K \models \psi_2)$ . From the induction hypothesis and the definition of subgraph satisfaction dependence then follows the correctness of the theorem.

For a nested condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ , it must hold that (1)  $\exists m_a \in \mathcal{M}_H^{Q'} : f \circ m_K = m_a \circ a \wedge m_a \models \psi' \wedge \nexists m'_a \in \mathcal{M}_{H'}^{Q'} : g \circ m_K = m'_a \circ a \wedge m'_a \models \psi'$  or (2)  $\nexists m_a \in \mathcal{M}_H^{Q'} : f \circ m_K = m_a \circ a \wedge m_a \models \psi' \wedge \exists m'_a \in \mathcal{M}_{H'}^{Q'} : g \circ m_K = m'_a \circ a \wedge m'_a \models \psi'$ .

Case 1: Let  $m_a : Q' \rightarrow H$  such that  $f \circ m_K = m_a \circ a \wedge m_a \models \psi'$ . It must then either hold that  $m'_a = \text{iso}_s \circ m_a$  is (1.1) not a valid graph morphism from  $Q'$  into  $H'$  or (1.2)  $m'_a \not\models \psi'$ . In case 1.1, since  $\text{iso}_s$  is an isomorphism, it then follows that  $m_a(Q') \cap H_p \neq \emptyset$  and thus  $f \circ m_K$  is subgraph satisfaction dependent. In case 1.2, it follows that there must be some graph morphism  $m_{K_a}$  such that  $m_a = f \circ m_{K_a}$  and  $m'_a = g \circ m_{K_a}$ . Consequently,  $m_a$  must be subgraph satisfaction dependent for  $\psi'$  or  $m'_a$  must be subgraph satisfaction dependent for  $\psi'$  by the induction hypothesis. Since in this case,  $g \circ m_K = m'_a \circ a$ , it follows that  $f \circ m_K$  is subgraph satisfaction dependent or  $g \circ m_K$  is subgraph satisfaction dependent.

Case 2: Argumentation works analogously to Case 1.

In the case where  $a$  is a partial graph morphism from some subgraph  $Q_p \subseteq Q$  into  $Q'$ , the argumentation works analogously.

From the correctness of the base case and the induction step follows the correctness of the theorem.  $\square$

**Theorem 4.12** (Consistent configurations for RETE nets localized via  $\text{localize}^{\text{sat}}$  yield all subgraph satisfaction dependent matches). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  an extended graph query, and  $\mathcal{C}^\Phi$  a consistent configuration for the RETE net  $(N^{\text{sat}}, p^{\text{sat}}) = \text{localize}^{\text{sat}}(Q, \psi)$ . It then holds that  $\forall m \in \mathcal{M}_H^Q : m$  is subgraph satisfaction dependent  $\Rightarrow m \in \mathcal{R}_X^\Phi(p^{\text{sat}}, \mathcal{C}^\Phi)$ .*

*Proof.* Follows directly from Lemma A.11.  $\square$

**Theorem 4.13** (Execution of localized RETE nets via  $\text{order}^{\text{sat}}$  yields consistent configurations). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  a graph query, and  $\mathcal{C}_0^\Phi$  an arbitrary starting configuration for the marking-sensitive RETE net  $(N^{\text{sat}}, p^{\text{sat}}) = \text{localize}^{\text{sat}}(Q, \psi)$ . Executing  $(N^{\text{sat}}, p^{\text{sat}})$  via  $O = \text{order}^{\text{sat}}(N^{\text{sat}})$  then yields a consistent configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^{\text{sat}}, H, H_p, \mathcal{C}_0^\Phi)$ .*

*Proof.* Follows directly from Lemma A.12.  $\square$

**Theorem 4.15** (For a subgraph-restricted graph modification, consistent configurations for RETE nets localized via  $\text{localize}^\Delta$  yield all matches that touch the relevant subgraph in the source or are subgraph satisfaction dependent in the source or target of the modification). *Let  $(Q, \psi)$  be an extended graph query and  $H \xleftarrow{f} K \xrightarrow{g} H'$  a subgraph-restricted graph modification modifying the graph  $H$  with relevant subgraph  $H_p \subseteq H$*

into the graph  $H'$  with relevant subgraph  $H'_p \subseteq H'$ . Furthermore, let  $\mathcal{C}^\Phi$  be a configuration that is consistent for  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  and consistent for  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'}) = \text{localize}^{\text{sat}}(Q, \psi)$  for host graph  $H'$  and relevant subgraph  $H'_p$ . It must then hold that  $\forall m \in \mathcal{M}_H^Q : m \models \psi \wedge (m(Q) \cap H_p \neq \emptyset \text{ or } m \text{ is subgraph satisfaction dependent or } \exists m' \in \mathcal{M}_{H'}^Q : f \circ g^{-1} \circ m' = m \wedge m' \text{ is subgraph satisfaction dependent}) \Rightarrow m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi)$ .

*Proof.* By the definition of  $\text{localize}^\Delta$ ,  $N^\Delta$  consists of  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(N^Q, p^Q)$  for a regular RETE net  $(N^Q, p^Q)$  with height  $h$  for the plain graph query  $Q$ ,  $[\phi > h]^\Phi$ ,  $(N^{\text{sat}}, p^{\text{sat}}) = \text{localize}^{\text{sat}}(Q, \psi)$ , and  $[\cup]^\Phi$ . In addition,  $N^\Delta$  comprises  $(N_{(Q,\psi)}^\Phi, p_{(Q,\psi)}^\Phi) = \text{localize}^\Psi(Q, \psi)$ ,  $RPS_l^\infty = RPS^\infty([\cup]^\Phi, N_{(Q,\psi)}^\Phi)$ ,  $[\times]^\Phi$ , and  $[\xrightarrow{f \circ g^{-1}}]^\Phi$ .

Let  $m \in \mathcal{M}_H^Q$  with  $m \models \psi \wedge (m(Q) \cap H_p \neq \emptyset \text{ or } m \text{ is subgraph satisfaction dependent or } \exists m' \in \mathcal{M}_{H'}^Q : f \circ g^{-1} \circ m' = m \wedge m' \text{ is subgraph satisfaction dependent})$ . We can then distinguish three cases:

Case 1: If  $m(Q) \cap H_p \neq \emptyset$ , it follows from Lemma 1 in [BG24b] that there must be a tuple  $(m, \infty) \in \mathcal{C}^\Phi(p_Q^\Phi)$  and thus also  $(m, \infty) \in \mathcal{C}^\Phi([\phi > h]^\Phi)$ .

Case 2: If  $m$  is subgraph satisfaction dependent, it follows from Theorem 4.12 that there must be a tuple  $(m, \phi) \in \mathcal{C}^\Phi(p^{\text{sat}})$ .

Case 3: If  $\exists m' \in \mathcal{M}_{H'}^Q : f \circ g^{-1} \circ m' = m \wedge m'$  is subgraph satisfaction dependent, it follows from Theorem 4.12 that there must be a tuple  $(m', \phi) \in \mathcal{C}^\Phi(p^{\text{sat}'})$  such that  $f \circ g^{-1} \circ m' = m$ . From the semantics of the marking-sensitive translation node then follows that there must be a tuple  $(m, \phi) \in \mathcal{C}^\Phi([\xrightarrow{f \circ g^{-1}}]^\Phi)$ .

In all three cases, there must hence be a tuple  $(m, \phi) \in \mathcal{C}^\Phi([\cup]^\Phi)$ . By Lemma A.6 and because  $m \models \psi$ , it must then hold that there is a tuple  $(m, \phi') \in \mathcal{C}^\Phi(p_{(Q,\psi)}^\Phi)$  and, by the semantics of the marking-sensitive semi-join,  $(m, \phi) \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi)$ .  $\square$

**Theorem 4.16** (Execution of localized RETE nets via  $\text{order}^\Delta$  yields consistent configurations). *Let  $(Q, \psi)$  be an extended graph query and  $H$  a host graph with relevant subgraph  $H_p$  and  $\mathcal{C}_0^\Phi$  an arbitrary starting configuration for the marking-sensitive RETE net  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$ . Executing  $(N^\Delta, p^\Delta)$  via  $O = \text{order}^\Delta(N^\Delta)$  then yields a consistent configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Delta, H, H_p, \mathcal{C}_0^\Phi)$ .*

*Proof.* Follows directly from Lemma A.14.  $\square$

**Theorem 4.18** (Removal of matches from the results of an extended graph query caused by a subgraph-restricted graph modification can be detected using RETE nets localized via  $\text{localize}^\Delta$ ). *Let  $(Q, \psi)$  be an extended graph query and  $H \xleftarrow{f} K \xrightarrow{g} H'$  a subgraph-restricted graph modification modifying the graph  $H$  with relevant subgraph  $H_p \subseteq H$  into the graph  $H'$  with relevant subgraph  $H'_p \subseteq H'$ . Furthermore, let  $\mathcal{C}^\Phi$  be a configuration that is consistent for  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$  and  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'}) = \text{localize}^{\text{sat}}(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  and consistent for  $(N^{\Delta'}, p^{\Delta'}) = \text{localize}^\Delta(Q, \psi)$  and  $(N_{(Q,\psi)}^{\text{sat}'}, p_{(Q,\psi)}^{\text{sat}'}) = \text{localize}^{\text{sat}}(Q, \psi)$  for host graph  $H'$  and relevant subgraph  $H'_p$ . It then holds that  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge \nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi\} = \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .*

*Proof.* By the definition of  $localize^\Delta$ ,  $N^\Delta$  consists of  $(N_Q^\Phi, p_Q^\Phi) = localize(N^Q, p^Q)$  for a regular RETE net  $(N^Q, p^Q)$  with height  $h$  for the plain graph query  $Q$ ,  $[\phi > h]^\Phi$ ,  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$ , and  $[\cup]^\Phi$ , which has a dependency on  $p_{(Q,\psi)}^{sat'}$  from  $(N_{(Q,\psi)}^{sat'}, p_{(Q,\psi)}^{sat'})$ . In addition,  $N^\Delta$  comprises  $(N_{(Q,\psi)}^\Phi, p_{(Q,\psi)}^\Phi) = localize^\Psi(Q, \psi)$ ,  $RPS_l^\infty = RPS^\infty([\cup]^\Phi, N_{(Q,\psi)}^\Phi)$ ,  $[\bowtie]$ , and  $[\xrightarrow{f \circ g^{-1}}]^\Phi$ .

Analogously,  $N^{\Delta'}$  consists of  $(N_Q^{\Phi'}, p_Q^{\Phi'}) = localize(N^Q, p^Q)$ ,  $[\phi > h]^{\Phi'}$ ,  $(N^{sat'}, p^{sat'}) = localize^{sat}(Q, \psi)$ , and  $[\cup]^{\Phi'}$  which has a dependency on  $p_{(Q,\psi)}^{sat}$  from  $(N_{(Q,\psi)}^{sat}, p_{(Q,\psi)}^{sat})$ . In addition,  $N^{\Delta'}$  comprises  $(N_{(Q,\psi)}^{\Phi'}, p_{(Q,\psi)}^{\Phi'}) = localize^\Psi(Q, \psi)$ ,  $RPS_l^{\infty'} = RPS^\infty([\cup]^{\Phi'}, N_{(Q,\psi)}^{\Phi'})$ ,  $[\bowtie]$ , and  $[\xrightarrow{g \circ f^{-1}}]^{\Phi'}$ .

We first show that  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge \nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi\} \subseteq \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ . Therefore, let  $m \in \mathcal{M}_H^Q$  such that  $m \models \psi \wedge \nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi$ .

We can distinguish two cases:

Case 1: If  $\nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m'$ , it follows that  $m(Q) \cap H_p \neq \emptyset$ , since otherwise,  $m' = iso_s \circ m$  would violate the assumption that  $\nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m'$ . It must then hold by Theorem 3.5 that there is a tuple  $(m, \infty) \in \mathcal{C}^\Phi(p_Q^\Phi)$  and thus also  $m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi)$  by Lemma A.6. By the assumption that  $\nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m'$ ,  $\mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'})$  then cannot contain a match  $m'$  such that  $m = f \circ g^{-1} \circ m'$  and thereby,  $m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

Case 2: If there exists a match  $m' \in \mathcal{M}_{H'}^Q$  such that  $m = f \circ g^{-1} \circ m'$ , it must hold that  $m' \not\models \psi$  for this match  $m'$ . By Theorem 4.11, we then know that either  $m$  or  $m'$  must be subgraph satisfaction dependent. By Theorem 4.15, it must then hold that  $m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi)$ . By Lemma A.6, it also follows that  $m' \notin \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'})$ . By the injectivity of  $f$  and  $g$ ,  $m'$  must furthermore be unique. Consequently,  $m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

Thus, it follows that  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge \nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi\} \subseteq \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

We now show that  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge \nexists m' \in \mathcal{M}_{H'}^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi\} \supseteq \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ . Therefore, let  $m \in \mathcal{R}_X^\Phi(p^\Delta, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

It then directly follows from Lemma A.6 that  $m \models \psi$ . We will now show that there cannot exist a corresponding match  $m'$  from  $Q$  into  $H'$  with  $m = f \circ g^{-1} \circ m'$  and  $m' \models \psi$ .

From the construction of  $N^\Delta$ , it follows that there must exist a tuple  $(m, \infty) \in \mathcal{C}^\Phi(p_Q^\Phi)$ , a tuple  $(m, \phi) \in \mathcal{C}^\Phi(p^{sat})$ , or a tuple  $(m', \phi) \in \mathcal{C}^{\Phi'}(p^{sat'})$ .

Case 1: If  $(m, \infty) \in \mathcal{C}^\Phi(p_Q^\Phi)$ , it must hold by Theorem 3.6 that  $m(Q) \cap H_p \neq \emptyset$ . Since there must be some element  $e$  in  $m(Q) \cap H_p$  for any match  $m'$  with  $m = f \circ g^{-1} \circ m'$  and because  $(g \circ f^{-1})(H_p) \subseteq H'_p$ , it must hold that  $m'(Q) \cap H'_p \neq \emptyset$ . Consequently,  $(m', \infty) \in \mathcal{C}^{\Phi'}(p_Q^{\Phi'})$  by Theorem 3.5. Since  $m \notin \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$  and thus  $m' \notin \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'})$ , it must follow from Lemma A.6 that  $m' \not\models \psi$ , since otherwise, it would hold that  $\{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

Case 2: If  $(m, \phi) \in \mathcal{C}^\Phi(p^{sat})$ , for any match  $m'$  with  $m = f \circ g^{-1} \circ m'$ , it holds that  $m' = g \circ f^{-1} \circ m'$  due to the injectivity of  $f$  and  $g$ . Therefore, it must hold that there is some tuple  $(m', \phi) \in \mathcal{C}^{\Phi'}([\xrightarrow{g \circ f^{-1}}]^\Phi)$ . Hence, it must hold that  $m' \not\models \psi$ , since otherwise, it would hold that  $\{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

Case 3: If  $(m', \phi) \in \mathcal{C}^{\Phi'}(p^{sat'})$  for some  $m'$  with  $m = f \circ g^{-1} \circ m'$ , it immediately follows that  $m' \not\models \psi$ , since otherwise, it would hold that  $\{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

Thus, it follows that  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge \nexists m' \in \mathcal{M}_H^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi\} \supseteq \mathcal{R}_X^\Phi(p^{\Delta}, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .

With inclusions in both directions, it holds that  $\{m \in \mathcal{M}_H^Q \mid m \models \psi \wedge \nexists m' \in \mathcal{M}_H^Q : m = f \circ g^{-1} \circ m' \wedge m' \models \psi\} = \mathcal{R}_X^\Phi(p^{\Delta}, \mathcal{C}^\Phi) \setminus \{m \in \mathcal{M}_H^Q \mid \exists m' \in \mathcal{R}_X^\Phi(p^{\Delta'}, \mathcal{C}^{\Phi'}) : m = f \circ g^{-1} \circ m'\}$ .  $\square$

**Theorem A.1** (RETE net localization via  $localize^{sat}$  introduces only a constant factor overhead on effective configuration size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  corresponding to  $(N, p)$ . It then holds that  $\sum_{n^{sat} \in V^{N^{sat}}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^{sat})} |m| \leq 18 \cdot |\mathcal{C}|_e$ .*

*Proof.* For nested queries of  $(Q, \psi)$  with the form  $(Q', \text{true})$ ,  $localize^{sat}$  only introduces a dummy node to  $N^{sat}$ . The result set for any dummy node in  $N^{sat}$  is always empty. Hence, the result sets of these nodes do not contribute to configuration size. Moreover,  $localize^{sat}$  does not introduce any additional nodes for nested queries with the form  $(Q', \neg\psi')$ .

For each nested query with the form  $(Q', \psi_1 \wedge \psi_2)$ ,  $N^{sat}$  contains one marking-sensitive union node  $[\cup]^\Phi$ , whereas  $N$  contains a RETE subnet  $(N_{Q'}, p_{Q'})$  computing matches for  $Q'$  as well as two semi-joins. Since the result set for  $[\cup]^\Phi$  can only contain at most one tuple for each match for  $Q'$  and the result set for  $p_{Q'}$  contains all such matches, it follows that  $\sum_{(m, \phi) \in \mathcal{C}^\Phi([\cup]^\Phi)} |m| \leq \sum_{m \in \mathcal{C}(p_{Q'})} |m|$ .

For each nested query with the form  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$ ,  $N^{sat}$  contains the subnets  $(N_{Q'}^\Phi, p_{Q'}^\Phi) = localize(Q')$  and  $(N_{Q''}^\Phi, p_{Q''}^\Phi) = localize(Q'')$ , a marking-sensitive union node  $[\cup]^\Phi$ , a marking-sensitive semi-join  $[\times]^\Phi$ , as well as the request projection structure  $RPS_l^\infty = RPS^\infty([\cup]^\Phi, N_{Q'}^\Phi)$ .  $N$  contains a RETE subnet  $(N_{Q'}, p_{Q'})$  computing matches for  $Q'$  and a semi-join  $[\times]$ .

By Theorem 3.10 and because the additional dependency on  $RPS_l^\infty$  does not affect the worst case of a fully populated net  $(N_Q^\Phi, p_Q^\Phi)$ , it holds that  $\sum_{n_{Q'}^\Phi \in V^{N_{Q'}^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n_{Q'}^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_{N_{Q'}}|_e$ . Also, by the semantics of the marking-sensitive semi-join and because the result set for  $p_{Q'}^\Phi$  contains each match for  $Q'$  at most once by Lemma 12 in [BG24b], it must hold that  $\sum_{(m, \phi) \in \mathcal{C}^\Phi([\times]^\Phi)} |m| \leq \sum_{m \in \mathcal{C}(p_{Q'})} |m| \leq |\mathcal{C}|_{N_{Q'}}|_e$ .

Any nested query with the form  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$  must have some nested query  $(Q'', \exists(a : Q'' \rightarrow Q''', \psi''))$  or  $(Q'', \text{true})$  in order for  $\psi$  by the definition of nested graph conditions. In either case,  $N$  has to contain a corresponding subnet  $(N_{Q''}^\Phi, p_{Q''}^\Phi)$  computing matches for  $Q''$ . It then holds for this subnet  $N_{Q''}$  by Theorem 3.10 that

$\sum_{n_{Q''}^\Phi \in V^{N_{Q''}^\Phi}} \sum_{(m,\phi) \in \mathcal{C}^\Phi(n_{Q''}^\Phi)} |m| \leq 7 \cdot |\mathcal{C}|_{N_{Q''}}|_e$ . It also must hold that  $\sum_{(m,\phi) \in \mathcal{C}^\Phi([\cup]^\Phi)} |m| \leq \sum_{m \in \mathcal{C}(p_{Q''})} |m| \leq |\mathcal{C}|_{N_{Q''}}|_e$  and  $\sum_{n_\Phi \in V^{RPS_i^\infty}} \sum_{(m,\phi) \in \mathcal{C}^\Phi(n_\Phi)} |m| \leq 2 \cdot |\mathcal{C}|_{N_{Q''}}|_e$ .

Since  $\psi$  is a tree of subconditions, it follows that each subnet  $(N_{Q''}, p_{Q''})$  must be counted again like this for at most one parent query  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$ . For any nested query of the form  $(Q', \text{true})$ ,  $\mathcal{C}^\Phi$  stores matches with a combined size of at most  $10 \cdot |\mathcal{C}|_{N_{Q'}}|_e$ . For any nested query of the form,  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$ ,  $\mathcal{C}^\Phi$  stores matches with a combined size of at most  $(8 + 10) \cdot |\mathcal{C}|_{N_{Q'}}|_e = 18 \cdot |\mathcal{C}|_{N_{Q'}}|_e$ .

In the worst case, it must thus hold that  $\sum_{n^{sat} \in V^{N^{sat}}} \sum_{(m,\phi) \in \mathcal{C}^\Phi(n^{sat})} |m| \leq 18 \cdot |\mathcal{C}|_e$ .  $\square$

**Corollary A.2** (RETE net localization via  $localize^{sat}$  introduces only a constant factor overhead on memory consumption). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  corresponding to  $(N, p)$ . Assuming that storing a match  $m$  requires an amount of memory in  $O(|m|)$  and storing an element from  $\bar{\mathbb{N}}$  requires an amount of memory in  $O(1)$ , storing  $\mathcal{C}^\Phi$  requires memory in  $O(|\mathcal{C}|_e)$ .*

*Proof.* Follows directly from Theorem A.1 and the assumptions.  $\square$

**Theorem A.3** (Execution time overhead introduced by RETE localization via  $localize^{sat}$  depends on query graph size compared to average match size). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}_0^\Phi$  the empty configuration for the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  corresponding to  $(N, p)$ . Executing  $(N^\Phi, p^\Phi)$  via  $execute(order^\Psi(N^\Phi), N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  then takes  $O(T \cdot |Q_{max}|)$  steps, with  $T = \sum_{n \in V^N} |\mathcal{C}(n)|$  and  $|Q_{max}|$  the size of the largest graph among  $Q$  and graphs in  $\psi$ .*

*Proof.* By the construction of  $order^{sat}(N^{sat})$ , each node in  $N^{sat}$  is either executed only once and none of its dependencies are executed after it or is part of some subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi) = localize(N^{Q'}, p^{Q'})$  for some subnet  $(N^{Q'}, p^{Q'})$  of  $(N, p)$  and only executed as part of a single execution of  $order(N_{Q'}^\Phi)$ .

By Theorem 3.12, the execution of  $order(N_{Q'}^\Phi)$  for  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  in isolation is in  $O(T_{Q'} \cdot |Q'|)$ , with  $T_{Q'} = \sum_{n \in V^{N^{Q'}}} |\mathcal{C}(n)|$ . Each subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  has at most one union node in one of its local navigation structures that has the marking assignment node of a request projection structure inserted by  $localize^{sat}$  as an additional dependency. By Lemma 17 in [BG24b], the time for executing all nodes in  $N^{sat}$  that are part of such a subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  is in  $O((T_Q^\Phi + T_{RPS}^\Phi) \cdot Q_{max})$ , with  $T_Q^\Phi = \sum_{n \in N_Q^\Phi} |\mathcal{C}(n)|$  and  $T_{RPS}^\Phi = \sum_{n \in N_{RPS}^\Phi} |\mathcal{C}(n)|$ , where  $N_Q^\Phi$  denotes the set of all RETE nodes in such subnets  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  and  $N_{RPS}^\Phi$  denotes the set of all RETE nodes in projection structures inserted by  $localize^\Psi$ .

All remaining nodes in  $(N^{sat}, p^{sat})$  are then marking-sensitive semi-joins, unions or projections, marking assignments, or marking filters that are only executed once. Moreover, each node in  $(N^{sat}, p^{sat})$  can be the dependency of at most one of these semi-joins or unions and at most one of these marking filters. By Lemmata 16, 17, 18, and 19 in [BG24b] and Lemma A.8, the execution of these nodes is thus in  $O((T_{sat}^\Phi + T_{RPS}^\Phi + T_Q^\Phi) \cdot Q_{max})$ ,

with  $T_{sat}^\Phi = \sum_{n \in N_{sat}^\Phi} |\mathcal{C}(n)|$ , where  $N_{sat}^\Phi$  denotes the set of all semi-join and anti-join nodes inserted by  $localize^{sat}$ .

Since  $T_{sat}^\Phi + T_{RPS}^\Phi + T_Q^\Phi \leq 25 \cdot T$  by Lemma A.13, it then follows that executing  $(N^{sat}, p^{sat})$  via  $execute(order^{sat}(N^{sat}), N^{sat}, H, H_p, \mathcal{C}_0^\Phi)$  takes  $O(T \cdot |Q_{max}|)$  steps.  $\square$

### A.5. Supplementary Lemmata.

**Lemma A.4** (RETE net localization via  $localize$  introduces an overhead of at most factor 5 on effective configuration size for the production node and at most factor 7 for all other nodes). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net with  $Q$  the associated query graph of  $p$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = localize(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . It then holds that  $\sum_{n^\Phi \in V^{N^\Phi}} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot \sum_{n \in V^N \setminus \{p\}} \sum_{m \in \mathcal{C}(n)} |m| + 5 \cdot \sum_{m \in \mathcal{C}(p)} |m|$ .*

*Proof.* For each edge input node  $[v \rightarrow w]$  in  $N$  with associated query subgraph  $Q_e$ ,  $N^\Phi$  contains the seven nodes of the corresponding local navigation structure  $LNS([v \rightarrow w])$ :  $[v]^\Phi$ ,  $[w]^\Phi$ ,  $[\cup]_v^\Phi$ ,  $[\cup]_w^\Phi$ ,  $[v \rightarrow_n w]^\Phi$ ,  $[w \leftarrow_n v]^\Phi$ , and  $[\cup]^\Phi$ . By the semantics of  $[v \rightarrow w]$ , it must hold that  $\mathcal{C}([v \rightarrow w]) = \mathcal{M}_H^{Q_e}$ .

By the semantics of the forward and backward local navigation nodes, we know that  $\mathcal{R}_X^\Phi([v \rightarrow_n w]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_e}$  and  $\mathcal{R}_X^\Phi([w \leftarrow_n v]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_e}$  and thus, by the semantics of the marking sensitive union node,  $\mathcal{R}_X^\Phi([\cup]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_e}$ .

Furthermore, by the semantics of the marking sensitive input nodes, it must hold that  $\mathcal{R}_X^\Phi([v]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_v}$  and  $\mathcal{R}_X^\Phi([w]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_w}$ , where  $Q_v$  and  $Q_w$  are the query subgraphs associated with  $v$  respectively  $w$ . By the semantics of the marking sensitive union nodes and the assumption regarding the common query subgraph associated with their dependencies, it must also hold that  $\mathcal{R}_X^\Phi([\cup]_v^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_v}$  and  $\mathcal{R}_X^\Phi([\cup]_w^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_w}$ .

By the semantics of the involved marking-sensitive RETE nodes, for each node in  $LNS([n_1 \rightarrow n_2])$ , each match can only be associated with at most one marking in the node's current result set in  $\mathcal{C}^\Phi$ . Furthermore, by assumption, it holds that  $|\mathcal{M}_H^{Q_v}| \leq |\mathcal{M}_H^{Q_e}|$  and  $|\mathcal{M}_H^{Q_w}| \leq |\mathcal{M}_H^{Q_e}|$ . Lastly, we know that  $|Q_v| = |Q_w| \leq \frac{|Q_e|}{2}$ .

It thus holds that  $\sum_{n^\Phi \in V^{LNS([v \rightarrow w])}} \sum_{(m, \phi) \in \mathcal{C}(n^\Phi)} |m| \leq 5 \cdot \sum_{m \in \mathcal{C}([v \rightarrow w])} |m|$ .

For each join node  $[\bowtie]$  in  $N$  with associated query subgraph  $Q_{\bowtie}$  and dependencies  $n_l$  and  $n_r$ ,  $N^\Phi$  contains the corresponding marking-sensitive join node  $[\bowtie]^\Phi$  with dependencies  $n_l^\Phi$  and  $n_r^\Phi$  corresponding to  $n_l$  respectively  $n_r$  with associated query subgraphs  $Q_l$  respectively  $Q_r$ , as well as the six nodes from the related request projection structures:  $[\phi > h]_l^\Phi$ ,  $[\pi_{Q_v}]_l^\Phi$ ,  $[\phi := h]_l^\Phi$ ,  $[\phi > h]_r^\Phi$ ,  $[\pi_{Q_v}]_r^\Phi$ , and  $[\phi := h]_r^\Phi$ .

By the semantics of join and marking sensitive join, we know that  $\mathcal{C}(n_l) \subseteq \mathcal{R}_X^\Phi(n_l^\Phi, \mathcal{C}^\Phi) \wedge \mathcal{C}(n_r) \subseteq \mathcal{R}_X^\Phi(n_r^\Phi, \mathcal{C}^\Phi) \Rightarrow \mathcal{C}([\bowtie]) \subseteq \mathcal{R}_X^\Phi([\bowtie]^\Phi, \mathcal{C}^\Phi)$ . For every edge input node  $[v \rightarrow w]$  in  $N$ , we know that  $\mathcal{C}([v \rightarrow w]) \subseteq \mathcal{R}_X^\Phi([\cup]^\Phi, \mathcal{C}^\Phi)$ , where  $[\cup]^\Phi$  is the root node of  $LNS([v \rightarrow w])$ . It is easy to see that thereby, it must hold that  $\mathcal{C}([\bowtie]) \subseteq \mathcal{R}_X^\Phi([\bowtie]^\Phi, \mathcal{C}^\Phi)$ . By Lemma 12 in [BG24b], it follows that  $\sum_{(m, \phi) \in \mathcal{C}([\bowtie]^\Phi)} |m| \leq \sum_{m \in \mathcal{C}([\bowtie])} |m|$ .

$[\phi > h]_l^\Phi$  is associated with the query subgraph  $Q_l$  and  $[\phi > h]_r^\Phi$  is associated with query subgraph  $Q_r$ .  $[\pi_{Q_v}]_l^\Phi$ ,  $[\phi := h]_l^\Phi$ ,  $[\pi_{Q_v}]_r^\Phi$ , and  $[\phi := h]_r^\Phi$  are all associated with a query subgraph  $Q_v$  that contains only a single vertex, whereas both  $Q_l$  and  $Q_r$  contain

at least one edge. Consequently,  $|Q_v| \leq \frac{|Q_l|}{2}$  and  $|Q_v| \leq \frac{|Q_r|}{2}$ . Hence, it must hold that  $\sum_{n^\Phi \in RPS_l} \sum_{(m, \phi) \in \mathcal{C}(n^\Phi)} |m| \leq 2 \cdot \sum_{m \in \mathcal{C}(n_l)} |m|$  and  $\sum_{n^\Phi \in RPS_r} \sum_{(m, \phi) \in \mathcal{C}(n^\Phi)} |m| \leq 2 \cdot \sum_{m \in \mathcal{C}(n_r)} |m|$ .

Due to the tree-like structure of  $(N, p)$ , the marking-sensitive union node at the top of each local navigation structure and each marking-sensitive join can be connected to at most one request projection structure. Moreover, the union node at the top of the local navigation structure or the marking-sensitive join corresponding to  $p$  is not connected to any request projection structure. It thus follows that  $\sum_{n^\Phi \in V^N} \sum_{(m, \phi) \in \mathcal{C}^\Phi(n^\Phi)} |m| \leq 7 \cdot \sum_{n \in V^N \setminus \{p\}} \sum_{m \in \mathcal{C}(n)} |m| + 5 \cdot \sum_{m \in \mathcal{C}(p)} |m|$ .  $\square$

**Lemma A.5** (RETE net localization via *localize* introduces an overhead of factor 7 on the number of computed matches for the production node and factor 10 for all other nodes). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a well-formed RETE net with  $Q$  the associated query graph of  $p$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the localized RETE net  $(N^\Phi, p^\Phi) = \text{localize}(N, p)$  for host graph  $H$  and relevant subgraph  $H_p$ . It then holds that  $\sum_{n^\Phi \in V^N} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \sum_{n \in V^N \setminus \{p\}} |\mathcal{C}(n)| + 7 \times |\mathcal{C}(p)|$ .*

*Proof.* For each edge input node  $[v \rightarrow w]$  in  $N$  with associated query subgraph  $Q_e$ ,  $N^\Phi$  contains the seven nodes of the corresponding local navigation structure  $LNS([v \rightarrow w])$ :  $[v]^\Phi$ ,  $[w]^\Phi$ ,  $[\cup]_v^\Phi$ ,  $[\cup]_w^\Phi$ ,  $[v \rightarrow_n w]^\Phi$ ,  $[w \leftarrow_n v]^\Phi$ , and  $[\cup]^\Phi$ . By the semantics of  $[v \rightarrow w]$ , it must hold that  $\mathcal{C}([v \rightarrow w]) = \mathcal{M}_H^{Q_e}$ .

By the semantics of the forward and backward local navigation nodes, we know that  $\mathcal{R}_X^\Phi([v \rightarrow_n w]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_e}$  and  $\mathcal{R}_X^\Phi([w \leftarrow_n v]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_e}$  and thus, by the semantics of the marking sensitive union node,  $\mathcal{R}_X^\Phi([\cup]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_e}$ .

Furthermore, by the semantics of the marking sensitive input nodes, it must hold that  $\mathcal{R}_X^\Phi([v]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_v}$  and  $\mathcal{R}_X^\Phi([w]^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_w}$ , where  $Q_v$  and  $Q_w$  are the query subgraphs associated with  $v$  respectively  $w$ . By the semantics of the marking sensitive union nodes and the assumption regarding the common query subgraph associated with their dependencies, it must also hold that  $\mathcal{R}_X^\Phi([\cup]_v^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_v}$  and  $\mathcal{R}_X^\Phi([\cup]_w^\Phi, \mathcal{C}^\Phi) \subseteq \mathcal{M}_H^{Q_w}$ .

By the semantics of the involved marking-sensitive RETE nodes, for each node in  $LNS([n_1 \rightarrow n_2])$ , each match can only be associated with at most one marking in the node's current result set in  $\mathcal{C}^\Phi$ . Furthermore, by assumption, it holds that  $\mathcal{M}_H^{Q_v} \leq \mathcal{M}_H^{Q_e}$  and  $\mathcal{M}_H^{Q_w} \leq \mathcal{M}_H^{Q_e}$ .

It thus holds that  $\sum_{n^\Phi \in V^{LNS([v \rightarrow w])}} |\mathcal{C}(n^\Phi)| \leq 7 \cdot |\mathcal{C}([v \rightarrow w])|$ .

For each join node  $[\bowtie]$  in  $N$  with associated query subgraph  $Q_{\bowtie}$  and dependencies  $n_l$  and  $n_r$ ,  $N^\Phi$  contains the corresponding marking-sensitive join node  $[\bowtie]^\Phi$  with dependencies  $n_l^\Phi$  and  $n_r^\Phi$  corresponding to  $n_l$  respectively  $n_r$  with associated query subgraphs  $Q_l$  respectively  $Q_r$ , as well as the six nodes from the related request projection structures:  $[\phi > h]_l^\Phi$ ,  $[\pi_{Q_v}]_l^\Phi$ ,  $[\phi > h]_l^\Phi$ ,  $[\phi > h]_r^\Phi$ ,  $[\pi_{Q_v}]_r^\Phi$ , and  $[\phi > h]_r^\Phi$ .

By the semantics of join and marking sensitive join, we know that  $\mathcal{C}(n_l) \subseteq \mathcal{R}_X^\Phi(n_l^\Phi, \mathcal{C}^\Phi) \wedge \mathcal{C}(n_r) \subseteq \mathcal{R}_X^\Phi(n_r^\Phi, \mathcal{C}^\Phi) \Rightarrow \mathcal{C}([\bowtie]) \subseteq \mathcal{R}_X^\Phi([\bowtie]^\Phi, \mathcal{C}^\Phi)$ . For every edge input node  $[v \rightarrow w]$  in  $N$ , we know that  $\mathcal{C}([v \rightarrow w]) \subseteq \mathcal{R}_X^\Phi([\cup]^\Phi, \mathcal{C}^\Phi)$ , where  $[\cup]^\Phi$  is the root node of  $LNS([v \rightarrow w])$ . It is easy to see that thereby, it must hold that  $\mathcal{C}([\bowtie]) \subseteq \mathcal{R}_X^\Phi([\bowtie]^\Phi, \mathcal{C}^\Phi)$ . By Lemma 12 in [BG24b], it follows that  $|\mathcal{C}([\bowtie]^\Phi)| \leq |\mathcal{C}([\bowtie])|$ .

It must then also hold that  $\sum_{n^\Phi \in RPS_l} |\mathcal{C}(n^\Phi)| \leq 3 \cdot |\mathcal{C}(n_l)|$  and  $\sum_{n^\Phi \in RPS_r} |\mathcal{C}(n^\Phi)| \leq 3 \cdot |\mathcal{C}(n_r)|$ .

Due to the tree-like structure of  $(N, p)$ , the marking-sensitive union node at the top of each local navigation structure and each marking-sensitive join can be connected to at most one request projection structure. Moreover, the union node at the top of the local navigation structure or the marking-sensitive join corresponding to  $p$  is not connected to any request projection structure. It thus follows that  $\sum_{n^\Phi \in V^{N^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \sum_{n \in V^N \setminus \{p\}} |\mathcal{C}(n)| + 7 \times |\mathcal{C}(p)|$ .  $\square$

**Lemma A.6** (Consistent configurations for modular extensions of RETE nets localized via  $localize^\Psi$  yield only matches that satisfy the NGC and all admissible matches with a partial match marked  $\infty$  in an extension point). *Let  $H$  be a graph,  $(Q, \psi)$  a graph query,  $(X^\Phi, p_X^\Phi)$  a modular extension of the localized RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \phi)$ , and  $\mathcal{C}$  a configuration for  $(X^\Phi, p_X^\Phi)$  that is consistent for all nodes in  $V^{N^\Phi}$ . Furthermore, let  $x \in \chi(N^\Phi)$ ,  $Q_v = (\{v\}, \emptyset, \emptyset, \emptyset)$  the query subgraph associated with  $x$ , and  $(m', \infty) \in \mathcal{C}(x)$  for some match  $m' : Q_v \rightarrow H$ . It then holds that  $\forall (m, \phi) \in \mathcal{C}(p^\Phi) : m \models \psi \wedge \forall m \in \mathcal{M}_H^Q : m(v) = m'(v) \wedge m \models \psi \Rightarrow (m, \infty) \in \mathcal{C}(p^\Phi)$ .*

*Proof.* We prove the lemma via structural induction over the nested graph condition  $\psi$ .

In the base case, that is, for a nested condition of the form  $\psi = \text{true}$ , satisfaction of the Lemma follows directly from Lemma 1 in the preprint version of our conference paper [BG24b].

We now proceed to showing that, under the induction hypothesis that the lemma holds for a nested graph condition of nesting depth  $d$ , the lemma holds for any nested graph condition of nesting depth  $d + 1$ .

For a nested condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ ,  $N^\Phi$  consists of the localized RETE net for  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = localize(Q)$ , the localized RETE net  $(N_{(Q', \psi')}^\Phi, p_{(Q', \psi')}^\Phi) = localize^\Psi(Q', \psi')$ , a request projection structure  $RPS_l^\infty = RPS^\infty(p_Q^\Phi, N_{(Q', \psi')}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\bowtie]^\Phi$  with left dependency  $p_Q^\Phi$  and right dependency  $p_{(Q', \psi')}^\Phi$ .

In the case where  $a$  is a partial graph morphism from some subgraph  $Q_p \subseteq Q$  into  $Q'$ , the argumentation works analogously.

By Lemma 1 in [BG24b], it immediately follows that  $\forall m \in \mathcal{M}_H^Q : m(v) = m'(v) \Rightarrow (m, \infty) \in \mathcal{C}(p_Q^\Phi)$ . Assuming that the Lemma holds for localized RETE net  $(N_{(Q', \psi')}^\Phi, p_{(Q', \psi')}^\Phi) = localize^\Psi(Q', \psi')$ , by the definition of  $RPS_l^\infty$  it must also hold that  $\forall (m, \infty) \in \mathcal{C}(p_Q^\Phi), m' \in \mathcal{M}_H^{Q'} : m = m' \circ a \wedge m' \models \psi' \Rightarrow (m', \infty) \in \mathcal{C}(p_{(Q', \psi')}^\Phi)$  and  $\forall (m', \phi') \in \mathcal{C}(p_{(Q', \psi')}^\Phi) : m' \models \psi'$ . From the semantics of the marking-sensitive semi-join then follows the satisfaction of the Lemma.

For a nested condition of the form  $\psi = \neg\psi'$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = localize(Q)$ , the localized RETE net  $(N_{(Q, \psi')}^\Phi, p_{(Q, \psi')}^\Phi) = localize^\Psi(Q, \psi')$ , a request projection structure  $RPS_l^\infty = RPS(p_Q^\Phi, N_{(Q, \psi')}^\Phi)$ , and a marking-sensitive anti-join node  $p^\Phi = [\triangleright]^\Phi$  with left dependency  $p_Q^\Phi$  and right dependency  $p_{(Q, \psi')}^\Phi$ .

By Lemma 1 in [BG24b], it immediately follows that  $\forall m \in \mathcal{M}_H^Q : m(v) = m'(v) \Rightarrow (m, \infty) \in \mathcal{C}(p_Q^\Phi)$ . Assuming that the Lemma holds for localized RETE net  $(N_{(Q, \psi')}^\Phi, p_{(Q, \psi')}^\Phi) = localize^\Psi(Q, \psi')$ , by the definition of  $RPS_l^\infty$  it must also hold that  $\forall (m, \infty) \in \mathcal{C}(p_Q^\Phi) : m \models$

$\psi' \Rightarrow (m, \infty) \in \mathcal{C}(p_{(Q, \psi')}^\Phi)$  and  $\forall(m, \phi) \in \mathcal{C}(p_{(Q, \psi')}^\Phi) : m \models \psi'$ . From the semantics of the marking-sensitive anti-join then follows the satisfaction of the Lemma.

For a nested condition of the form  $\psi = \psi_1 \wedge \psi_2$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net  $(N_{(Q, \psi_1)}^\Phi, p_{(Q, \psi_1)}^\Phi) = \text{localize}^\Psi(Q, \psi_1)$ , a request projection structure  $RPS_1^\infty = RPS^\infty(p_Q^\Phi, N_{(Q, \psi_1)}^\Phi)$ , a marking-sensitive semi-join node  $[\bowtie]_1^\Phi$  with left dependency  $p_Q^\Phi$  and right dependency  $p_{(Q', \psi')}^\Phi$ , the localized RETE net  $(N_{(Q, \psi_2)}^\Phi, p_{(Q, \psi_2)}^\Phi) = \text{localize}^\Psi(Q, \psi_2)$ , a request projection structure  $RPS_2^\infty = RPS^\infty([\bowtie]_1^\Phi, N_{(Q, \psi_2)}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\bowtie]_2^\Phi$  with left dependency  $[\bowtie]_1^\Phi$  and right dependency  $p_{(Q, \psi_2)}^\Phi$ .

By Lemma 1 in [BG24b], it immediately follows that  $\forall m \in \mathcal{M}_H^Q : m(v) = m'(v) \Rightarrow (m, \infty) \in \mathcal{C}(p_Q^\Phi)$ . Assuming that the lemma holds for localized RETE net  $(N_{(Q, \psi_1)}^\Phi, p_{(Q, \psi_1)}^\Phi) = \text{localize}^\Psi(Q, \psi_1)$ , by the definition of  $RPS_1^\infty$  it must also hold that  $\forall(m, \infty) \in \mathcal{C}(p_Q^\Phi) : m \models \psi_1 \Rightarrow (m, \infty) \in \mathcal{C}(p_{(Q, \psi_1)}^\Phi)$  and  $\forall(m_1, \phi_1) \in \mathcal{C}(p_{(Q, \psi_1)}^\Phi) : m_1 \models \psi_1$ . Assuming the lemma also holds for  $(N_{(Q, \psi_2)}^\Phi, p_{(Q, \psi_2)}^\Phi) = \text{localize}^\Psi(Q, \psi_2)$ , from the semantics of the marking-sensitive semi-join and the definition of  $RPS_2^\infty$  it then also follows that  $\forall(m, \infty) \in \mathcal{C}(p_Q^\Phi) : m \models \psi_1 \wedge m \models \psi_2 \Rightarrow (m, \infty) \in \mathcal{C}(p_{(Q, \psi_2)}^\Phi)$  and  $\forall(m_2, \phi_2) \in \mathcal{C}(p_{(Q, \psi_2)}^\Phi) : m_2 \models \psi_2$ . From the semantics of the marking-sensitive semi-join then follows the satisfaction of the Lemma.  $\square$

**Lemma A.7** (Execution of modular extensions of localized RETE nets via  $\text{order}^\Psi$  yields consistent configurations for the base RETE net). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  a graph query, and  $(X^\Phi, p_X^\Phi)$  a modular extension of the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = \text{localize}^\Psi(Q, \psi)$ . Furthermore, let  $\mathcal{C}_0^\Phi$  be an arbitrary configuration for  $(X^\Phi, p_X^\Phi)$ . Executing  $(X^\Phi, p_X^\Phi)$  via  $O = \text{order}^\Psi(N^\Phi)$  then yields a configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  that is consistent for all nodes in  $N^\Phi$ .*

*Proof.* We show the correctness of the Lemma via structural induction over the nested graph condition  $\psi$ .

In the base case, where  $\psi = \text{true}$ ,  $O$  is given by  $O = \text{localize}(Q)$ . By Lemma 13 in [BG24b], it follows directly that the lemma holds.

We now proceed to showing that, under the induction hypothesis that the lemma holds for a nested graph condition of nesting depth  $d$ , the lemma holds for any nested graph condition of nesting depth  $d + 1$ .

For a nested condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ ,  $O$  is given by  $O = \text{order}(N_Q^\Phi) \circ \text{toposort}^{-1}(RPS_l^\infty) \circ \text{order}^\Psi(N_{(Q', \psi')}^\Phi) \circ [\bowtie]^\Phi$ . By Lemma 13 in [BG24b], it follows that  $\mathcal{C}_1^\Phi = \text{execute}(\text{order}(N_Q^\Phi), X^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  is consistent for all nodes in  $N_Q^\Phi$ . Since no node in  $N_Q^\Phi$  depends on a node in  $RPS_l^\infty$ ,  $\mathcal{C}_2^\Phi = \text{execute}(\text{toposort}^{-1}(RPS_l^\infty), X^\Phi, H, H_p, \mathcal{C}_1^\Phi)$  must then be consistent for all nodes in  $N_Q^\Phi \cup RPS_l^\infty$ . By the same argument and the induction hypothesis,  $\mathcal{C}_3^\Phi = \text{execute}(\text{order}^\Psi(N_{(Q', \psi')}^\Phi), X^\Phi, H, H_p, \mathcal{C}_2^\Phi)$  is consistent for all nodes in  $N_Q^\Phi \cup RPS_l^\infty \cup N_{(Q', \psi')}^\Phi$  and  $\mathcal{C}_4^\Phi = \text{execute}([\bowtie]^\Phi, X^\Phi, H, H_p, \mathcal{C}_3^\Phi) = \text{execute}(O, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  is consistent for all nodes in  $N^\Phi$ .

By analogous argumentation, the lemma then also holds for a nested condition of the form  $\psi = \neg\psi'$ .

For a nested condition of the form  $\psi = \psi_1 \wedge \psi_2$ ,  $O$  is given by  $order^\Psi(N^\Phi) = order(N_Q^\Phi) \circ toposort^{-1}(RPS_1^\infty) \circ order^\Psi(N_{(Q,\psi_1)}^\Phi) \circ [\times]_1^\Phi \circ toposort^{-1}(RPS_2^\infty) \circ order^\Psi(N_{(Q,\psi_2)}^\Phi) \circ [\times]_2^\Phi$ . By the same argumentation as for the case where  $\psi = \exists(a : Q \rightarrow Q', \psi')$ , the configuration  $\mathcal{C}_1^\Phi = execute(O_1, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  is consistent for all nodes in  $N_Q^\Phi \cup RPS_1^\infty \cup N_{(Q,\psi_1)}^\Phi \cup [\times]_1^\Phi$ , where  $O_1 = order(N_Q^\Phi) \circ toposort^{-1}(RPS_1^\infty) \circ order^\Psi(N_{(Q,\psi_1)}^\Phi) \circ [\times]_1^\Phi$ . By the induction hypothesis and the construction of  $N^\Phi$ , it then holds that  $\mathcal{C}_2^\Phi = execute(O_2, N^\Phi, H, H_p, \mathcal{C}_1^\Phi) = execute(O, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  is consistent for all nodes in  $N^\Phi$ , where  $O_2 = toposort^{-1}(RPS_2^\infty) \circ order^\Psi(N_{(Q,\psi_2)}^\Phi) \circ [\times]_2^\Phi$ .

From the correctness of the base case and the induction step follows the correctness of the lemma.  $\square$

**Lemma A.8** (Execution time of semi-join nodes is linear in the number of matches for the semi-join's left dependency). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $[\times]^\Phi \in V^{N^\Phi}$  a marking-sensitive semi-join node with left dependency  $n_l^\Phi$  and right dependency  $n_r^\Phi$ , and  $\mathcal{C}_0^\Phi$  a configuration such that  $\mathcal{C}_0^\Phi([\times]^\Phi) = \emptyset$ . Executing  $[\times]^\Phi$  via  $\mathcal{C}_1^\Phi = execute([\times]^\Phi, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  then takes  $O(|\mathcal{C}_0^\Phi(n_l^\Phi)|)$  steps.*

*Proof.* In the case where  $\mathcal{C}_0^\Phi([\times]^\Phi) = \emptyset$ ,  $[\times]^\Phi$  can be executed by enumerating all elements  $(m_l, \phi_l) \in \mathcal{C}_0^\Phi(n_l^\Phi)$ , checking whether it holds that  $\exists(m_r, \phi_r) \in \mathcal{C}^\Phi(n_r^\Phi) : m_l|_{Q \cap} = m_r|_{Q \cap}$ , and adding them to  $\mathcal{C}_0^\Phi([\times]^\Phi)$  if the condition holds. Assuming efficient indexing structures that allow lookup and insertion times linear in match size, this takes  $O(|\mathcal{C}_0^\Phi(n_l^\Phi)|)$  steps.  $\square$

**Lemma A.9** (Execution time of anti-join nodes is linear in the number of matches for the semi-join's left dependency). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $[\triangleright]^\Phi \in V^{N^\Phi}$  a marking-sensitive anti-join node with left dependency  $n_l^\Phi$  and right dependency  $n_r^\Phi$ , and  $\mathcal{C}_0^\Phi$  a configuration such that  $\mathcal{C}_0^\Phi([\triangleright]^\Phi) = \emptyset$ . Executing  $[\triangleright]^\Phi$  via  $\mathcal{C}_1^\Phi = execute([\triangleright]^\Phi, N^\Phi, H, H_p, \mathcal{C}_0^\Phi)$  then takes  $O(|\mathcal{C}_0^\Phi(n_l^\Phi)|)$  steps.*

*Proof.* In the case where  $\mathcal{C}_0^\Phi([\triangleright]^\Phi) = \emptyset$ ,  $[\triangleright]^\Phi$  can be executed by enumerating all elements  $(m_l, \phi_l) \in \mathcal{C}_0^\Phi(n_l^\Phi)$ , checking whether it holds that  $\nexists(m_r, \phi_r) \in \mathcal{C}^\Phi(n_r^\Phi) : m_l|_{Q \cap} = m_r|_{Q \cap}$ , and adding them to  $\mathcal{C}_0^\Phi([\triangleright]^\Phi)$  if the condition holds. Assuming efficient indexing structures that allow lookup and insertion times linear in match size, this takes  $O(|\mathcal{C}_0^\Phi(n_l^\Phi)|)$  steps.  $\square$

**Lemma A.10** (RETE net localization via  $localize^\Psi$  introduces an overhead of at most factor 10 on the number of computed matches). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^\Phi, p^\Phi) = localize^\Psi(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  corresponding to  $(N, p)$ . It then holds that  $\sum_{n^\Phi \in V^{N^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^N} |\mathcal{C}(n)|$ .*

*Proof.* We show the correctness of the theorem via structural induction over the nested graph condition  $\psi$ .

In the base case where  $\psi = true$ , it holds that  $localize^\Psi(Q, \psi) = localize(N, p)$ . From Lemma A.5, it then immediately follows that  $\sum_{n^\Phi \in V^{N^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^N} |\mathcal{C}(n)|$ .

We now proceed to showing that, under the induction hypothesis that the lemma holds for a nested graph condition of nesting depth  $d$ , the lemma holds for any nested graph condition of nesting depth  $d + 1$ .

For a nested condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ ,  $N^\Phi$  consists of the localized RETE net for  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net  $(N_{(Q', \psi')}^\Phi, p_{(Q', \psi')}^\Phi) = \text{localize}^\Psi(Q', \psi')$ , a request projection structure  $RPS_l^\infty = RPS^\infty(p_Q^\Phi, N_{(Q', \psi')}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\bowtie]^\Phi$ .

By the construction described in [BG23],  $(N, p)$  then consists of a RETE net for the plain graph query  $Q$ ,  $(N_Q, p_Q)$ , the RETE net for the extended graph query  $(Q', \psi')$ ,  $(N', p')$ , and a semi-join  $p = [\bowtie]$  with dependencies  $p_Q$  and  $p'$ .

By Lemma A.5, we know that  $\sum_{n^\Phi \in V^{N_Q^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \sum_{n \in V^{N_Q} \setminus \{p_Q\}} |\mathcal{C}(n)| + 7 \cdot |\mathcal{C}(p_Q)|$ , since  $\text{localize}(Q) = \text{localize}(N_Q, p_Q)$ .

In the worst case, the extension points of  $(N_{(Q', \psi')}^\Phi, p_{(Q', \psi')}^\Phi)$  already contain all possible single-vertex matches into  $H$  even without the additional input via  $RPS_l^\infty$ . Therefore, by the induction hypothesis, it must hold that  $\sum_{n^\Phi \in V^{N_{(Q', \psi')}^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^{N'}} |\mathcal{C}(n)|$ .

From Lemma 12 in [BG24b], it follows that  $\sum_{n_Q^\Phi \in V^{RPS_l^\infty}} |\mathcal{C}^\Phi(n_Q^\Phi)| \leq 3 \cdot |\mathcal{C}(p_Q)|$ .

Finally, from Theorem 4.3 and Lemma 12 in [BG24b], it follows that  $|\mathcal{C}^\Phi(p^\Phi)| \leq |\mathcal{C}([\bowtie])|$ . Consequently, it must hold that  $\sum_{n^\Phi \in V^{N^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^N} |\mathcal{C}(n)|$ .

For a nested condition of the form  $\psi = \neg\psi'$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net  $(N_{(Q, \psi')}^\Phi, p_{(Q, \psi')}^\Phi) = \text{localize}^\Psi(Q, \psi')$ , a request projection structure  $RPS_l^\infty = RPS(p_Q^\Phi, N_{(Q, \psi')}^\Phi)$ , and a marking-sensitive anti-join node  $p^\Phi = [\triangleright]^\Phi$ .

By the construction described in [BG23],  $(N, p)$  then consists of a RETE net for the plain graph query  $Q$ ,  $(N_Q, p_Q)$ , the RETE net for the extended graph query  $(Q, \psi')$ ,  $(N', p')$ , and an anti-join  $p = [\triangleright]$  with dependencies  $p_Q$  and  $p'$ .

By analogous argumentation as for the existential case, it must then also hold that  $\sum_{n^\Phi \in V^{N^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^N} |\mathcal{C}(n)|$ .

For a nested condition of the form  $\psi = \psi_1 \wedge \psi_2$ ,  $N^\Phi$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ , the localized RETE net  $(N_{(Q, \psi_1)}^\Phi, p_{(Q, \psi_1)}^\Phi) = \text{localize}^\Psi(Q, \psi_1)$ , a request projection structure  $RPS_1^\infty = RPS^\infty(p_Q^\Phi, N_{(Q, \psi_1)}^\Phi)$ , a marking-sensitive semi-join node  $[\bowtie]_1^\Phi$ , the localized RETE net  $(N_{(Q, \psi_2)}^\Phi, p_{(Q, \psi_2)}^\Phi) = \text{localize}^\Psi(Q, \psi_2)$ , a request projection structure  $RPS_2^\infty = RPS^\infty([\bowtie]_1^\Phi, N_{(Q, \psi_2)}^\Phi)$ , and a marking-sensitive semi-join node  $p^\Phi = [\bowtie]_2^\Phi$ .

By the construction described in [BG23],  $(N, p)$  then consists of a RETE net for the plain graph query  $Q$ ,  $(N_Q, p_Q)$ , the RETE net for the extended graph query  $(Q, \psi_1)$ ,  $(N_1, p_1)$ , the RETE net for the extended graph query  $(Q, \psi_2)$ ,  $(N_2, p_2)$ , a semi-join  $[\bowtie]_1$  with dependencies  $p_Q$  and  $p_1$ , and a semi-join  $p = [\bowtie]_2$  with dependencies  $[\bowtie]_1$  and  $p_2$ .

As in the existential case, it holds that  $\sum_{n^\Phi \in V^{N_Q^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \sum_{n \in V^{N_Q} \setminus \{p_Q\}} |\mathcal{C}(n)| + 7 \cdot |\mathcal{C}(p_Q)|$ . Furthermore, it must hold that  $\sum_{n^\Phi \in V^{N_{(Q, \psi_1)}^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^{N_1}} |\mathcal{C}(n)|$  and  $\sum_{n^\Phi \in V^{N_{(Q, \psi_2)}^\Phi}} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^{N_2}} |\mathcal{C}(n)|$  by the induction hypothesis.

From Lemma 12 in [BG24b], it follows that  $\sum_{n_Q^\Phi \in V^{RPS_1^\infty}} |\mathcal{C}^\Phi(n_Q^\Phi)| \leq 3 \cdot |\mathcal{C}(p_Q)|$ .

We furthermore know by Theorem 4.3 that  $\mathcal{C}^\Phi([\bowtie]_1^\Phi)$  can only contain matches for  $Q$  that satisfy  $\psi_1$  and  $\mathcal{C}([\bowtie]_1)$  contains all such matches. From Lemma 12 in [BG24b], it thus follows that  $|\mathcal{C}^\Phi([\bowtie]_1^\Phi)| \leq |\mathcal{C}([\bowtie]_1)|$ . Consequently, it must hold that  $\sum_{n_Q^\Phi \in V^{RPS_2^\infty}} |\mathcal{C}^\Phi(n_Q^\Phi)| \leq$

$3 \cdot |\mathcal{C}([\times]_1)|$ . By Theorem 4.3 and Lemma 12 in [BG24b], it also follows that  $|\mathcal{C}^\Phi([\times]_2^\Phi)| \leq |\mathcal{C}([\times]_2)|$ .

We thus know that  $\sum_{n^\Phi \in V^N} |\mathcal{C}^\Phi(n^\Phi)| \leq 10 \cdot \sum_{n \in V^N} |\mathcal{C}(n)|$ .

From the correctness of base case and induction step then follows the correctness of the lemma.  $\square$

**Lemma A.11** (Consistent configurations for modular extensions of RETE nets localized via  $localize^{sat}$  yield all subgraph satisfaction dependent matches for an extended graph query with marking  $\infty$ ). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  an extended graph query, and  $\mathcal{C}^\Phi$  a consistent configuration for the modular extension  $(X^{sat}, p^{sat})$  of  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$ . It then holds that  $\forall m \in \mathcal{M}_H^Q : m$  is subgraph satisfaction dependent  $\Rightarrow (m, \infty) \in \mathcal{C}^\Phi(p^{sat})$ .*

*Proof.* We show the correctness of the lemma via structural induction over  $\psi$ .

In the base case of  $\psi = \text{true}$ , the lemma trivially holds, since no match is subgraph satisfaction dependent in this case.

We now proceed to showing that, under the induction hypothesis that the lemma holds for a nested graph condition of nesting depth  $d$ , the lemma holds for any nested graph condition of nesting depth  $d + 1$ .

For a nested condition of the form  $\psi = \neg\psi'$ , the correctness of the lemma follows directly from the induction hypothesis.

For a nested condition of the form  $\psi = \psi_1 \wedge \psi_2$ , the correctness of the lemma follows directly from the induction hypothesis and the semantics of the marking-sensitive union node.

For a nested condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ ,  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$  consists of the localized RETE net for the plain pattern  $Q$ ,  $(N_Q^\Phi, p_Q^\Phi) = localize(Q)$ , the localized RETE net for the plain pattern  $Q'$ ,  $(N_{Q'}^\Phi, p_{Q'}^\Phi) = localize(Q')$ , the RETE net  $(N_{(Q', \psi')}^{sat}, p_{(Q', \psi')}^{sat}) = localize^{sat}(Q', \psi')$ , a marking-sensitive union node  $[\cup]^\Phi$  with dependencies  $p_{Q'}^\Phi$  and  $p_{(Q', \psi')}^{sat}$ , a request projection structure  $RPS_r^\infty = RPS^\infty([\cup]^\Phi, N_Q^\Phi)$ , and a marking-sensitive semi-join node  $p^{sat} = [\times]^\Phi$  with left dependency  $p_Q^\Phi$  and right dependency  $[\cup]^\Phi$ .

By Lemma 1 in [BG24b], the induction hypothesis, and the semantics definition of the marking-sensitive union node, it must then hold that  $\forall m' \in \mathcal{M}_H^{Q'} : m'(Q') \cap H_p \neq \emptyset$  or  $m'$  is subgraph satisfaction dependent  $\Rightarrow (m', \infty) \in \mathcal{C}^\Phi([\cup]^\Phi)$ .

From the construction of  $RPS_r^\infty$  and Lemma 1 in [BG24b], it follows that  $\forall m \in \mathcal{M}_H^Q : (\exists m' \in \mathcal{M}_H^{Q'} : m = m' \circ a \wedge (m'(Q') \cap H_p \neq \emptyset \text{ or } m' \text{ is subgraph satisfaction dependent})) \Rightarrow (m, \infty) \in \mathcal{C}^\Phi(p_Q^\Phi)$ .

From the semantics of the marking-sensitive semi-join then follows the correctness of the lemma.

From the satisfaction of the base case and the induction step follows the correctness of the lemma.  $\square$

**Lemma A.12** (Execution of modular extensions of localized RETE nets via  $order^{sat}$  yields consistent configurations for the base RETE net). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  a graph query, and  $(X^{sat}, p_X^{sat})$  a modular extension of the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = localize^{sat}(Q, \psi)$ . Furthermore, let  $\mathcal{C}_0^\Phi$  be an arbitrary configuration*

for  $(X^{sat}, p_X^{sat})$ . Executing  $(X^{sat}, p_X^{sat})$  via  $O = \text{order}^{sat}(N^{sat})$  then yields a configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^{sat}, H, H_p, \mathcal{C}_0^\Phi)$  that is consistent for all nodes in  $N^{sat}$ .

*Proof.* We show the correctness of the lemma via structural induction over  $\psi$ .

In the base case of  $\psi = \text{true}$ , the lemma trivially holds.

We now proceed to showing that, under the induction hypothesis that the lemma holds for a nested graph condition of nesting depth  $d$ , the lemma holds for any nested graph condition of nesting depth  $d + 1$ .

For a nested condition of the form  $\psi = \neg\psi'$ , the correctness of the lemma follows directly from the induction hypothesis.

For a nested condition of the form  $\psi = \psi_1 \wedge \psi_2$ , the correctness of the lemma follows directly from the induction hypothesis.

For a nested condition of the form  $\psi = \exists(a : Q \rightarrow Q', \psi')$ , the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = \text{localize}^{sat}(Q, \psi)$  consists of  $(N_Q^\Phi, p_Q^\Phi) = \text{localize}(Q)$ ,  $(N_{Q'}^\Phi, p_{Q'}^\Phi) = \text{localize}(Q')$ ,  $(N_{(Q', \psi')}^\Phi, p_{(Q', \psi')}^\Phi) = \text{localize}^{sat}(Q', \psi')$ ,  $[\cup]^\Phi$ ,  $RPS_r^\infty = RPS([\cup]^\Phi, N_Q^\Phi)$ , and  $p^{sat} = [\times]^\Phi$ . The execution order for  $(N^{sat}, p^{sat})$  is then given by  $\text{order}^{sat}(N^{sat}) = \text{order}(N_Q^\Phi) \circ \text{order}^{sat}(N_{(Q', \psi')}^\Phi) \circ [\cup]^\Phi \circ \text{toposort}(RPS_r^\infty)^{-1} \circ \text{order}(N_Q^\Phi) \circ [\times]^\Phi$ .

By Lemma 13 in [BG24b],  $\mathcal{C}_1^\Phi = \text{execute}(\text{order}(N_Q^\Phi), N^{sat}, H, H_p, \mathcal{C}_0^\Phi)$  is consistent for all nodes in  $N_Q^\Phi$ .  $\mathcal{C}_2^\Phi = \text{execute}(\text{order}^{sat}(N_{(Q', \psi')}^\Phi), N^{sat}, H, H_p, \mathcal{C}_1^\Phi)$  is then also consistent for all nodes in  $N_{Q'}^\Phi$  by the induction hypothesis. By Lemma 13 in [BG24b],  $\mathcal{C}^\Phi = \text{execute}([\cup]^\Phi \circ \text{toposort}(RPS_r^\infty)^{-1} \circ \text{order}(N_Q^\Phi) \circ [\times]^\Phi, N^{sat}, H, H_p, \mathcal{C}_2^\Phi)$  must then be consistent for all nodes in  $N^{sat}$ .

From the satisfaction of the base case and the induction step follows the correctness of the lemma.  $\square$

**Lemma A.13** (RETE net localization via  $\text{localize}^{sat}$  introduces an overhead of at most factor 25 on the number of computed matches). *Let  $H$  be an edge-dominated graph,  $H_p \subseteq H$ ,  $(N, p)$  a RETE net created via the procedure described in [BG23] for the extended graph query  $(Q, \psi)$ ,  $\mathcal{C}$  a consistent configuration for  $(N, p)$  for host graph  $H$ , and  $\mathcal{C}^\Phi$  a consistent configuration for the marking-sensitive RETE net  $(N^{sat}, p^{sat}) = \text{localize}^{sat}(Q, \psi)$  for host graph  $H$  and relevant subgraph  $H_p$  corresponding to  $(N, p)$ . It then holds that  $\sum_{n^{sat} \in V^{N^{sat}}} |\mathcal{C}^\Phi(n^{sat})| \leq 25 \cdot \sum_{n \in V^N} |\mathcal{C}(n)|$ .*

*Proof.* For nested queries of  $(Q, \psi)$  with the form  $(Q', \text{true})$ ,  $\text{localize}^{sat}$  only introduces a dummy node to  $N^{sat}$ . The result set for any dummy node in  $N^{sat}$  is always empty. Hence, the result sets of these nodes do not contribute to configuration size. Moreover,  $\text{localize}^{sat}$  does not introduce any additional nodes for nested queries with the form  $(Q', \neg\psi')$ .

For each nested query with the form  $(Q', \psi_1 \wedge \psi_2)$ ,  $N^{sat}$  contains one marking-sensitive union node  $[\cup]^\Phi$ , whereas  $N$  contains a RETE subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  computing matches for  $Q'$  as well as two semi-joins. Since the result set for  $[\cup]^\Phi$  can only contain at most one tuple for each match for  $Q'$  and the result set for  $p_{Q'}^\Phi$  contains all such matches, it follows that  $|\mathcal{C}^\Phi([\cup]^\Phi)| \leq |\mathcal{C}(p_{Q'}^\Phi)|$ .

For each nested query with the form  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$ ,  $N^{sat}$  contains the subnets  $(N_{Q'}^\Phi, p_{Q'}^\Phi) = \text{localize}(Q')$  and  $(N_{Q''}^\Phi, p_{Q''}^\Phi) = \text{localize}(Q'')$ , a marking-sensitive union node  $[\cup]^\Phi$ , a marking-sensitive semi-join  $[\times]^\Phi$ , as well as the request projection structure  $RPS_r^\infty = RPS^\infty([\cup]^\Phi, N_{Q'}^\Phi)$ .  $N$  contains a RETE subnet  $(N_{Q'}^\Phi, p_{Q'}^\Phi)$  computing matches for  $Q'$  and a semi-join  $[\times]$ .

By Lemma A.5 and because the additional dependency on  $RPS_l^\infty$  does not affect the worst case of a fully populated RETE net  $(N_Q^\Phi, p_Q^\Phi)$ , it must hold that  $\sum_{n_{Q'}^\Phi \in V^{N_{Q'}^\Phi}} |\mathcal{C}^\Phi(n_{Q'}^\Phi)| \leq 10 \cdot \sum_{n_{Q'} \in V^{N_{Q'}}} |\mathcal{C}(n_{Q'})|$ . Also, by the semantics of the marking-sensitive semi-join and because the result set for  $p_{Q'}^\Phi$  contains each match for  $Q'$  at most once by Lemma 12 in [BG24b], it must hold that  $\sum_{(m, \phi) \in \mathcal{C}^\Phi([\times]^\Phi)} |m| \leq \sum_{m \in \mathcal{C}(p_{Q'}^\Phi)} |m| \leq \sum_{n_{Q'} \in V^{N_{Q'}}} |\mathcal{C}(n_{Q'})|$ .

Any nested query with the form  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$  must have some nested query  $(Q'', \exists(a : Q'' \rightarrow Q''', \psi''))$  or  $(Q'', \text{true})$  in order for  $\psi$  by the definition of nested graph conditions. In either case,  $N$  has to contain a corresponding subnet  $(N_{Q''}, p_{Q''})$  computing matches for  $Q''$ . It then holds for this subnet  $N_{Q''}$  by Theorem 3.10 that  $\sum_{n_{Q''}^\Phi \in V^{N_{Q''}^\Phi}} |\mathcal{C}^\Phi(n_{Q''}^\Phi)| \leq 10 \cdot \sum_{n_{Q''} \in V^{N_{Q''}}} |\mathcal{C}(n_{Q''})|$ . It also must hold that  $|\mathcal{C}^\Phi([\cup]^\Phi)| \leq |\mathcal{C}(p_{Q''})| \leq \sum_{n_{Q''} \in V^{N_{Q''}}} |\mathcal{C}(n_{Q''})|$  and  $\sum_{n^\Phi \in V^{RPS_l^\infty}} |\mathcal{C}^\Phi(n^\Phi)| \leq 3 \cdot \sum_{n_{Q''} \in V^{N_{Q''}}} |\mathcal{C}(n_{Q''})|$ .

Since  $\psi$  is a tree of subconditions, it follows that each subnet  $(N_{Q''}, p_{Q''})$  must be counted again like this for at most one parent query  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$ . For any nested query of the form  $(Q', \text{true})$ ,  $\mathcal{C}^\Phi$  stores at most  $14 \cdot \sum_{n_{Q''} \in V^{N_{Q'}}} |\mathcal{C}(n_{Q'})|$  matches. For any nested query of the form,  $(Q', \exists(a : Q' \rightarrow Q'', \psi'))$ ,  $\mathcal{C}^\Phi$  stores at most  $(11 + 14) \cdot \sum_{n_{Q''} \in V^{N_{Q'}}} |\mathcal{C}(n_{Q'})|$  matches.

In the worst case, it must thus hold that  $\sum_{n^{sat} \in V^{N^{sat}}} |\mathcal{C}^\Phi(n^{sat})| \leq 25 \cdot \sum_{n \in V^N} |\mathcal{C}(n)|$ .  $\square$

**Lemma A.14** (Execution of modular extensions of localized RETE nets via  $order^\Delta$  yields consistent configurations for the base RETE net). *Let  $H$  be a graph,  $H_p \subseteq H$ ,  $(Q, \psi)$  a graph query, and  $(X^\Delta, p_X^\Delta)$  a modular extension of the marking-sensitive RETE net  $(N^\Delta, p^\Delta) = \text{localize}^\Delta(Q, \psi)$ . Furthermore, let  $\mathcal{C}_0^\Phi$  be an arbitrary configuration for  $(X^\Delta, p_X^\Delta)$ . Executing  $(X^\Delta, p_X^\Delta)$  via  $O = order^\Delta(N^\Delta)$  then yields a configuration  $\mathcal{C}^\Phi = \text{execute}(O, N^\Delta, H, H_p, \mathcal{C}_0^\Phi)$  that is consistent for all nodes in  $N^\Delta$ .*

*Proof.* Follows directly from Lemma 13 in [BG24b] and Lemmata A.12 and A.7.  $\square$

## APPENDIX B. QUERIES

Information regarding the metamodel, that is, the type graph for the synthetic and real abstract syntax graph scenarios can be found in [BCDM14]. For a visualization of the type graph for the LDBC scenario, see [AAA<sup>+</sup>24].

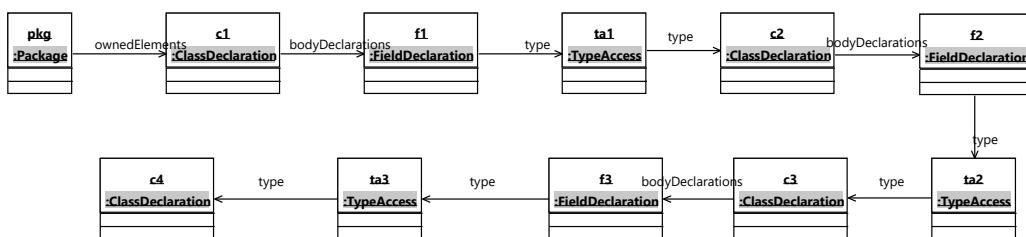


FIGURE 24. Simple path query for the synthetic abstract syntax graph scenario

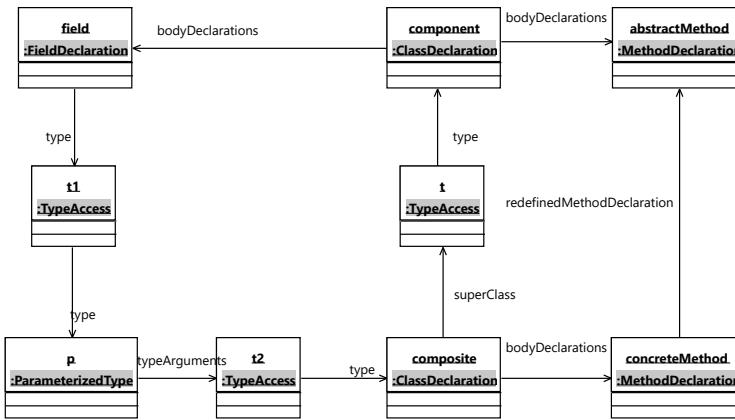


FIGURE 25. Composite query for the real abstract syntax graph scenario

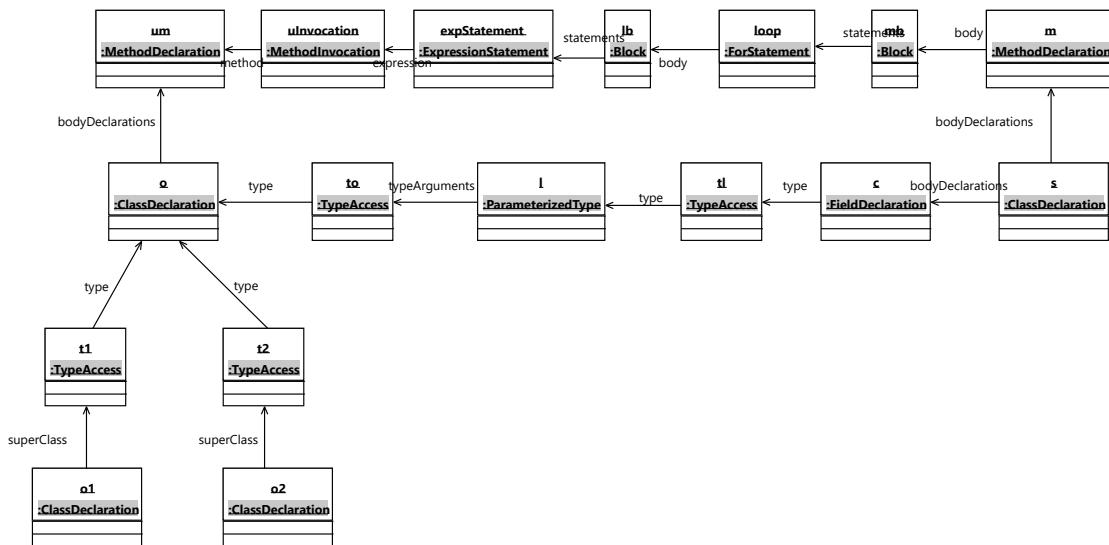


FIGURE 26. Observer query for the real abstract syntax graph scenario

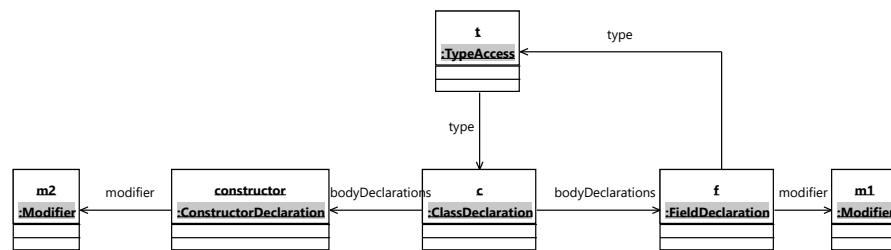


FIGURE 27. Singleton query for the real abstract syntax graph scenario

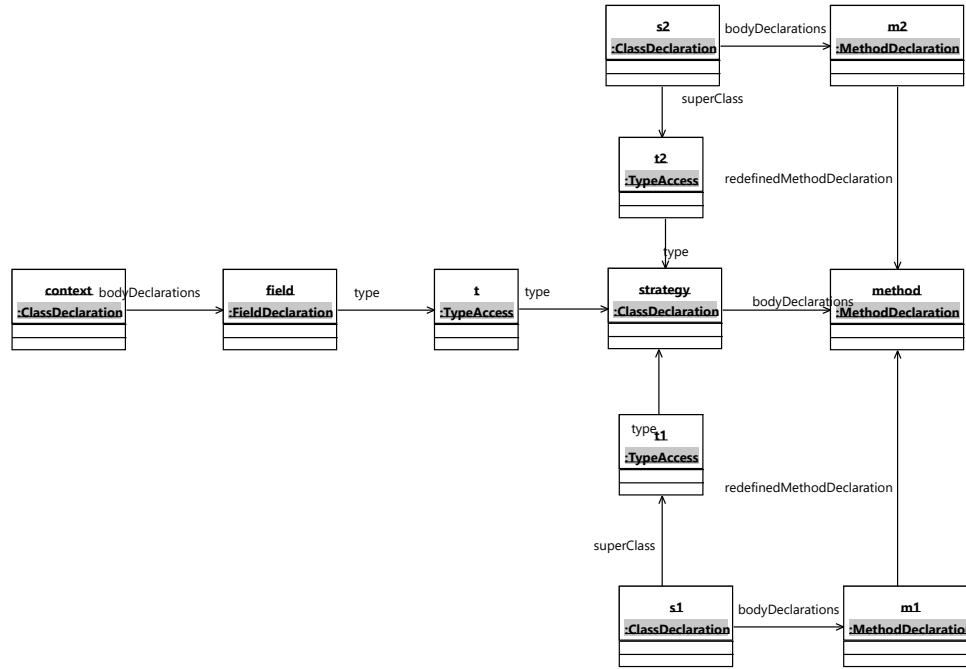


FIGURE 28. Strategy query for the real abstract syntax graph scenario

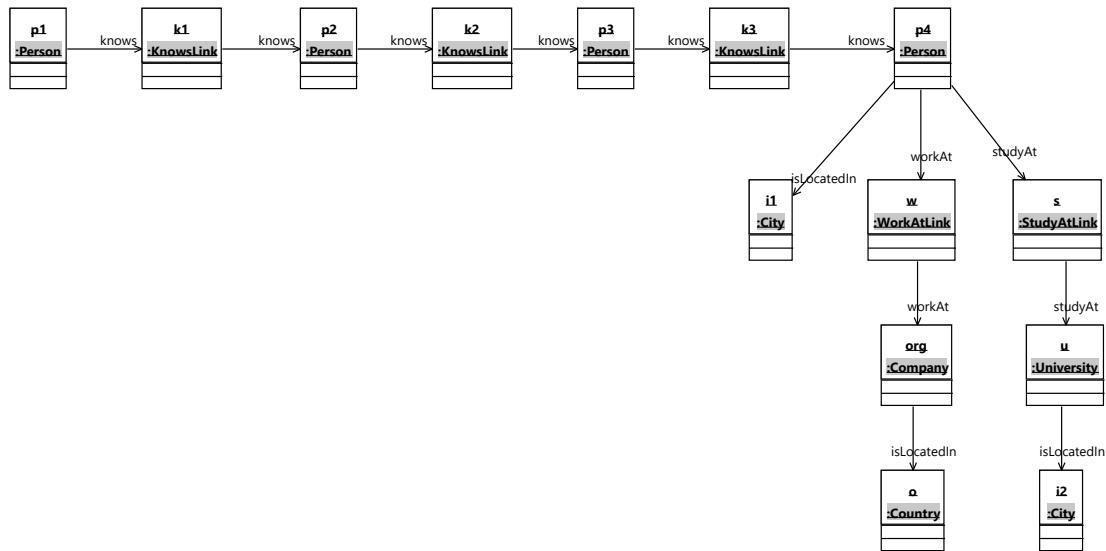


FIGURE 29. Query "Interactive 1" for the LDBC scenario

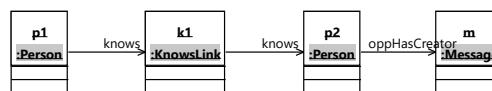


FIGURE 30. Query "Interactive 2" for the LDBC scenario

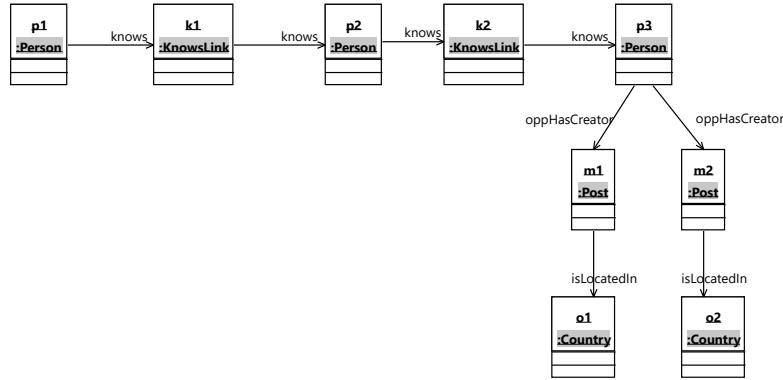


FIGURE 31. Query “Interactive 3” for the LDBC scenario

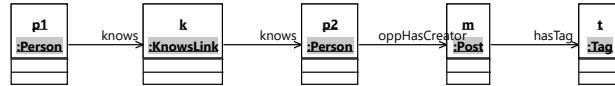


FIGURE 32. Query “Interactive 4” for the LDBC scenario

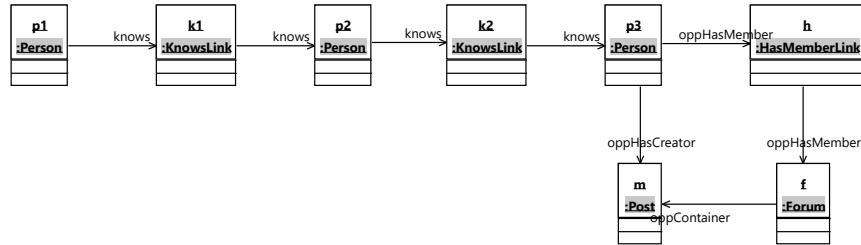


FIGURE 33. Query “Interactive 5” for the LDBC scenario

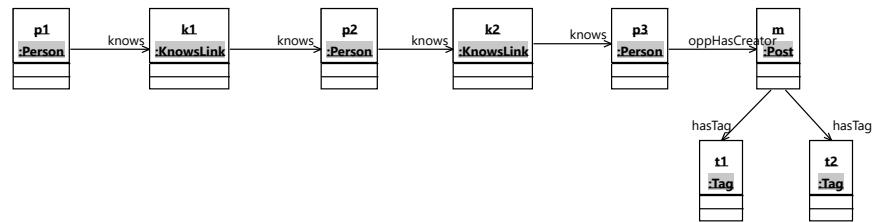


FIGURE 34. Query “Interactive 6” for the LDBC scenario

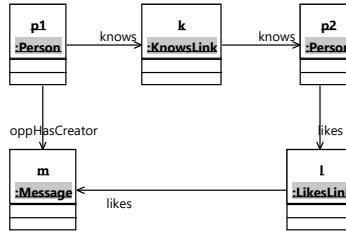


FIGURE 35. Query “Interactive 7” for the LDBC scenario

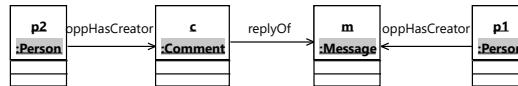


FIGURE 36. Query “Interactive 8” for the LDBC scenario

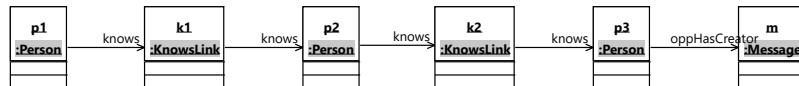


FIGURE 37. Query “Interactive 9” for the LDBC scenario

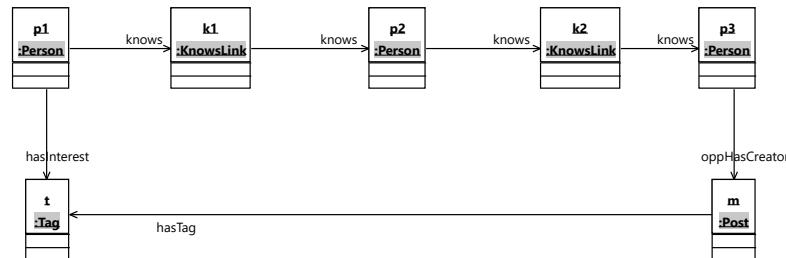


FIGURE 38. Query “Interactive 10” for the LDBC scenario

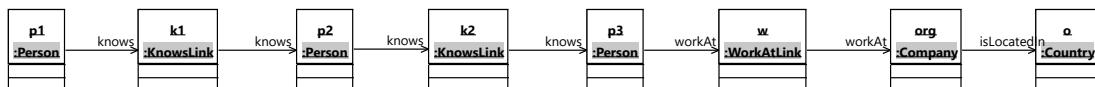


FIGURE 39. Query “Interactive 11” for the LDBC scenario

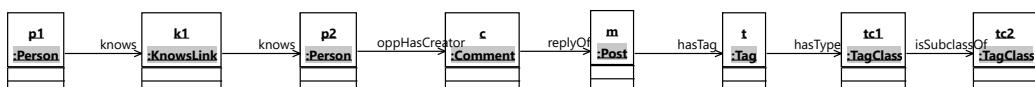


FIGURE 40. Query “Interactive 12” for the LDBC scenario

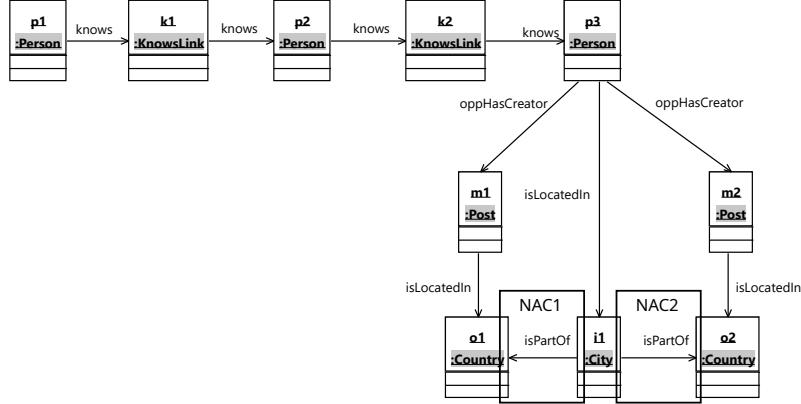


FIGURE 41. Query “Interactive 3 NGC” for the LDBC scenario. The query includes a nested graph condition of the form  $\neg\exists(a_1 : Q \rightarrow Q'_1, \text{true}) \wedge \neg\exists(a_2 : Q \rightarrow Q'_2, \text{true})$ , with the parts of  $Q'_1$  and  $Q'_2$  that are not in the image of  $a_1$  respectively  $a_2$  inside the boxes labeled “NAC1” and “NAC2”

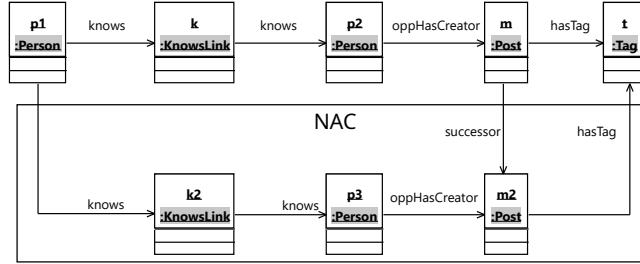


FIGURE 42. Query “Interactive 4 NGC” for the LDBC scenario. The query includes a nested graph condition of the form  $\neg\exists(a : Q \rightarrow Q', \text{true})$ , with the part of  $Q'$  that is not in the image of  $a$  inside the box labeled “NAC”