



ELECTRONIC COMMUNICATIONS OF THE EASST

Volume 4

**Proc. of 2nd International Workshop on Graph and Model
Transformation 2006**

ISSN 1863-2122

Search Trees for Distributed Graph Transformation Systems

Ulrike Ranger and Mathias Lüstraeten

12 pages, 2006

Guest Editors: Gabor Karsai, Gabriele Taentzer

Managing Editors: Tiziana Margaria, Julia Padberg, Gabriele Taentzer

Homepage of ELECTRONIC COMMUNICATIONS OF THE EASST: <http://www.easst.org/eceasst>

Search Trees for Distributed Graph Transformation Systems

Ulrike Ranger and Mathias Lüstraeten*

*RWTH Aachen University, Department of Computer Science 3,
Ahornstrasse 55, 52074 Aachen, Germany
[ranger|matlue]@i3.informatik.rwth-aachen.de

***Abstract.** Graph transformation systems, like PROGRES and Fujaba, can be used for modeling software systems of various domains, and support the automatic generation of executable code. A graph transformation rule is executed only if the pattern of the transformation's left-hand side is found in the graph. The search for the pattern has an exponential worst-case complexity. In many cases, the average complexity can be reduced using search tree algorithms in the code generation phase. When modeling distributed graph transformations, the communication overhead between the coupled applications largely affects the pattern matching performance. Therefore, we present an approach for adapting existing search tree algorithms for the efficient search of distributed graph patterns. Our algorithm divides the distributed graph pattern into several sub-patterns such that every sub-pattern affects solely the graph of exactly one coupled application. The results of these sub-patterns are used to determine the match of the entire graph pattern.*

Keywords: Graph Transformations, Search Trees, Distributed Systems

1 Introduction

Graph transformation systems (GTS), like PROGRES [Sch91] and Fujaba [FNTZ00], can be used to model software systems in a visual way. Additionally, they facilitate the generation of executable code like C or Java for the modeled software system. Several large projects of various domains have been developed using GTS, but they lack support for specifying distributed systems. In our project, we extend GTS for appropriate concepts including the visual specification of distributed graph transformations.

To support distributed graph transformations, we have to consider three aspects: First, the syntax and semantics of distributed graph transformations has to be defined (*specification level*). Second, a concept for the generation of efficient code has to be developed considering special requirements like communication costs (*code generation level*). Third, a runtime environment must be designed and implemented, which supports the execution of the generated applications (*runtime level*). This paper focuses on the code generation level. The specification level is described in [RS06].

Our code generation approach for distributed graph transformations is based on the existing *search tree* algorithms of PROGRES and Fujaba. Search trees allow to reduce the complexity for searching graph patterns specified within a graph transformation, as this search has an exponential worst-case complexity.

Regarding distributed graph transformations, also the high communication costs within a distributed system have to be considered within the search trees. Therefore, we have to adapt the code generation algorithm by the following approach: The distributed graph pattern specified within a graph transformation is divided into several sub-patterns, such that every sub-pattern affects solely the graph of exactly one coupled application. These sub-patterns are sent to the coupled applications, thus reducing the communication costs in comparison of querying the applications for every single remote element of the pattern. As sub-patterns may depend on formerly queried pattern elements, the dependencies to other sub-patterns have to be analyzed. Thus, the sub-patterns are executed with the former results as parameters, and their results determine the match of the entire distributed graph pattern.

The paper is structured as follows: In section 2 we shortly introduce distributed graph transformations considering a simple Publishing Trade as an example. Search trees, which are used to generate efficient executable code for local graph transformations, are described in section 3. Section 4 presents our approach for adapting the presented search tree algorithm to distributed graph transformations. A summary and an outlook on future work are given in section 5.

2 Specifying Distributed Graph Transformation Systems

In this section, we shortly introduce the visual modeling of distributed graph transformation systems at specification level, which is described in [RS06]. This comprises the graph schema of the distributed system, which will be described in section 2.1, and the visual specification of the dynamic behavior at runtime presented in section 2.2.

In our approach, a distributed system is modeled by different specifications, each modeling one module of the software system. At runtime, each specification is executed separately in an own application storing the according host graph¹. These applications are coupled at runtime by executing distributed graph transformations, which affect several applications simultaneously. The specifier only has to develop appropriate transformations for the distributed system, as the GTS generates adequate executable code, and the runtime environment automatically performs their execution.

2.1 Graph Schema of a Distributed Graph Transformation Systems

To illustrate our concepts, we use a simple *Publishing Trade* as example, whose static structure is depicted in figure 1. Here, we assume that a book publisher already has an existing module *Publisher*, managing all Books and their Authors². As the Publisher has decided to sell his Books in an online shop, a new module *Publisher Shop* for this purpose shall be developed, using the existing specification of the *Publisher* module. The new *Publisher Shop* manages the necessary data for selling the Books of the

¹Until now, our approach has the restriction that every specification is executed exactly once within a distributed system. We are developing appropriate concepts to fill this gap including the introduction of *roles* for executed applications and the usage of adequate *role restrictions* within graph transformations.

²For the sake of simplicity, every BOOK is written by exactly one Author.

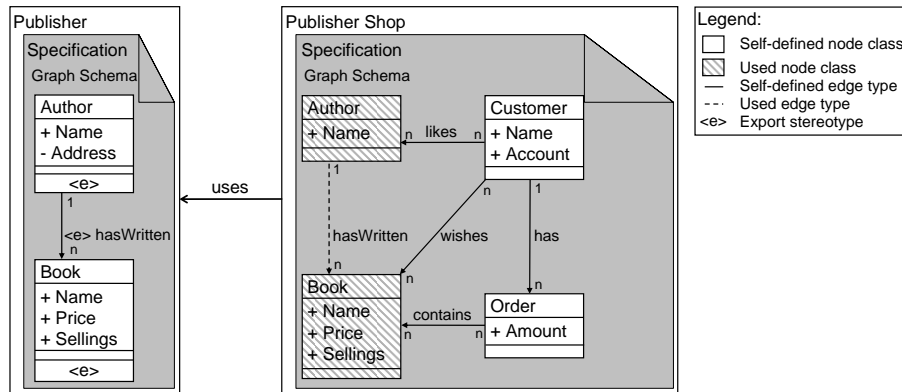


Figure 1: Graph schema for the distributed Publishing Trade

Publisher including Customers and their according Orders. Additionally, a Customer can have a wish list for desired Books, which is modeled by the wishes-edge. For advertising Books, the Publisher Shop memorizes the Authors liked by the Customers (likes-edge) when the Customer has bought a Book.

The Publisher and the Publisher Shop are modeled by different specifications shown in figure 1. As the Publisher manages all information about the Books and the Authors, the Publisher Shop needs access to these data. Therefore, the Publisher defines an *export interface* by using the stereotype `<e>`. Each graph schema element marked with this stereotype forms the interface of the specification. An exported node class consists of the class name and the public attributes and methods defined within the node class.

The exported specification interface can be *used* by other specifications. Although the used graph schema elements are *read-only*, they can be used within the specification in nearly the same way as local, self-defined schema elements. In this way, e.g. edges can be defined between self-defined and used node classes or just between used node classes. To distinguish between used and self-defined schema elements, the used elements are depicted as striped rectangles resp. dashed arrows within the specification. With this mechanism, the Publisher Shop uses the interface of the Publisher and defines new edges between self-defined and used node classes, e.g. the wishes-edge between the self-defined Customer node class, and the used Book node class realizing the wish list.

2.2 Distributed Graph Transformations

The dynamic behavior of a distributed system has to be specified within *distributed graph transformations*. In distributed transformations, edges and nodes of used types can be applied in the same way as of self-defined types but they address remote objects instead of local objects. At runtime, every node and edge only exists once within a distributed system, i.e. in the host graph of the application, which is based on the specification defining its type. Coupled applications only store *reference nodes* on remote nodes instead of copying remote nodes with their data into the local host graph. They are explicitly inserted into the runtime graphs storing the location of their according remote nodes. As reference nodes are only helper structures for accessing remote objects, they are implicitly managed by the runtime level and

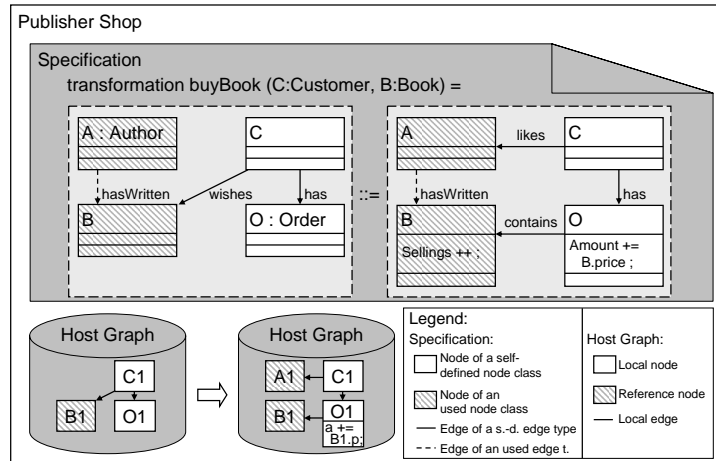


Figure 2: Graph transformation BuyBook

are not regarded in the graph transformations. Their usage supports the realization of self-defined edges incident to remote nodes. In contrast to reference nodes, we do not store reference edges.

A remote node (and thus an appropriate reference node in the local host graph) can be created within a distributed graph transformation by using the according node only in the transformation’s right-hand side. This concept corresponds to creating a local node of a self-defined type. Analogously, the deletion of a remote node (and thus of the corresponding reference node) is triggered by using the according node only on the left-hand side of the transformation.

Furthermore, a reference node is automatically created in the local host graph if an appropriate node is used in the left-hand side of the transformation, and no adequate reference node is locally available, although an according remote node exists. In this case, the remote application defining the node’s type is searched for an adequate node, and an according reference node is created. Additionally, all reference nodes are automatically deleted if the actual remote node is deleted. Operations specified on remote objects, like the deletion of a node or an attribute assignment, are propagated transparently to the remote applications³. To ensure the consistency of the host graphs by remote graph transformations, we will introduce a *rule engine* (cf. section 5).

Figure 2 shows an example of the distributed graph transformation buyBook of the Publisher Shop. In this transformation, a Customer buys a Book (given as input parameters), which has been on the Customer’s wish list. This leads to the deletion of the wishes-edge as it is no longer needed. Furthermore, the two edges likes and contains are created. The attribute Sellings counting the number of sales of B is incremented. Below the transformation, two example host graphs of the Publisher Shop are depicted showing the host graph before and after the execution of the transformation. According to the transformation edges are created and deleted and a new reference node A1 for the Author A is inserted. In the Publisher’s host graph only the attribute Sellings of node B1 is modified, which is not shown in figure 2.

³For propagating remote operations and the management of references, we develop an appropriate plug-in for the database DRAGOS [Böh04], which is used by PROGRES (and soon by Fujaba) prototypes for storing the graphs.

3 Search Trees

After modeling a software system with a GTS, appropriate code has to be generated. As the matching of the left-hand side of a graph transformation has an exponential worst-case complexity, PROGRES and Fujaba use *search trees* [Zün96] for their code generation. In the course of this section, we will describe these search trees using the code generation plug-in CodeGen2 [GSR05] of Fujaba as an example. As we present the current mechanism, we ignore the remote nature of operations concerning a coupled application.

3.1 General Structure of Search Trees

To generate executable and efficient code for graph transformations, every transformation is translated into a search tree. A search tree is a tree, having *operations* as nodes and *precondition*-relations between an operation and its child operations. The search tree contains all operations, which have to be performed for the modeled graph transformation in the right execution order. In this paper, we use a more general notion of a search tree, since not only search operations are used. CodeGen2 uses over 20 different operations, covering constraint checks and operations for creating and deleting nodes.

In general, an operation describes one runtime semantics of one or more modeling elements within a graph transformation. There is no *1-to-1*-mapping between operations and modeling elements. One modeling element may result in several different operations in the search tree, because its runtime semantics is precisely defined in the code generation phase. On the other hand, one operation may also cover several modeling elements, as the semantics of one modeling element may be determined only in combination with other elements.

For translating the graph transformation into a search tree, every node of the transformation's left-hand side (*graph pattern*) has to be identified in the host graph. As a first step, a search tree covering all bound nodes⁴ is created. The second step regards all unbound nodes calculating a spanning forest for the graph pattern. For every bound node, the forest contains a tree having the bound node as root node. Every unbound node has to be searched in the host graph by traversing one edge of the pattern incident to a bound node. The found unbound node together with its edge is inserted in the tree of the corresponding bound node. After building the spanning forest for all pattern nodes, for all unbound nodes of the spanning forest a *search operation* representing the unbound node and the traversed edge is inserted into the search tree. This operation is placed as child of the operation, which binds the edge's source node.

Note that Fujaba does not allow isolated unbound nodes in graph patterns, because it uses the heap of the JAVA runtime environment instead of a graph database. Other graph transformation systems like PROGRES support this feature by querying the underlying database for nodes of a specific type. For the sake of simplicity, we do not regard the search of isolated nodes and patterns although this is covered by our approach using DRAGOS [Böh04].

After inserting the search operations into the search tree, the remaining modeling elements of the graph transformation (like edges and attribute assertions) have to be inserted. The position of such an operation

⁴A *bound node* is a node, which already has been uniquely identified in the host graph, for example nodes given as input parameters of a transformation. In contrast, *unbound nodes* are nodes, which have to be searched in the host graph. If an appropriate node for the unbound node is found, this node becomes a bound node.

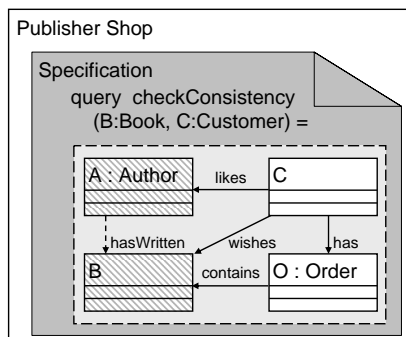


Figure 3: Graph query checkConsistency

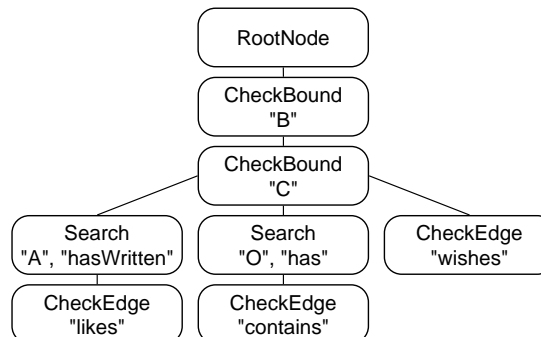


Figure 4: Search tree for query checkConsistency

is determined by the following overall invariant of the tree: All preconditions of an operation must be fulfilled by preceding operations. To satisfy this invariant, the search tree may have to be reorganized according to the invariant. The complete algorithm for the search tree generation is shown in [GSR05].

To illustrate the generation of search trees, we introduce the query `checkConsistency` depicted in figure 3. This query may be used for consistency checks, e.g. to test if the `wishes`-edge has been deleted when the Book has been bought by the Customer and is thus part of an Order.

Figure 4 shows a possible search tree for `checkConsistency`. The generation of the tree starts with two `CheckBound`-operations for parameters B and C, which are inserted as children of the `RootNode`. A `CheckBound`-operation checks the validity of bound nodes, i.e. that they are not equal to `null`.

To determine all search operations needed to cover all unbound pattern nodes, a spanning forest has to be computed, consisting of two trees: One tree with B as root node and the other tree with C as root node. For our example, we assume that the spanning forest is given by all nodes with the `hasWritten`- and the `has`-edge. According to this assumption, the search operations for A and O using the `hasWritten`- resp. `has`-edge are inserted into the search tree.

To cover the remaining edges, namely `contains`, `wishes` and `likes`, `CheckEdge`-operations have to be inserted. A `CheckEdge`-operation must have at least those operations in their parent hierarchy, which cover the source and the target node of the edge. To fulfill this invariant, the tree may have to be reorganized before the `CheckEdge`-operations can be inserted.

In our example, the `CheckBound`-operation of node C (`CheckC`) is moved below the `CheckBound`-operation of node B (`CheckB`) and all children of `CheckB` are moved below `CheckC`⁵. After the reorganization, the `CheckEdge`-operation of the edge `wishes` can be inserted into the search tree considering the invariant. All other `CheckEdge`-operations can be inserted without any further reorganization.

3.2 Cost Model

A graph transformation may have several different valid search trees, because there are several ways to cover all modeling elements by operations. To evaluate the different search trees, the code generation uses a *cost model* presented in the following.

⁵These reorganizations are possible, because siblings in the search tree are independent of each other.

index	source	edge	target	cost
1	C	has	O	25
2	B	contains	O	25
3	B	hasWritten	A	1
4	C	likes	A	25

Table 1: Cost table for search operations of query checkConsistency

Search operations may have different runtime costs, which can be estimated by exploiting the cardinalities of the graph schema, for example. Search operations, which use a *to-n*-edge for matching an unbound node, are very costly compared to those using a *to-1*-edge. Considering a *to-1*-edge, the right candidate node can be directly determined, whereas a search operation using a *to-n*-edge leads to n possible candidate nodes.

However, the costs for search operations can only be estimated, since the exact cardinality for a *to-n*-edge in the host graph is not known during the code generation phase. Therefore, CodeGen2 and the PROGRES code generation use default values for the expected costs. [VVF05] presents an approach for adapting the edge cardinalities by analysis of sample instance graphs leading to very precise cost estimations. So far, this approach is not considered by CodeGen2 and thus not regarded in this paper.

Even by considering only the static graph schema for cost estimations, the runtime behavior can be greatly improved. Before generating the search tree, the costs for each operation of the graph transformation are estimated. Afterwards, a minimum spanning forest of the graph pattern is computed, leading to a preliminary search tree. This preliminary tree is incrementally extended by operations until all modeling elements are covered choosing the cheapest operation in each step.

In our example checkConsistency (cf. figures 3, 4), there are four possible search operations (cf. table 1), which have to be considered for the search tree generation. According to the target cardinalities of the edges, the costs for the search operations are either 1 or n . Note that this cost model is highly simplified and n has to be a real value greater than 1. CodeGen2 contains different *cost strategies* for assigning concrete cost values. The default cost strategy for search operations based on *to-n*-edges assumes a maximal cardinality of 50. As result, it computes a cost of 25, since this is the expected value having 1 as lower and 50 as upper cardinality.

To minimize the costs of the search tree and thus get the best runtime efficiency of the generated code, the algorithm chooses the cheapest search operations until all unbound nodes are covered. In our example, search operation 3 is chosen first due to the smallest costs. Since node O can be covered by two operations with the same costs, the algorithm chooses non-deterministically between the search operation 1 and 2. Given that operation 1 is chosen, the search tree depicted in figure 4 is computed.

4 Search Trees for Distributed Graph Transformations

In section 2 we presented our syntax for modeling distributed graph transformations. To execute such a transformation, it has to be translated into executable code. Therefore, we extend the CodeGen2-

index	source	edge	target	cost
1	C	has	O	25
2	B	contains	O	25
3	B	hasWritten	A	51
4	C	likes	A	25

Table 2: Modified cost table for search operations of query checkConsistency

approach presented in section 3, which uses search trees to generate efficient code.

In the following, only the code generation for queries is shown, because queries have the greatest impact on the runtime efficiency, and transforming operations can only be executed if the search of the pattern has been successful. A distributed query requires *remote operations*, like remote check operations and remote search operations. Each remote operation affects exactly one application which is called *home application*. After searching the distributed pattern, the transforming operations are performed, which require no modifications of CodeGen2.

4.1 Remote Offset

Remote operations have additional costs compared to their local correspondents, which results from the communication delay between the participating applications. The additional cost is called *remote offset* and depends on the network infrastructure. The remote offset is configurable for each pair of applications, which participate in the communication. Its calculation is integrated into CodeGen2, using a different cost strategy for remote operations.

Revisit the query consistencyCheck introduced in section 3. Every search tree that is computed without regarding a remote offset uses the hasWritten-edge as first search operation. Thus, a remote operation is performed even it is not necessary, since A can also be found via the local edge likes.

In many cases, a remote operation is more costly than every local operation. Considering this in our cost model, we introduce a remote offset for every remote operation, which is larger than the maximal cardinality of any local edge. Thus, assuming a cost of 25 for a search operation using a *to-n*-edge, we set the remote offset to 50. The modified costs are depicted in table 2.

Figure 5 shows the search tree for the query checkConsistency, which is computed according to the modified cost table (cf. table 2). Search operation 4 is chosen first instead of search operation 3 due to the remote offset. The only remote operation in the modified search tree is the CheckEdge-operation of likes, which is depicted striped in the search tree.

4.2 Boundary Nodes

In many cases, there exists a number of operations, which have the same home application and are dependent on each other. Instead of formulating queries for every operation, operations may be processed *en bloc* leading to one single query. Thus, the number of communications can be heavily reduced.

The crucial question is how to cut the graph pattern into separate sub-patterns (*remote blocks*), such that

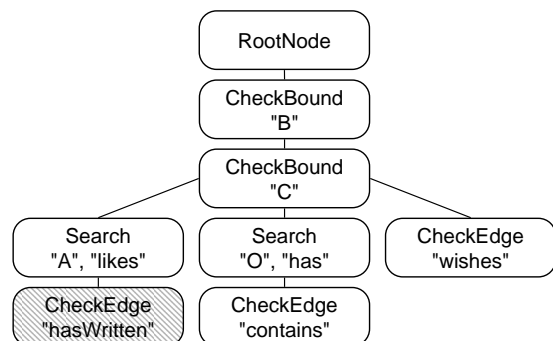


Figure 5: Modified search tree for query checkConsistency

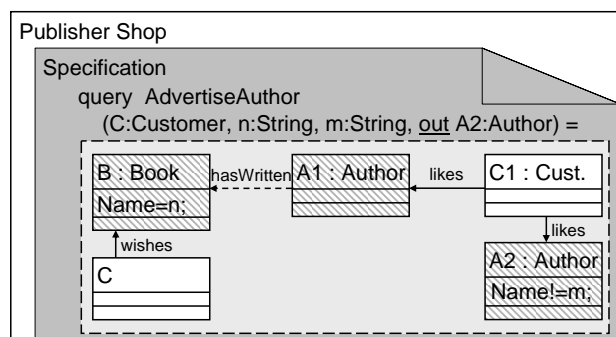


Figure 6: Graph query AdvertiseAuthor

every sub-pattern concerns solely one home application. To illustrate this, the graph query `AdvertiseAuthor` is depicted in figure 6. The query can be used to advertise an `Author` to a given `Customer C`, using the `likes`-relation of another `Customer` with the same interests.

For searching an unbound node like `C1`, the search is *parameterized* by the edge’s source node needed for the search operation, i.e. `A1`. Due to this parametrization, `C1` can only be searched if `B` and `A1` are found. Furthermore `A2` can be searched only if `B`, `A1` and `C1` are found. Consequently, it is not always possible to create only one remote block per affected home application.

To define the boundary of a block, we introduce the notion of *boundary nodes*. A boundary node is a node which needs at least one remote operation for its identification, and is incident to a local edge. In figure 6, `B` and `A2` are boundary nodes due to the remote attribute check. Without the attribute check, the nodes `B` and `A2` can be covered without remote operations. A block is defined as the largest connected graph pattern, which is limited by its boundary nodes and is covered only by remote operations concerning the same home application. Revisit the query `AdvertiseAuthor`, there exists two blocks: Block 1 contains the attribute condition of node `B`, the node `A1` and the `hasWritten`-edge. Block 2 contains the attribute check of `A2`. These blocks are sent to their home application separately preserving this order and are parameterized by `B` resp. `A2`.

4.3 Realization of Remote Blocks

In the following, we describe how the creation of remote blocks can be integrated into the search tree approach of `CodeGen2` (cf. section 3) considering the remote offset. We develop an algorithm for creating remote blocks using the generated search tree as input. This algorithm uses the following two propositions: Let R be the set of all remote operations concerning the same home application forming one remote block. Then

1. every element of R is contained in a common subtree with root RN^6
2. every element of R is reachable by RN by only traversing other elements of R .

We do not provide a formal proof of these propositions, but want to discuss their plausibility.

⁶The following abbreviations are used: RN for root node, and RO for remote operations having RN as direct predecessor.

To 1: Since a remote block is connected, all elements of R are somehow dependent on each other. Due to the tree structure of the search tree, these operations must have a common root operation. This root operation must cover a boundary node. Then, the elements of R are contained in the subtrees having a child of RN as their root.

To 2: This is valid, because all remote operations inside the remote block are dependent on each other.

In the following, we present the algorithm depicted in listing 1, which calculates the remote blocks. It rearranges all remote operations in a chain within the search tree, from which the blocks for the remote applications and their parameterization can be directly derived. The algorithm needs two sets, which are defined as follows: Let RO be the set of remote operations, which have RN as direct predecessor and the same home application as RN . If RO is empty, RO is recomputed containing all remote operations based on the home application of one direct child of RN . Let LO be a set of subtrees within the search tree, which aggregates all direct subtrees of RN whose root node is not contained in RO . The algorithm `linearize` (cf. listing 1) starts with the root node of the search tree as RN and $LO = \emptyset$.

```

linearize(Node RN, TreeSet LO){
  RO:=all subtrees of RN regarding one HA
  LO:=LO+all remaining subtrees of RN
  if (RO!=∅){
    for (i=2 to |RO|) move ROi below RO1
    RN:=root node of RO1
    linearize(RN, LO)
  } else{
    for (each l ∈ LO){
      insert l below RN
      RN:=root node of l
      linearize(RN, ∅)
    }
  }
}

```

Listing 1: Pseudocode of `linearize`

If $RO \neq \emptyset$, all subtrees contained in RO except of RO_1 are moved below the root node of RO_1 . Then the algorithm is called with the root node of RO_1 as new parameter RN and LO . If $RO = \emptyset$, all subtrees contained in LO are inserted below RN and the algorithm is called recursively for all of their root nodes as RN and $LO := \emptyset$.

Due to proposition 1, every subtree containing a remote operation of block RB is considered. According to proposition 2, every remote operation is considered by the algorithm. All needed rearrangements are possible, because siblings – and hence disjoint subtrees – in the search tree are independent of each other. The algorithm terminates, as every node of the search tree is once used as parameter RN .

The algorithm `linearize` does not improve the worst-case complexity of a distributed pattern matching. A worst-case scenario consists of many remote blocks each covering only one remote operation. On the other hand, a best-case scenario consists of a chain of m remote objects, which are connected by $to-n$ -edges, and the chain is connected by a $to-1$ -edge with a local object. Without optimization, at most $n^{(m-1)}$ communication steps are necessary to determine the match of the remote block, since $m - 1$ links have to be traversed and there exist n possible candidates for each node in the pattern. With the block creation, the whole query can be done within one communication step.

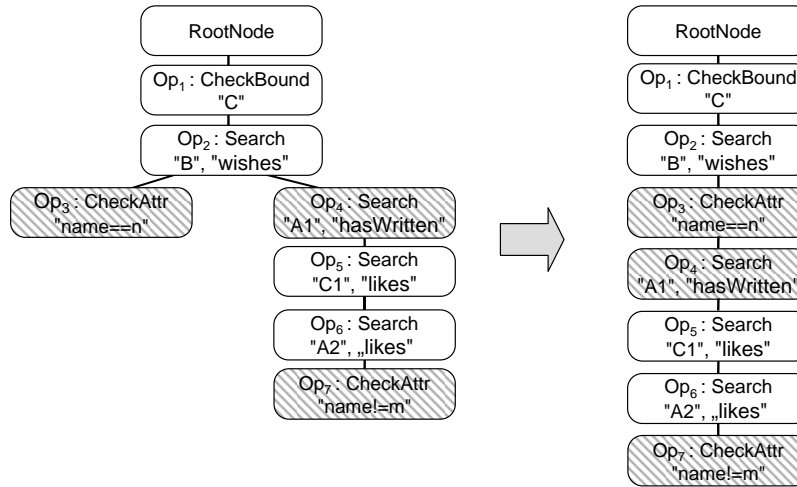


Figure 7: Search tree for query AdvertiseAuthor

We will now illustrate the behavior of `linearize` considering the query `blockExample`, whose search tree is shown in figure 7. The algorithm starts with $RN = \text{RootNode}$ and $LO = \emptyset$. Since all subtrees of RN start with a local operation, $RO = \emptyset$ and $LO = \{ST(\text{Op}_1)\}$ ⁷. Due to $RO = \emptyset$, `linearize` is called with $RN = \{\text{Op}_1\}$ and $LO = \emptyset$. Now, RO and LO are recomputed, resulting in $RO = \emptyset$ and $LO = \{ST(\text{Op}_2)\}$. `linearize` is called again with $RN = \text{Op}_2$ and $LO = \emptyset$. RO is recomputed as $\{ST(\text{Op}_3), ST(\text{Op}_4)\}$ and LO remains \emptyset . The search tree is rearranged, i.e. $ST(\text{Op}_4)$ is moved below Op_3 . Afterwards, `linearize` is called with $RN = \text{Op}_3$ and $LO = \emptyset$. The algorithm has already computed the desired chain, and thus performs no further rearrangements of the search tree. Now it contains two chains of remote operations representing two remote blocks (cf. section 4.2).

5 Conclusion

In this paper, we presented our approach for generating efficient code for distributed graph transformations. Our approach is based on existing search tree algorithms used by the code generations of `PROGRES` and `Fujaba`. The usage of search trees is advantageous for reducing the average complexity of searching graph patterns, as this search has an exponential worst-case complexity. To reduce also the communication costs arising in distributed systems, we extended the existing code generations: We have integrated a remote offset for remote operations in the cost table giving priority to the execution of local operations. Additionally, we extract sub-patterns, which are sent to the appropriate remote applications. Thus, the remote applications are not queried for every single element of the graph pattern, decreasing the communication overhead.

As the coupled specifications export only schema parts, distributed graph transformations may be modeled, which violate certain local constraints. Thus, the execution of distributed graph transformations

⁷ $ST(\text{Op}_1)$ stands for subtree with root node Op_1 .

may lead to inconsistent host graphs within the distributed system. Therefore, we will introduce a *rule engine* as next step, which will facilitate the execution of *repair actions* [Win00]. Furthermore, this engine will be able to prevent the creation or deletion of a certain node as described in [HEET99].

References

- [Böh04] B. Böhlen. Specific graph models and their mappings to a common model. In J. L. Pfaltz, M. Nagl, and B. Böhlen, editors, *Applications of Graph Transformations with Industrial Relevance, 2nd Intern. Workshop, AGTIVE 2003*, volume 3062 of *Lect. Notes in Comp. Sci.*, pages 45–60. Springer-Verlag, 2004.
- [FNTZ00] T. Fischer, J. Niere, L. Torunski, and A. Zündorf. Story diagrams: A new graph rewrite language based on the Unified Modelling Language and Java. In H. Ehrig, G. Engels, H.-J. Kreowski, and G. Rozenberg, editors, *Theory and Application of Graph Transformations, 6th Intern. Workshop, TAGT'98*, volume 1764 of *Lect. Notes in Comp. Sci.*, pages 296–309. Springer-Verlag, 2000.
- [GSR05] L. Geiger, C. Schneider, and C. Reckord. Template- and model-based code generation for MDA-tools. In H. Giese and A. Zündorf, editors, *Proc. of the Fujaba Days 2005*, volume 259 of *Technical Report*. University of Paderborn, Germany, 2005.
- [HEET99] R. Heckel, H. Ehrig, G. Engels, and G. Taentzer. A view-based approach to system modeling based on open graph transformation systems. In H. Ehrig, G. Engels, H.-J. Kreowski, and G. Rozenberg, editors, *Handbook on Graph Grammars and Computing by Graph Transformation: Applications, Languages, and Tools*, volume 2, pages 639–668. World Scientific, Singapore, 1999.
- [RS06] U. Ranger and E. Schultchen. Specifying distributed graph transformation systems. 2006. Submitted to the 3rd International Workshop on Graph-Based Tools (GraBaTs 2006).
- [Sch91] A. Schürr. *Operationales Spezifizieren mit programmierten Graphersetzungssystemen*. PhD-Thesis, RWTH Aachen, 1991.
- [VVF05] G. Varró, D. Varró, and K. Friedl. Adaptive graph pattern matching for model transformations using model-sensitive search plans. In G. Karsai and G. Taentzer, editors, *Proc. International Workshop on Graph and Model Transformation*, volume 125 of *Electr. Notes in Theor. Comp. Sci.* Elsevier Science, 2005.
- [Win00] A. J. Winter. *Visuelles Programmieren mit Graphtransformationen*. PhD-Thesis, RWTH Aachen, 2000.
- [Zün96] A. Zündorf. Graph pattern matching in PROGRES. In J. E. Cuny, H. Ehrig, G. Engels, and G. Rozenberg, editors, *Graph Grammars and Their Application to Computer Science, 5th Intern. Workshop*, volume 1073 of *Lect. Notes in Comp. Sci.*, pages 454–468. Springer-Verlag, 1996.