

Decomposition of Distributed Nonmonotonic Multi-Context Systems*

Seif El-Din Bairakdar, Minh Dao-Tran, Thomas Eiter, Michael Fink and Thomas Krennwallner

Institute of Information Systems, Vienna University of Technology
Favoritenstrasse 9–11, A-1040 Vienna, Austria
{bairakdar,dao,eiter,fink,tkren}@kr.tuwien.ac.at

Abstract

Multi-Context Systems (MCS) are formalisms that enable the interlinkage of single knowledge bases, called contexts, via bridge rules. Recently, the evaluation of heterogeneous, non-monotonic MCS was considered in Dao-Tran et al. (2010), where a fully distributed algorithm was described. In this paper, we continue this line of work and present a decomposition technique for MCS which analyzes the topology of an MCS. It applies pruning techniques to get economically small representations of context dependencies. Orthogonal to this, we characterize minimal interfaces for information exchange between contexts, such that data transmissions can be minimized. We then present a novel evaluation algorithm that operates on a query plan which is compiled with topology pruning and interface minimization. The effectiveness of the optimization techniques is demonstrated by a prototype implementation, which uses an off-the-shelf SAT solver and shows encouraging experimental results.

Introduction

In the last years, there has been increasing interest in systems comprising multiple knowledge bases. The rise of distributed systems and the World Wide Web fostered this development, and to date, several formalisms are available that accommodate multiple, possibly distributed knowledge bases. One formalism are Multi-Context Systems (MCS), which consist of several theories (the contexts) that are interlinked with bridge rules which allow to add knowledge to a context depending on knowledge in other contexts. For instance, the bridge rule $a \leftarrow (2 : b)$ of a context C_1 means that C_1 should conclude a if context C_2 believes b . MCS have applications in various areas, such as argumentation, data integration, or multi-agent systems. There, each context may model the beliefs of an agent while the bridge rules model an agent's perception of the environment, i.e., other contexts.

Among the various MCS proposals (e.g., McCarthy, 1993, Giunchiglia and Serafini, 1994, and Ghidini and Giunchiglia, 2001), the general MCS framework of Brewka and Eiter (2007) is of special interest, as it generalizes previous approaches in contextual reasoning and allows for *heterogeneous and nonmonotonic* MCS, i.e., with different, possibly

nonmonotonic logics in its contexts (thus furthering heterogeneity), and bridge rules may use default negation (to deal, e.g., with incomplete information). Hence, nonmonotonic MCS interlinking monotonic context logics are possible. This MCS framework can conveniently capture the following scenario, which we use as a running example.

Example 1 A group of four scientists, Ms. 1, Mr. 2, Mr. 3, and Ms. 4, just finished their conference visit and are now arranging a trip back home. They can choose between going by train or by car (which is usually slower than the train); and if they use the train, they should bring along some food. Moreover, Mr. 3 and Ms. 4 have additional information from home that might affect their decision.

Mr. 3 has a daughter, Ms. 6. He is fine with either transportation option, but if Ms. 6 is sick then he wants to use the fastest vehicle to get home. Ms. 4 just got married, and her husband, Mr. 5, wants her to come back as soon as possible. He urges her to try to come home even sooner, while Ms. 4 tries to yield to her husband's plea.

If they go by train, Mr. 3 is responsible for buying provisions. He might choose either sandwiches or chocolate peanuts. The options for beverages are coke or juice. Mr. 2 is a modest person as long as he gets home. He agrees to any choice that Mr. 3 and Ms. 4 select for vehicle but he dislikes coke. Ms. 1 is the leader of the group and prefers to go by car, but if Mr. 2 and 3 want to go by train then she would not object. The only problem is that Ms. 1 is allergic to peanuts.

Mr. 3 and Ms. 4 do not want to bother the others with their personal situation and communicate just their preferences, which is sufficient for reaching an agreement. Ultimately, Ms. 1 decides which option to take based on the information she gets from Mr. 2 and Mr. 3. \square

Similar scenarios have already been investigated in the realm of multi-agent systems (see, e.g., Buccafurri and Caminiti (2008) on social answer set programming). We do not aim at introducing a new semantics for such scenarios; our example is meant to be a plain showcase application of MCS. We stress that MCS have potential as a host for KR formalisms, just like answer set programs have; however, in this paper we concentrate on efficient MCS evaluation.

Dao-Tran et al. (2010) introduced a distributed algorithm, called DMCS, to compute the semantics of an MCS, which is given in terms of equilibria. Roughly, an equilibrium is a

*This research has been supported by the Austrian Science Fund (FWF) project P20841 and by the Vienna Science and Technology Fund (WWTF) project ICT 08-020.

collection of local models (belief sets) for the individual contexts that is compatible with the bridge rules. The principle of the algorithm is, starting from context C_k (the root), that models will be processed at each context. The belief import relationship of the MCS given by the bridge rules is used to navigate the system; models returned from invoked neighbors are combined with the local beliefs and passed back to the invoking contexts. DMCS uses a parameter for projecting models to relevant variables to reduce data payload.

Experiments for an instantiation of DMCS with answer set programming contexts revealed some scalability issues which can be tracked down to the following problems:

1. contexts are unaware of context dependencies in the system beyond their neighbors, and thus treat each neighbor in a generic way. Specifically, cyclic dependencies remain undetected until a context, seeing the invocation chain, requests models from a context in the chain. Furthermore, a context C_k does not know whether a neighbor C_i already requests models from another neighbor C_j which then would be passed to C_k ; hence, C_k makes possibly a superfluous request to C_j .
2. a context C_i returns the combination of its local models with the models received from all neighboring contexts. As contexts may have multiple models, the number of models can become huge as the size of the system respectively neighbors increases. In fact, this is one of the main performance obstacles.

In this work, we address the issue of optimization; there is an urgent need for this in order to increase the scalability of distributed MCS evaluation. Resorting to methods from graph theory, we aim at decomposing, pruning, and improved cycle breaking for dependencies in multi-context systems.

Focusing on 1, we describe a decomposition method using biconnected components of inter-context dependencies. Based on this we can break cycles and prune acyclic parts before evaluating the system and create an acyclic query plan. To address 2, we foster a partial view of the system, which is often sufficient to reach a satisfactory answer. In Example 1, for instance, we could mask out the beliefs of Mr. 5 and Ms. 6 to compute a partial equilibrium within the scientist group. In this way, we have the possibility of a trade-off between partial information and performance. Concretely, we define a set of variables for each import dependency in the system that is used to project the models of each context to the bare minimum for performing a meaningful computation. In this manner, we can omit needless information and circumvent excessive model combinations.

Based on these ideas, we have designed a new evaluation algorithm DMCSOPT, which intertwines decomposition and pruning with variable projection. For evaluation, we adapted our prototype implementation of DMCS and ran some experiments. The results show a major improvement compared to DMCS; here we can handle systems with up to 700 contexts. This demonstrates that our optimization techniques are effective and bring MCS closer to applications.

Preliminaries

We recall some basic notions of heterogeneous nonmonotonic multi-context systems (Brewka and Eiter 2007).

A *logic* is, viewed abstractly, a tuple $L = (\mathbf{KB}_L, \mathbf{BS}_L, \mathbf{ACC}_L)$, where

- \mathbf{KB}_L is a set of well-formed knowledge bases, each being a set (of formulas),
- \mathbf{BS}_L is a set of possible belief sets, each being a set (of formulas), and
- $\mathbf{ACC}_L: \mathbf{KB}_L \rightarrow 2^{\mathbf{BS}_L}$ assigns each $kb \in \mathbf{KB}_L$ a set of acceptable belief sets.

This covers many (non-)monotonic KR formalisms like description logics, default logic, answer set programs, etc.

For example, a (propositional) *ASP logic* L may be such that \mathbf{KB}_L is the set of answer set programs over a (propositional) alphabet \mathcal{A} , $\mathbf{BS}_L = 2^{\mathcal{A}}$ contains all subsets of atoms, and \mathbf{ACC}_L assigns each $kb \in \mathbf{KB}_L$ the set of all its answer sets (see Gelfond and Lifschitz, 1991, for details).

Definition 1 A multi-context system (MCS) $M = (C_1, \dots, C_n)$ consists of contexts $C_i = (L_i, kb_i, br_i)$, $1 \leq i \leq n$, where $L_i = (\mathbf{KB}_i, \mathbf{BS}_i, \mathbf{ACC}_i)$ is a logic, $kb_i \in \mathbf{KB}_i$ is a knowledge base, and br_i is a set of L_i -bridge rules of the form

$$s \leftarrow (c_1 : p_1), \dots, (c_j : p_j), \text{not } (c_{j+1} : p_{j+1}), \dots, \text{not } (c_m : p_m) \quad (1)$$

where $1 \leq c_k \leq n$, p_k is an element of some belief set of L_{c_k} , $1 \leq k \leq m$, and $kb \cup \{s\} \in \mathbf{KB}_i$ for each $kb \in \mathbf{KB}_i$.

Informally, bridge rules allow to modify the knowledge base by adding s , depending on the beliefs in other contexts.

The semantics of an MCS M is defined in terms of particular *belief states*, which are sequences $S = (S_1, \dots, S_n)$ of belief sets $S_i \in \mathbf{BS}_i$. Intuitively, S_i should be a belief set of the knowledge base kb_i ; however, also the bridge rules br_i must be respected. To this end, kb_i is augmented with the conclusions of all $r \in br_i$ that are applicable.

Formally, r of form (1) is *applicable in* S , if $p_i \in S_{c_i}$, for $1 \leq i \leq j$, and $p_k \notin S_{c_k}$, for $j+1 \leq k \leq m$. Let $app(R, S)$ denote the set of all bridge rules $r \in R$ that are applicable in S . Furthermore, $head(r)$ denotes the part s , and $B(r) = \{(c_k : p_k) \mid 1 \leq k \leq m\}$, for any r of form (1).

Definition 2 A *belief state* $S = (S_1, \dots, S_n)$ of a multi-context system M is an equilibrium iff for all $1 \leq i \leq n$, $S_i \in \mathbf{ACC}_i(kb_i \cup \{head(r) \mid r \in app(br_i, S)\})$.

In the rest of this paper, we assume that contexts C_i have finite belief sets S_i that are represented by truth assignments $v_{S_i}: \Sigma_i \rightarrow \{0, 1\}$ to a finite set Σ_i of propositional atoms such that $p \in S_i$ iff $v_{S_i}(p) = 1$ (as in Brewka and Eiter, 2007, such S_i may serve as kernels that correspond 1-1 to infinite belief sets). Furthermore, we assume that the Σ_i are pairwise disjoint and that $\Sigma = \bigcup_i \Sigma_i$.

Example 2 The scenario in Example 1 can be encoded as an MCS $M = (C_1, \dots, C_6)$, where all L_i are ASP logics and $kb_1 = \{car_1 \leftarrow \text{not } train_1; \perp \leftarrow \text{peanuts}_1\}$ and $br_1 = \left\{ \begin{array}{l} train_1 \leftarrow (2 : train_2), (3 : train_3) \\ peanuts_1 \leftarrow (3 : \text{chocolate-peanuts}_3) \end{array} \right\}$;

$kb_2 = \{\perp \leftarrow \text{not } car_2, \text{not } train_2\}$ and $br_2 = \left\{ \begin{array}{l} car_2 \leftarrow (3 : car_3), (4 : car_4) \\ train_2 \leftarrow (3 : train_3), (4 : train_4), \text{not } (3 : coke_3) \end{array} \right\}$;

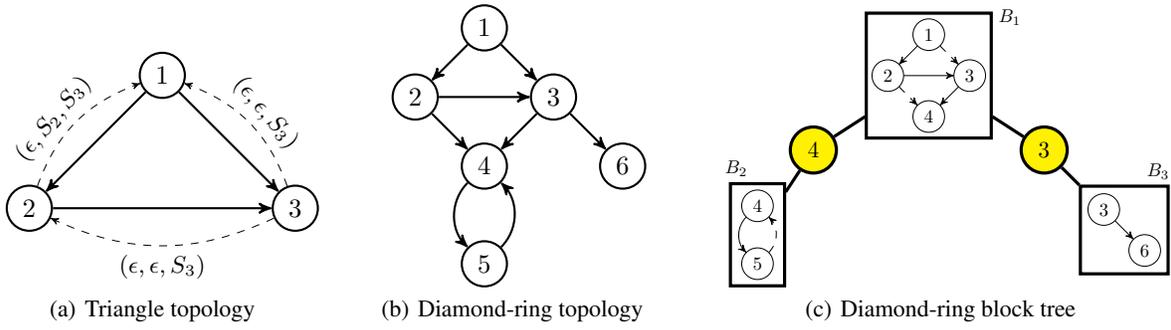


Figure 1: Topologies and Decomposition of Scientist Group Example

$$kb_3 = \left\{ \begin{array}{l} car_3 \vee train_3 \leftarrow; \quad train_3 \leftarrow urgent_3 \\ sandwiches_3 \vee chocolate_peanuts_3 \leftarrow train_3 \\ coke_3 \vee juice_3 \leftarrow train_3 \end{array} \right\}$$

and $br_3 = \{urgent_3 \leftarrow (6 : sick_6); train_3 \leftarrow (4 : train_4)\}$;

$$kb_4 = \{car_4 \vee train_4 \leftarrow\} \text{ and}$$

$$br_4 = \{train_4 \leftarrow (5 : want_sooner_5)\};$$

$$kb_5 = \{want_sooner_5 \leftarrow soon_5\} \text{ and}$$

$$br_5 = \{soon_5 \leftarrow (4 : train_4)\};$$

$$kb_6 = \{sick_6 \vee fit_6 \leftarrow\} \text{ and } br_6 = \emptyset.$$

The dependencies between contexts in M are shown in Figure 1(b). M has three equilibria, namely:

$$S = (\{train_1\}, \{train_2\}, \{train_3, urgent_3, juice_3, sandwiches_3\}, \{train_4\}, \{soon_5, want_sooner_5\}, \{sick_6\});$$

$$T = (\{train_1\}, \{train_2\}, \{train_3, juice_3, sandwiches_3\}, \{train_4\}, \{soon_5, want_sooner_5\}, \{fit_6\});$$

$$U = (\{car_1\}, \{car_2\}, \{car_3\}, \{car_4\}, \emptyset, \{fit_6\}).$$

Partial Equilibria. We recall partial equilibria, which informally are equilibria of a sub-MCS generated by a context C_k .

Definition 3 (Import Closure) Let $M = (C_1, \dots, C_n)$ be an MCS. The import neighborhood of a context C_k is the set

$$In(k) = \{c_i \mid (c_i : p_i) \in B(r), r \in br_k\}.$$

Moreover, the import closure $IC(k)$ of C_k is the smallest set S such that (i) $k \in S$ and (ii) for all $i \in S$, $In(k) \subseteq S$.

Based on the import closure, we then define:

Definition 4 (Partial Belief States and Equilibria)

Let $M = (C_1, \dots, C_n)$ be an MCS, and let $\epsilon \notin \bigcup_{i=1}^n \mathbf{BS}_i$. A partial belief state of M is a sequence $S = (S_1, \dots, S_n)$, such that $S_i \in \mathbf{BS}_i \cup \{\epsilon\}$, for $1 \leq i \leq n$.

A partial belief state $S = (S_1, \dots, S_n)$ of M is a partial equilibrium of M w.r.t. a context C_k iff $i \in IC(k)$ implies $S_i \in \mathbf{ACC}_i(kb_i \cup \{head(r) \mid r \in app(br_i, S)\})$, and if $i \notin IC(k)$, then $S_i = \epsilon$, for all $1 \leq i \leq n$.

Example 3 In our example, $(\epsilon, \epsilon, \{car_3, coke_3, chocolate_peanuts_3\}, \{train_4\}, \{soon_5, want_sooner_5\}, \{fit_6\})$ is a partial equilibrium w.r.t. context C_3 . \square

For combining partial belief states $S = (S_1, \dots, S_n)$ and $T = (T_1, \dots, T_n)$, we define their *join* $S \bowtie T$ as the partial belief state (U_1, \dots, U_n) with (i) $U_i = S_i$, if $T_i = \epsilon \vee S_i = T_i$, and (ii) $U_i = T_i$, if $T_i \neq \epsilon \wedge S_i = \epsilon$, for all $1 \leq i \leq n$. Note that $S \bowtie T$ is void, if some S_i, T_i are from \mathbf{BS}_i but different.

The *join* of two sets \mathcal{S} and \mathcal{T} of partial belief states is then naturally defined as $\mathcal{S} \bowtie \mathcal{T} = \{S \bowtie T \mid S \in \mathcal{S}, T \in \mathcal{T}\}$.

Given a (partial) belief state S and set $V \subseteq \Sigma$ of variables, the *restriction* of S to V , denoted $S|_V$, is given by $S = (S_1|_V, \dots, S_n|_V)$, where $S_i|_V = S_i \cap V$ if $S_i \neq \epsilon$, and $\epsilon|_V = \epsilon$; for a set of (partial) belief states \mathcal{S} , we let $\mathcal{S}|_V = \{S|_V \mid S \in \mathcal{S}\}$.

Definition 5 The import interface of context C_k in M is $V(k) = \{p_i \mid (c_i : p_i) \in B(r), r \in br_k\}$, and its recursive import interface is $V^*(k) = V(k) \cup \{p \in V(j) \mid j \in IC(k)\}$.

There are two extremal cases: 1. $V = V^*(k)$. Then, partial equilibria projected to V can be basically used for consistency-checking on the import closure of C_k . 2. $V = \Sigma$. Here, the projection to V yields partial equilibria w.r.t. C_k . By providing a fixed interface V such that $V^*(k) \subseteq V \subseteq \Sigma$, problem-specific knowledge (e.g. query variables) and infrastructure information can be exploited to focus computations to relevant projections.

Example 4 (cont'd) The recursive import interface of C_1 in M is $V^*(1) = \{car_3, car_4, chocolate_peanuts_3, coke_3, sick_6, train_2, train_3, train_4, want_sooner_5\}$. \square

Decomposition of Nonmonotonic MCS

Reconsider our running example, with all contexts C_i and atoms referring to them removed, for $i > 3$. Then, C_1 has bridge rules with atoms of form $(2 : p_2)$ and $(3 : p_3)$ in the body, and C_2 with atoms $(3 : p_3)$. That is, C_1 depends on both C_2 and C_3 , while C_2 depends on C_3 (see Figure 1(a)). A straightforward approach to evaluate this modified MCS is to ask in C_1 for the belief sets of C_2 and C_3 . But as C_2 also depends on C_3 , we would need another query from C_2 to C_3 to evaluate C_2 w.r.t. the belief sets of C_3 . This shows that there is some evident redundancy in this approach, as C_3 will need to compute its belief sets twice. Simple caching strategies could mellow out the second belief state building in C_3 ; nonetheless, when C_1 asks C_3 , the context will transmit back its belief states, thus consuming network resources.

Moreover, when C_2 asks for the partial equilibria of C_3 , it will receive a set of partial equilibria that covers the belief sets of C_3 and in addition all contexts in the import closure $IC(3)$. This is excessive from the view of C_1 , as it only needs to know the truth of $(2 : p_2)$ and $(3 : p_3)$. However, C_1 needs the belief states of both C_2 and C_3 in reply of C_2 : if C_2 only reports its own belief sets (which are consistent w.r.t. C_3),

then C_1 has no chance to align the belief sets received from C_2 with those received from C_3 . Realizing that C_2 also reports the belief sets of C_3 , no call to C_3 must be made.

In the following, we present an optimization strategy which hinges on this observation. It pursues two orthogonal goals: (i) to prune dependencies in an MCS and cut superfluous transmissions, belief state building, and joining of belief states; and (ii) to minimize information in transmissions.

Graph-Theoretic Concepts

We start with defining the topology of an MCS.

Definition 6 *The topology of an MCS $M = (C_1, \dots, C_n)$ is the digraph $G_M = (V, E)$, where $V = \{1, \dots, n\}$ and $(i, j) \in E$ iff some rule in br_i has an atom $(j:p)$ in the body.*

The first optimization technique is built up by three operations on graphs. To get a coarse view on the MCS, we will decompose the topology into its *biconnected components*. The latter form a *tree representation* of the MCS, and we can apply edge removal techniques in each component.

In the following, we will use standard terminology from graph theory (see, e.g., Bondy and Murty, 2008) and assume that graphs are directed by default. For any graph G and set $S \subseteq E(G)$ of edges, we denote by $G \setminus S$ the subgraph of G that has no edges from S . For a vertex $v \in V(G)$, we denote by $G \setminus v$ the subgraph of G induced by $V(G) \setminus \{v\}$. A graph is weakly connected if replacing every directed edge by an undirected edge yields a connected graph. A vertex c of a weakly connected graph G is a *cut vertex*, if $G \setminus c$ is disconnected. A *biconnected graph* is a weakly connected graph without cut vertices. A *block* in a graph G is a maximal biconnected subgraph of G . Let $T(G) = (\mathcal{B} \cup \mathcal{C}, \mathcal{E})$ denote the undirected bipartite graph, called *block tree of graph G* , where \mathcal{B} is the set of blocks of G , \mathcal{C} is the set of cut vertices of G , and $(B, c) \in \mathcal{E}$ with $B \in \mathcal{B}$ and $c \in \mathcal{C}$ iff $c \in V(B)$. Note that $T(G)$ is a rooted tree for any weakly connected graph G ; for arbitrary graphs, it is a forest.

Example 5 The topology G_M of M in Example 2 is shown in Figure 1(b). It has two cut vertices, viz. 3 and 4; thus $T(G_M)$ contains the blocks B_1, B_2 , and B_3 , which are subgraphs of G_M induced by $\{1, 2, 3, 4\}$, $\{4, 5\}$, and $\{3, 6\}$, respectively. The block tree $T(G_M)$ is shown in Figure 1(c).

Optimization

Pruning. In acyclic topologies, like the triangle presented in the previous section, we can exploit a minimal graph representation to avoid unnecessary calls between contexts. Namely, the *transitive reduction* of the graph G_M ; recall that the transitive reduction of a digraph G is the graph G^- with the smallest set of edges whose transitive closure equals the one of G . Note that G^- is unique if G is acyclic.

Another essential part of our optimization strategy is to break cycles by removing edges from topologies. To this end, we use ear decompositions of cyclic graphs. A block may have multiple cycles which are not necessarily strongly connected, thus we first decompose cyclic blocks to its strongly connected components. The topological sort of these components yield a sequence of nodes r_1, \dots, r_s that are used as

entry points to each component. The next step is to break cycles. An *ear decomposition* of a strongly connected graph G rooted at a node r is a sequence $P = \langle P_0, \dots, P_m \rangle$ of subgraphs of G such that (i) $G = P_0 \cup \dots \cup P_m$, (ii) P_0 is a simple cycle (i.e., has no repeated edges or vertices) with $r \in V(P_0)$, and (iii) each P_i ($i > 0$) is a non-trivial path (without cycles) whose endpoints are in $P_0 \cup \dots \cup P_{i-1}$, but the other nodes are not. Let $cb(G, P)$ be the set of edges containing (l, r) from P_0 , and from each P_i , $i > 0$, the last edge (l, t) in P_i .

Example 6 Block B_1 of $T(G_M)$ is acyclic; hence, we apply a transitive reduction to obtain B_1^- with edges $E(B_1^-) = \{(1, 2), (2, 3), (3, 4)\}$. The block B_2 is cyclic, and $\langle B_2 \rangle$ is the only ear decomposition rooted at 4; removing $cb(B_2, \langle B_2 \rangle) = \{(5, 4)\}$, we obtain B_2' with edges $E(B_2') = \{(4, 5)\}$. The block B_3 is acyclic and already reduced. Figure 1(c) shows the final result (dashed edges are removed).

The graph-theoretic concepts introduced here, in particular the transitive reduction of acyclic blocks and the ear decomposition of cyclic blocks, are used to implement the first optimization of MCS evaluation outlined above. Intuitively, given the transitive reduction B^- of an acyclic block $B \in \mathcal{B}$, and a total order on $V(B^-)$ that extends B^- , one can evaluate the respective contexts in reverse order for computing partial equilibria at some context C_k : the first context simply computes its local belief sets which—represented as a set of partial belief states \mathcal{S}_0 —constitutes an initial set of partial belief states \mathcal{T}_0 . In any iterative Step i , \mathcal{T}_{i-1} is updated by joining it with the local belief sets \mathcal{S}_i of the context under consideration. Given \mathcal{T}_k (after updating with \mathcal{S}_k) for context C_k , it holds that $\mathcal{T}_k|_{V^*(k)}$ is the set of partial equilibria at C_k (restricted to contexts in $V(B^-)$).

For cyclic blocks, one can in principle proceed as above; however, any context C_k accessing beliefs from a context C_i that precedes it in the given total order, has to temporarily consider all possible belief sets for C_i in \mathcal{S}_k . As a consequence, the above relation to partial equilibria can only be established after visiting all contexts that have been temporarily considered in previous steps.

Refined recursive import. Next, we define the second part of our optimization strategy which handles minimization of information needed for transmission between two neighboring contexts C_i and C_j . For this purpose, we refine the notion of recursive import interface in a context w.r.t. a particular neighbor, and a given (sub-)graph.

Definition 7 *Given an MCS $M = (C_1, \dots, C_n)$ and a subgraph G of G_M , for an edge $(i, j) \in E(G)$, the recursive import interface of C_i to C_j w.r.t. G is $V^*(i, j)_G = \{p \in V^*(i) \mid p \in \Sigma_\ell, j \text{ reaches } \ell \text{ in } G\}$.*

Intuitively, if a context is a cut vertex c in G_M , one can drop all entries \mathcal{S}_i ($i \neq c$) from the partial belief states computed at c , and pass this result to the parent block of c in $T(G_M)$, without compromising the computation of compatible (restricted) belief sets at the parent. Recursive import interfaces w.r.t. blocks in G_M reflect this property, which can be exploited for minimizing the information transmitted.

Algorithms. Algorithms 1 and 2 combine the optimization techniques outlined above. Intuitively, `OptimizeTree` takes

Algorithm 1: OptimizeTree(T, c_p, c_r)

Input: $T = (\mathcal{B} \cup \mathcal{C}, \mathcal{E})$: block tree, c_p, c_r : context ids
Output: F : removed edges from $\bigcup \mathcal{B}$, v : labels for $\bigcup \mathcal{B}$
 $B' := \emptyset, F := \emptyset, v := \emptyset$
if $c_p = c_r$ **then** $B' := \{B \in \mathcal{B} \mid c_r \in V(B)\}$
else $B' := \{B \in \mathcal{B} \mid (B, c_p) \in \mathcal{E}\}$
foreach $B \in B'$ **do** // sibling blocks B of parent c_p
 $F := \text{OptimizeBlock}(B, c_p)$ // prune block
 $\mathcal{C}' := \{c \in \mathcal{C} \mid (B, c) \in \mathcal{E} \wedge c \neq c_p\}$ // children of B
 $B' := B \setminus F$
 foreach $(i, j) \in E(B')$ **do** // setup interface of B'
 $v(i, j) := V^*(i, j)_{B'} \cup \bigcup_{c \in \mathcal{C}'} V^*(c_r)_{\Sigma_c} \cup \bigcup_{(l, m) \in F} V^*(c_p)_{\Sigma_m}$
 foreach $c \in \mathcal{C}'$ **do** // accumulate children
 $(F', v') := \text{OptimizeTree}(T \setminus B, c, c_p)$
 $F := F \cup F', v := v \cup v'$
return (F, v)

Algorithm 2: OptimizeBlock(G : graph, r : context id)

if G is cyclic **then** // ear decomp. of strong components
 $F := \text{CycleBreaker}(G, r)$
else $F := \emptyset$
Let G^- be the transitive reduction of $G \setminus F$
return $E(G) \setminus E(G^-)$ // removed edges from G

a block tree T as input together with parent cut vertex c_p and root cut vertex c_r . It traverses T in a DFS-way and calls `OptimizeBlock` on every block. The result of the latter calls are removed edges F ; after all blocks have been processed, the final result of `OptimizeTree` is a pair of all edges removed from blocks in T , and a labelling v for the remaining edges. `OptimizeBlock` takes a graph G and calls subroutine `CycleBreaker` for cyclic G , which decomposes G into its strongly connected components, creates an ear decomposition P for each component G_c , and breaks cycles by removing edges $cb(G_c, P)$. For the resulting acyclic subgraph of G (or if G was already acyclic), `OptimizeBlock` computes the transitive reduction G^- . All edges removed from G are returned. `OptimizeTree` continues computing the labelling v for the remaining edges, building on the refined notion of recursive import interface, but keeping relevant interface variables of child cut vertices and removed edges. It can be shown that:

Proposition 1 For any context C_k in an MCS M , the algorithm `OptimizeTree`($T(G_M), k, k$) returns a pair (E, v) such that (i) the subgraph G of $G_M \setminus E$ induced by $IC(k)$ is acyclic, and (ii) for all $(i, j) \in E(G)$, $v(i, j) = V^*(i, j)_G$.

Given G_M , the block tree graph $T(G_M)$ can be constructed in linear time; transitive reductions thereof can be computed in quadratic time. Since no other operation of the algorithm exceeds this bound, the following holds.

Proposition 2 For any context C_k in an MCS M , the algorithm `OptimizeTree`($T(G_M), k, k$) runs in time polynomial (quadratic) in the size of $T(G_M)$ resp. G_M .

Given the topology of an MCS, we need to represent a stripped version of it which contains both the minimal depen-

dencies between contexts and interface variables that need to be transferred between contexts. This representation will be a *query plan* that can be used for execution processing. Syntactically, query plans have the following form.

Definition 8 (Query Plan) A query plan of an MCS M w.r.t. context C_k is any labeled subgraph Π of G_M induced by $IC(k)$ with $E(\Pi) \subseteq E(G_M)$, and edge labels $v: E(G) \rightarrow 2^{\Sigma}$.

In particular, for any MCS M and context C_k of M , the graph $\Pi_k = (V(G), E(G) \setminus E, v)$ is a query plan of M w.r.t. C_k , where G is the subgraph of G_M induced by $IC(k)$ and $(E, v) = \text{OptimizeTree}(T(G_M), k, k)$. This query plan is in fact effective; we show how to use it for MCS evaluation.

MCS Evaluation with Query Plans

Given an MCS M and a starting context C_k , we aim at finding all projected partial equilibria of M w.r.t. C_k in a distributed way. To this end, we design an algorithm `DMCSOPT` that is based on the algorithm `DMCS` in (Dao-Tran et al. 2010), but exploits properties of the optimization techniques described above. As a by-product, we obtain a simplification, because explicit cycle breaking is not needed. At each context node, an instance of `DMCSOPT` runs independently and communicates with other instances for exchanging sets of partial belief states. This provides a method for distributed model building, such that `DMCSOPT` can be deployed to any MCS where appropriate solvers for the respective context logics are available. The main feature of `DMCSOPT` is that it computes projected partial equilibria based on a query plan. This can be exploited for specific tasks like, e.g., local query answering or consistency checking. When computing projected partial equilibria, the information communicated between contexts is minimized, keeping communication cost low.

In the sequel, we present a basic version of the algorithm, abstracting from low-level implementation issues. The idea is as follows: we start with context C_k and traverse a given query plan by expanding the outgoing edges of that plan at each context, like in a depth-first search, until a leaf context is reached. A leaf context C_i simply computes its local belief sets, transforms all belief sets into partial belief states, and returns this result to its parent. If the leaf C_i contains $(j : p)$ in bodies of bridge rules such that there is no context C_j to visit in the query plan—this means we broke a cycle by removing the last edge to C_j —, all possible truth assignments to the import interface to C_j are considered.

The result of any context C_i is a set of partial belief states, which amounts to the join, i.e., the consistent combination, of its local belief sets with the results of its neighbors; the final result is obtained from C_k . To keep re-computation and recombination of belief states with local belief sets at a minimum, partial belief states are cached in every context.

Algorithm 3 shows our distributed algorithm, `DMCSOPT`, with its instance at context C_k that runs in a background process (or daemon in Unix). On input of the id c of a predecessor context (which the process awaits), it proceeds based on an (acyclic) query plan Π_k . The algorithm maintains a cache $cache(k)$ at C_k , which is kept persistent by the background process. It uses the following primitives:

Algorithm 3: DMCSOPT(c) at $C_k = (L_k, kb_k, br_k)$

Input: c : context identifier of a direct predecessor
Data: Π_k : query plan with label v , $cache(k)$: cache
Output: set of accumulated partial belief states

(a) **if** $cache(k)$ is not empty **then** $\mathcal{S} := cache(k)$
else
 $\mathcal{T} := \{(\epsilon, \dots, \epsilon)\}$
 (b) **foreach** $(k, i) \in E(\Pi_k)$ **do**
 $\mathcal{T} := \mathcal{T} \bowtie C_i.DMCSOPT(k)$
 (c) **if** there is $(i:p)$ in br_k s.t. $(k, i) \notin E(\Pi_k)$ and
 $T_i \in \mathcal{T}$ s.t. $T_i = \epsilon$ **then** $\mathcal{T} := guess(v(c, k)) \bowtie \mathcal{T}$
 (d) **foreach** $T \in \mathcal{T}$ **do** $\mathcal{S} := \mathcal{S} \cup solve(T)$
 $cache(k) := \mathcal{S}$
 (e) **if** $(c, k) \in E(\Pi_k)$ **then** $\mathcal{S} := \mathcal{S}|_{v(c, k)}$ // C_k is non-root
return \mathcal{S}

Algorithm 4: Isolve(S) at $C_k = (L_k, kb_k, br_k)$

Input: S : partial belief state
Output: set of locally acceptable partial belief states
 $\mathbf{T} := ACC_k(kb_k \cup \{head(r) \mid r \in app(br_k, S)\})$
return $\{(S_1, \dots, S_{k-1}, T_k, S_{k+1}, \dots, S_n) \mid T_k \in \mathbf{T}\}$

- function $C_i.DMCSOPT(c)$: send context id c to the process at context C_i and wait for its return result.
- function $guess(V)$: guess all possible truth assignments for the interface variables V .
- function $Isolve(S)$ (Algorithm 4): given a partial belief state S , augment kb_k with the heads of all bridge rules in br_k that are applicable w.r.t. S ($=: kb'_k$), compute the local belief sets by $ACC(kb'_k)$, and combine them with S ; return the resulting set of partial belief states.

The steps of Algorithm 3 are explained as follows:

- (a) check the cache for an appropriate partial belief state;
 (b) get neighbor contexts from the query plan, request partial belief states from all neighbors and join them;
 (c) if there are $(i:p)$ in the bridge rules br_k such that $(k, i) \notin E(\Pi_k)$, and no neighbor delivered the belief sets for C_i in step (b) (i.e., $T_i = \epsilon$), we have to call $guess$ on the interface $v(c, k)$ and join the result with \mathcal{T} : intuitively, this happens when edges had been removed from cycles.
 (d) compute local belief states given the imported partial belief states collected from neighbors;
 (e) project to the variables in $v(c, k)$ for non-root contexts.

The following proposition shows that DMCSOPT is sound and complete.

Proposition 3 *Let C_k be a context of an MCS M , let Π_k be the query plan as defined above and let $V = \{p \in v(k, j) \mid (k, j) \in E(\Pi_k)\}$. Then,*

- for each $S' \in C_k.DMCSOPT(k)$, there exists a partial equilibrium S of M w.r.t. C_k such that $S' = S|_V$;
- for each partial equilibrium S of M w.r.t. C_k , there exists an $S' \in C_k.DMCSOPT(k)$ such that $S' = S|_V$.

Implementation and Experimental Results

We present initial results for a SAT-solver based prototype implementation of DMCSOPT under Ubuntu Linux 9.10, writ-

topology / parameter	DMCSOPT	#	DMCS	#
$D_1 / (13, 9, 4, 4)$	1.365	16	1.909	2048
$D_2 / (13, 9, 4, 4)$	2.310	42	5.485	4256
$D_3 / (25, 9, 4, 4)$	38.622	6	—	—
$D_4 / (25, 9, 4, 4)$	14.763	17	—	—
$D_5 / (28, 9, 4, 4)$	130.900	16	—	—
$D_6 / (28, 9, 4, 4)$	99.102	16	—	—
$R_2 / (10, 9, 4, 4)$	0.246	2	3.262	43008
$R_3 / (13, 9, 4, 4)$	0.253	6	1.833	31488
$R_4 / (13, 9, 4, 4)$	0.447	6	3.148	25056
$R_5 / (301, 9, 4, 4)$	12.380	4	—	—
$R_6 / (301, 9, 4, 4)$	12.718	20	—	—
$Z_1 / (13, 9, 4, 4)$	2.296	6	33.549	1152
$Z_2 / (13, 9, 4, 4)$	4.844	30	66.557	5160
$Z_3 / (151, 9, 4, 4)$	36.071	96	—	—
$Z_4 / (151, 9, 4, 4)$	29.529	8	—	—
$Z_5 / (301, 9, 4, 4)$	96.887	40	—	—
$Z_6 / (301, 9, 4, 4)$	116.060	8	—	—

Table 1: Runtime in secs, timeout 180 secs (—)

ten in C++ (available at <http://www.kr.tuwien.ac.at/research/systems/dmcs/>). The host system was using a Pentium Core2 Duo 2.53GHz processor with 4GB RAM. We compare response times of DMCSOPT to DMCS.

We used the development version of *clasp* (2010-01-31) as a SAT solver, which accepts DIMACS CNF input (Gebser et al. 2007). Specifically, all generated instantiations of multi-context systems have contexts with ASP logics. We use the translation defined in (Dao-Tran et al. 2010) to create SAT instances at contexts C_k and *clasp* to compute all models.

For initial experimentation, we created random MCS instances of ordinary and zig-zag diamond stack, ring, and binary tree topologies. A diamond stack combines multiple diamonds in a row (stacking m diamonds in a tower of $3m+1$ contexts). Ordinary diamonds have, in contrast to zig-zag diamonds like block B_1 in Figure 1(c), no connection between the two middle contexts. Other topologies are currently under consideration, including variants of those presented here.

A parameter setting (n, s, b, r) specifies (i) the number n of contexts, (ii) the local alphabet size $|\Sigma_i| = s$ (each C_i has a random ASP program on s atoms with 2^k answer sets, $0 \leq k \leq s/2$), (iii) the maximum interface size b (number of atoms exported), (iv) and the maximum number r of bridge rules per context, each having ≤ 2 body literals.

Table 1 shows some experimental results for parameter settings $(n, 9, 4, 4)$, where n varies between 10 and 301. Each row X_i ($X \in \{D, R, Z\}$) displays pure computation time (no output) for ordinary diamond stacks (D), zig-zag diamond stacks (Z), and rings (R), where the # columns show the numbers of projected partial equilibria computed at C_1 (initiated by sending the request 1 to C_1 for DMCSOPT with a fixed query plan Π_1 , respectively $V^*(1)$ to C_1 for DMCS).

The optimizations that can be applied in the topologies are quite diverse. In ordinary diamond stacks and in binary trees, we cannot remove edges, as the topologies are equal to their transitive reductions. But we can refine the import interface at each sub-diamond (every fourth context is a cut

vertex), thus the partial belief states eventually computed just contain entries for the first four contexts. The refinement of the import interface in binary trees is even more drastic, as every non-leaf context is a cut vertex, and we can restrict to import interfaces between two neighboring contexts. In the ring topology, we can remove the last edge closing the cycle to context C_1 . As the resulting topology is a spanning tree, the refinement of the import interface is restricted to neighboring contexts including the import interface of the removed edge. In zig-zag diamond stacks, we remove in each block two edges to obtain the transitive reduction and update the recursive import interface accordingly.

Evaluating the MCS instances with DMCSOPT compared to DMCS yields a drastic improvement in response time. Stacking multiple diamonds in a tower models hard instances with many joins. This is reflected in the ratio of running time to result size in DMCS. Still, DMCSOPT could handle much larger instances. The ring topology shows a similar increase in scalability. Thanks to the refined interface, the system size can be increased dramatically; the runs for $n=301$ took only a few seconds. Also for zig-zag diamond stacks, the optimizations effect that DMCSOPT can run substantially larger systems, with hundreds of contexts; DMCS has an early breakdown at $n=16$. Comparing ordinary diamonds (D_i) to zig-zag diamonds (Z_i), one can notice a large gap in the size n of the MCS that can be handled. This is explained by the transitive reduction that can be applied to zig-zag diamonds, essentially resulting in a chain of contexts such that each context can take the partial belief states of its single neighbor and simply add its local beliefs. Diamonds D_i cannot be further optimized and additionally need to join the results from their neighbors. We omit detailed outcomes for binary tree topology tests here. However, we noticeably could evaluate an instance with even $n=700$ contexts in 68.061 seconds ($\# = 48$) with our parameter setting; setting the timeout to 12 minutes, DMCS runs out of memory for instances with $n=22$ during belief state joining.

Related Work and Conclusion

Roelofsen, Serafini, and Cimatti (2004) described evaluation of monotone MCS with classical theories using SAT solvers for the contexts in parallel. They used a (co-inductive) fix-point strategy to check MCS satisfiability, where a centralized process iteratively combines results of the SAT solvers. Apart from being not truly distributed, an extension to nonmonotonic MCS is non-obvious; also, no caching was used.

Serafini and Tamin (2005) and Serafini, Borgida, and Tamin (2005) defined distributed tableaux algorithms for reasoning in distributed ontologies. They can be used to decide consistency of distributed description logic knowledge bases, provided that the distributed TBox is acyclic. The DRAGO system is an implementation of this approach.

Adjiman et al. (2006) presented a framework of peer-to-peer inference systems. Local theories of propositional clause sets share atoms, and a special algorithm can be used for consequence finding. As we pursue the dual problem of model building, application for our needs is not straightforward.

Similar to this, Bikakis and Antoniou (2008) developed a distributed algorithm for query evaluation in a multi-context

system framework based on defeasible logic. Contexts in this setup are built using defeasible rules, and the query evaluation procedure can determine for a given literal l three values: whether l is (not) a logical conclusion of the MCS, or whether it cannot be proved that l is a logical conclusion. Again, applying this approach to model building is not easy.

Baget and Tognetti (2001) use biconnected components to decompose constraint satisfaction problems. The decomposition is used to localize the computation of a single solution in the components of undirected constraint graphs. Likened to our approach, we are based on directed dependencies, which gives us the chance to use a query plan for MCS evaluation.

We have presented techniques and algorithms for decomposing, pruning, and cycle breaking of dependencies in non-monotonic multi-context systems. Based on this, we have devised an algorithm, which uses a query plan to compute all partial equilibria of such a system. A prototypical implementation of this approach shows promising experimental results. They are a substantial improvement and encourage to research further algorithms and methods for evaluation of distributed MCS, such that efficient platforms for distributed nonmonotonic reasoning applications will become available.

References

- Adjiman, P.; Chatalic, P.; Goasdoué, F.; Rousset, M.-C.; and Simon, L. 2006. Distributed Reasoning in a Peer-to-Peer Setting: Application to the Semantic Web. *J. Artif. Intell. Res.* 25:269–314.
- Baget, J.-F., and Tognetti, Y. 2001. Backtracking through biconnected components of a constraint graph. In *IJCAI'01*, 291–296.
- Bondy, A., and Murty, U. S. R. 2008. *Graph Theory*. Springer.
- Bikakis, A., and Antoniou, G. 2008. Distributed Defeasible Reasoning in Multi-Context Systems. In *NMR'08*, 200–206.
- Brewka, G., and Eiter, T. 2007. Equilibria in heterogeneous nonmonotonic multi-context systems. In *AAAI'07*, 385–390.
- Buccafurri, F., and Caminiti, G. 2008. Logic programming with social features. *Theory Pract. Log. Program.* 8(5–6):643–690.
- Dao-Tran, M.; Eiter, T.; Fink, M.; and Krennwallner, T. 2010. Distributed nonmonotonic multi-context systems. In *KR'10*.
- Gebser, M.; Kaufmann, B.; Neumann, A.; and Schaub, T. 2007. Conflict-driven answer set solving. In *IJCAI'07*, 386–392.
- Gelfond, M., and Lifschitz, V. 1991. Classical negation in logic programs and disjunctive databases. *New Gener. Comput.* 9:365–385.
- Ghidini, C., and Giunchiglia, F. 2001. Local models semantics, or contextual reasoning=locality+compatibility. *Artif. Intell.* 127(2):221–259.
- Giunchiglia, F., and Serafini, L. 1994. Multilanguage hierarchical logics or: how we can do without modal logics. *Artif. Intell.* 65(1):29–70.
- McCarthy, J. 1993. Notes on formalizing context. In *IJCAI'93*.
- Roelofsen, F.; Serafini, L.; and Cimatti, A. 2004. Many hands make light work: localized satisfiability for multi-context systems. In *ECAI'04*, 58–62.
- Serafini, L., and Tamin, A. 2005. Drago: Distributed reasoning architecture for the semantic web. In *ESWC'05*, 361–376.
- Serafini, L.; Borgida, A.; and Tamin, A. 2005. Aspects of distributed and modular ontology reasoning. In *IJCAI'05*, 570–575.