

Reactive Multi-Context Systems: Heterogeneous Reasoning in Dynamic Environments

Gerhard Brewka^a, Stefan Ellmauthaler^a, Ricardo Gonçalves^b, Matthias Knorr^{b,*},
João Leite^b, Jörg Pührer^{a,*}

^a*Institute of Computer Science, Leipzig University, Germany*

^b*NOVA LINCS & Departamento de Informática, Universidade NOVA de Lisboa, Portugal*

Abstract

Managed multi-context systems (mMCSs) allow for the integration of heterogeneous knowledge sources in a modular and very general way. They were, however, mainly designed for static scenarios and are therefore not well-suited for dynamic environments in which continuous reasoning over such heterogeneous knowledge with constantly arriving streams of data is necessary. In this paper, we introduce reactive multi-context systems (rMCSs), a framework for reactive reasoning in the presence of heterogeneous knowledge sources and data streams. We show that rMCSs are indeed well-suited for this purpose by illustrating how several typical problems arising in the context of stream reasoning can be handled using them, by showing how inconsistencies possibly occurring in the integration of multiple knowledge sources can be handled, and by arguing that the potential non-determinism of rMCSs can be avoided if needed using an alternative, more skeptical well-founded semantics instead with beneficial computational properties. We also investigate the computational complexity of various reasoning problems related to rMCSs. Finally, we discuss related work, and show that rMCSs do not only generalize mMCSs to dynamic settings, but also capture/extend relevant approaches w.r.t. dynamics in knowledge representation and stream reasoning.

Keywords: Heterogeneous knowledge, Stream reasoning, Knowledge integration, Reactive systems, Dynamic systems

1. Introduction

Fueled by initiatives such as the Semantic Web, Linked Open Data, and the Internet of Things, among others, the wide and increasing availability of machine-processable data and knowledge has prepared the ground and called for a new class of dynamic,

*Corresponding author

Email addresses: brewka@informatik.uni-leipzig.de (Gerhard Brewka),
ellmauthaler@informatik.uni-leipzig.de (Stefan Ellmauthaler), rjrg@fct.unl.pt
(Ricardo Gonçalves), mkn@fct.unl.pt (Matthias Knorr), jleite@fct.unl.pt (João Leite),
puehrer@informatik.uni-leipzig.de (Jörg Pührer)

rich, knowledge-intensive applications. Such new applications require automated reasoning based on the integration of several heterogeneous knowledge bases – possibly overlapping, independently developed, and written in distinct languages with different semantic assumptions – together with data/event streams produced by sensors and detectors, to be able to support automation and problem-solving, to enforce traceable and correct decisions, and to facilitate the internalization of relevant dynamic data and knowledge into such heterogeneous knowledge bases.

Consider a scenario where Dave, an elderly person suffering from dementia, lives alone in an apartment equipped with various sensors, e.g., smoke detectors, cameras, and body sensors measuring relevant body functions (e.g., pulse, blood pressure, etc.). An assisted living application in such a scenario could leverage the information continuously received from the sensors, together with Dave’s medical records stored in a relational database, a biomedical health ontology with information about diseases, their symptoms and treatments, represented in some description logic, some action policy rules represented as a non-monotonic logic program, to name only a few, and use it to detect relevant events, suggest appropriate action, and even raise alarms, while keeping a history of relevant events and Dave’s medical records up to date, thus allowing him to live on his own despite his condition. After detecting that Dave left the room while preparing a meal, the system could alert him in case he does not return soon, or even turn the stove off in case it detects that Dave fell asleep, not wanting to wake him up because his current treatment/health status values rest over immediate nutrition. Naturally, if Dave is not gone long enough, and no sensor shows any potential problems (smoke, gas, fire, etc.), then the system should seamlessly take no action.

Given the requirements posed by novel applications such as the one just described, the availability of a vast number of knowledge bases – written using many different formalisms – and the relevance of streams of data/events produced by sensors/detectors, modern research in knowledge representation and reasoning faces two fundamental problems: dealing with the **integration** of heterogeneous data and knowledge, and dealing with the **dynamics** of such novel knowledge-based systems.

Integration. The first problem stems from the availability of knowledge bases written in many different languages and formats developed over the last decades, from the rather basic ones, such as relational databases or the more recent triplestores, to the more expressive ones, such as ontology languages (e.g., description logics), temporal and modal logics, non-monotonic logics, or logic programs under answer set semantics, to name just a few. Each of these formalisms was developed for different purposes and with different design goals in mind. Whereas some of these formalisms could be combined to form a new, more expressive formalism, with features from its constituents – such as dl-programs [16] and Hybrid MKNF [35] which, to different extent, combine description logics and logic programs under answer set semantics – in general this is simply not feasible, either due to the mismatch between certain assumptions underlying their semantics, or because of the high price to pay, often in terms of complexity, sometimes even in terms of decidability. It is nowadays widely accepted that there simply is no such thing as a single universal, general purpose knowledge representation language.

What seems to be needed is a principled way of integrating knowledge expressed

in different formalisms. Multi-context systems (MCSs) provide a general framework for this kind of integration. The basic idea underlying MCSs is to leave the diverse formalisms and knowledge bases untouched, and to use so-called bridge rules to model the flow of information among different parts of the system. An MCS consists of reasoning units – called contexts for historical reasons [27] – where each unit is equipped with a collection of bridge rules. In a nutshell, the bridge rules allow contexts to “listen” to other contexts, that is to take into account beliefs held in other contexts.

Bridge rules are similar to logic programming rules (including default negation), with an important difference: they provide means to access other contexts in their bodies. Bridge rules not only allow for a fully declarative specification of the information flow, but they also allow information to be modified instead of being just passed along as is. Using bridge rules we may translate a piece of information into the language/format of another context, pass on an abstraction of the original information, leaving out unnecessary details, select or hide information, add conclusions to a context based on the absence of information in another one, and even use simple encodings of preferences among parent contexts.

MCSs went through several development steps until they reached their present form. Advancing work in [26, 34] aiming to integrate different inference systems, monotonic heterogeneous multi-context systems were defined in [27], with reasoning within as well as across monotonic contexts. The first, still limited attempts to include non-monotonic reasoning were done in [39, 12], where default negation in the rules is used to allow for reasoning based on the *absence* of information from a context.

The non-monotonic MCSs of [9] substantially generalize previous approaches, by accommodating *heterogeneous* and both *monotonic* and *non-monotonic* contexts. Hence, they are capable of integrating, among many others, “typical” monotonic logics like description logics or temporal logics, and non-monotonic formalisms like Reiter’s default logic, logic programs under answer set semantics, circumscription, defeasible logic, or theories in autoepistemic logic. The semantics of nonmonotonic MCSs is defined in terms of equilibria: a belief set for each context that is acceptable for its knowledge base augmented by the heads of its applicable bridge rules.

More recently, the so-called managed MCSs (mMCSs) [10] addressed a limitation of MCSs in the way they integrate knowledge between contexts. Instead of simply *adding* the head of an applicable bridge rule to the context’s knowledge base, which could cause some inconsistency, mMCSs allow for operations other than addition, such as, for instance, *revision* and *deletion*, hence dealing with the problem of consistency management within contexts.

Dynamics. The second problem stems from the shift from static knowledge-based systems that assume a one-shot computation, usually triggered by a user query, to open and dynamic scenarios where there is a need to react and evolve in the presence of incoming information.

Indeed, traditional knowledge-based systems – including the different variants of MCSs mentioned above – focus entirely on static situations, which is the right thing for applications such as for instance expert systems, configuration or planning problems, where the available background knowledge changes rather slowly, if at all, and where all that is needed is the solution of a new instance of a known problem. However, the

new kinds of applications we consider are becoming more and more important, and these require continuous online reasoning, including observing and reacting to events.

There are some examples of systems developed with the purpose of reacting to streams of incoming information, such as Reactive ASP [23, 22], C-SPARQL [6], Ontology Streams [32] and ETALIS [3], to name only a few. However, they are very limited in the kind of knowledge that can be represented, and the kind of reasoning allowed, hence unsuitable to address the requirements of the applications we envision, such as those that need to integrate heterogeneous knowledge bases. Additionally, reacting to the streams of incoming information is only part of the dynamic requirements of our applications. In many cases, the incoming information is processed only once, perhaps requiring complex reasoning using various knowledge bases to infer the right way to react, and does not have to be dealt with again – e.g., concluding that nothing needs to be done after determining that the tachycardia is caused by the decongestant recently taken by Dave. In other cases, it is important that these observations not only influence the current reaction of the system – do nothing in the previous example – but, at the same time, be able to change the knowledge bases in a more permanent way, i.e., allowing for the internalization of knowledge. For example, relevant observations regarding Dave’s health status should be added to his medical records, such as for example that he had an episode of tachycardia caused by a decongestant, and, in the future, maybe even revise such episode if it is found that Dave had forgotten to take the decongestant after all. Other more sophisticated changes in the knowledge bases include, for example, an update to the biomedical health ontology whenever new treatments are found or the revision of the policy rules whenever some exceptions are found. EVOLP [2] extends logic programming under answer set semantics with the possibility to specify its evolution, through successive updates, in reaction to external observations. It is nevertheless limited to a single knowledge representation formalism and to a single operation (update).

Our Approach. In this paper, we aim to address both challenges in a single, uniform approach. We develop a system that allows us to integrate heterogeneous knowledge bases with streams of incoming information and to use them for continuous online reasoning, reacting, and evolving the knowledge bases by internalizing relevant knowledge. In a nutshell, our work follows the multi-context systems tradition, but adds what is needed to also deal with dynamics.

To this end, we introduce *reactive Multi-Context Systems (rMCSs)*. These systems build upon mMCSs and thus provide their functionality for integrating heterogeneous knowledge sources, admitting also operations for manipulating their knowledge bases. In addition, rMCSs can handle continuous streams of input data. Equilibria remain the fundamental underlying semantic notion, but the focus now lies on the dynamic evolution of the systems. Given an initial configuration of knowledge bases, that is, an initial knowledge base for each context, a specific input stream will lead to a corresponding stream of equilibria, generated by respective updates of the knowledge bases.

Contrary to existing MCSs which possess only one type of bridge rules modeling the information flow to be taken into account when computing equilibria (or the operations that need to be applied in case of mMCSs), rMCSs have an additional, different type of bridge rules, distinguished by the occurrence of the operator **next** in the head.

Whereas the head of a managing bridge rule in mMCS specifies how a specific knowledge base should temporarily be changed for the purpose of determining the current equilibrium, the head of a new kind of bridge rule in rMCS with the operator **next** in the head specifies how the knowledge base should permanently be changed as the system *moves* to the next state. These new bridge rules allow us to decouple the computation of equilibria from specifying the evolution of knowledge base configurations.

It is often the case that the information *currently* available from some stream needs to be taken into account by some context when computing the semantics at that point. The way such information is taken into account is specified by a bridge rule with a management operation in its head, just as in mMCSs. For example, we may have some default value for some variable stored in a context, but, under some conditions (*cond1*), prefer to use the value that is read from some sensor (observed in some stream) whenever one is available. This could be specified by the bridge rule (where *set* is a management operation with the obvious meaning):

$$\text{set}(\text{value}(X)) \leftarrow st::\text{meter}(X), \text{cond1}.$$

Whenever *meter*(*X*) is observed in the stream, and *cond1* holds, the *value*(*X*) would be matched, but only while *meter*(*X*) is observed and *cond1* holds, i.e., the management operation's effects are not permanent. In subsequent states, if there was no observation for *meter*(*X*), or *cond1* no longer holds, then *value*(*X*) would return to its original value. However, sometimes we may wish to specify some permanent change to the knowledge base of some context that persists beyond the state at which the pre-conditions were true, which is why a new form of bridge rule is needed. In the previous example, if under certain conditions (*cond2*) we want to internalize the current observation in a permanent way, i.e., permanently update the default value, then the following new kind of bridge rule would be specified:

$$\text{next}(\text{set}(\text{value}(X))) \leftarrow st::\text{meter}(X), \text{cond2}.$$

The intuitive reading of this bridge rule is that whenever at some state both *meter*(*X*) is true in the stream *st* and condition *cond2* holds, then the bridge rule's context should be changed according to the operation *set*(*value*(*X*)) when moving to the next state. These two different behaviors could not be achieved with a single kind of bridge rule since one of them aims at temporarily affecting the system at current state while the other one aims at permanently affecting the system from the subsequent state. Interestingly, the two kinds could be combined to specify that some current effect should be made permanent.

This new kind of bridge rule further allows the specification of the evolution of some context based on current conditions, which was not possible in mMCS. Suppose, for example, a single context that maintains a counter *ctr*(.) keeping track of how often the context was updated. With the new kind of bridge rule, we could simply write:

$$\text{next}(\text{set}(\text{ctr}(N + 1))) \leftarrow \text{ctr}(N).$$

Assuming that the management operation *set* has the intuitively expected meaning, using the bridge rule

$$\text{set}(\text{ctr}(N + 1)) \leftarrow \text{ctr}(N).$$

would not work, as it would lead to the non-existence of equilibria. In mMCSs, one might think of a workaround that would overload the management function, but that would also require that the language of the context be extended, which might not be possible and is clearly at odds with the very idea of declarative knowledge representation. Additional examples showing the benefits of the **next** operator will be discussed in Sect. 3.2.

The main contributions of this paper can be summarized as follows:

- We extend multi-context systems with the concepts relevant to deal with dynamics, ultimately leading to the notion of *equilibria streams*.
- We introduce a new, second type of bridge rules using the operator **next**. This allows us to separate the computation of equilibria from the specification of the context evolution which provides a lot more modeling flexibility and leads to much more natural problem representations.
- We study various forms of inconsistency, generalize existing inconsistency handling techniques to rMCSs and introduce several new ones.
- For a subclass of rMCSs we define well-founded semantics as a skeptical, deterministic alternative to the equilibria-based semantics. The new semantics gives rise to so-called grounded equilibria streams.
- We investigate the complexity of various problems related to rMCSs and provide a number of complexity results.

The paper is organized as follows. In Section 2, we introduce reactive MCSs, our framework for reactive reasoning in the presence of heterogeneous knowledge sources. In particular, we show how to integrate data streams into mMCSs and how to model the dynamics of our systems, based on two types of bridge rules. Section 3 illustrates how several typical problems arising in the context of stream reasoning can be handled using our framework. Reasoning based on multiple knowledge sources that need to be integrated faces the problem of potential inconsistencies. Section 4 discusses various methods for handling inconsistencies, with a special focus on non-existence of equilibria. In particular, we show how methods developed for mMCSs can be generalized to rMCSs. Nondeterminism in rMCSs is discussed in Section 5. One way of avoiding nondeterminism is by applying an alternative, skeptical semantics. We show how such a semantics, called well-founded semantics, can be defined for rMCSs, and what the effect of using this semantics instead of the original one is. The complexity of various problems related to rMCSs is investigated in Section 6. Section 7 discusses related work, with a special focus on two of the most relevant approaches w.r.t. stream reasoning, namely LARS (Logic-based framework for Analyzing Reasoning over Streams) [7] and STARQL [37]. Section 8 concludes and points out directions for future work.

This paper combines and unifies the results of [11] and [29], two papers by different subsets of the authors describing independent adaptations of multi-context systems for dynamic environments. The approach developed here generalizes these earlier approaches and substantially improves on the presentation of the underlying concepts.

2. Reactive Multi-Context Systems

Reactive multi-context systems (rMCSs) make use of basic ideas from managed multi-context systems (mMCSs) [10] which extend multi-context systems (MCSs) as defined by Brewka and Eiter [9] by management capabilities. In particular, similar to mMCSs, we will make use of a management function and bridge rules that allow for conflict resolution between contexts as well as a fine-grained declarative specification of the information flow between contexts. To not unnecessarily burden the reader with repetitive material on these common components, we abstain from recalling the details of mMCSs first. It will be clear from the presentation when new concepts/ideas specific to rMCSs will be presented.

2.1. Specifying the Components of an rMCS

Similar as for previous notions of MCSs, we build on an abstract notion of a *logic*, which is a triple $L = \langle KB, BS, \mathbf{acc} \rangle$, where KB is the *set of admissible knowledge bases* of L , BS is the *set of possible belief sets*, whose elements are called *beliefs*; and $\mathbf{acc} : KB \rightarrow 2^{BS}$ is a function describing the semantics of L by assigning to each knowledge base a set of *acceptable belief sets*.¹

Example 1. *We illustrate how different formalisms we use in examples throughout the paper can be represented by the notion of a logic.*

First, consider the case of description logics (DLs) as (commonly decidable) fragments of first-order logic [5]. Given a DL language \mathcal{L} , we consider the logic $L_d = \langle KB_d, BS_d, \mathbf{acc}_d \rangle$ where KB_d is the set of all well-formed DL knowledge bases over \mathcal{L} , also called ontologies, BS_d is the set of deductively closed subsets of \mathcal{L} , and \mathbf{acc}_d maps every $kb \in KB_d$ to $\{E\}$, where E is the set of formulas in \mathcal{L} entailed by kb .

As an example for a non-deterministic formalism, consider logic programs under the answer set semantics [25]. Given a set of ground, i.e., variable-free, atoms A , we use the logic $L_a = \langle KB_a, BS_a, \mathbf{acc}_a \rangle$ such that KB_a is the set of all logic programs over A . The set of possible belief sets is given by the set $BS_a = 2^A$ of possible answer sets and the function \mathbf{acc}_a maps every logic program to the set of its answer sets.

Given a set E of entries, a simple logic for storing elements from E can be realized by the logic $L_s = \langle KB_s, BS_s, \mathbf{acc}_s \rangle$, such that $KB_s = BS_s = 2^E$, and \mathbf{acc}_s maps every set $E' \subseteq E$ to $\{E'\}$. Such L_s can, e.g., be used to represent a simple database logic. We will call a logic of this type a storage logic.

In addition to a logic that captures language and semantics of a formalism to be integrated in an rMCS, a context also describes how a knowledge base belonging to the logic can be manipulated.

Definition 1 (Context). *A context is a triple $C = \langle L, OP, \mathbf{mng} \rangle$ where*

¹To ease readability, throughout the paper, we will often use the following convention when writing symbols: single entities are lower-case, while sets of entities and structures with different components are upper-case; in addition, sequences of those are indicated in `sans serif`, while notions with a temporal dimension are written in calligraphic letters (only upper-case, such as \mathcal{S} or \mathcal{I}); finally, operators and functions are **bold**.

- $L = \langle KB, BS, \mathbf{acc} \rangle$ is a logic,
- OP is a set of operations,
- $\mathbf{mng} : 2^{OP} \times KB \rightarrow KB$ is a management function.

For an indexed context C_i we will write $L_i = \langle KB_i, BS_i, \mathbf{acc}_i \rangle$, OP_i , and \mathbf{mng}_i to denote its components. Note that we leave the exact nature of the operations in OP unspecified – they can be seen as mere labels whose semantics is determined by the management function – and that we use a deterministic management function instead of a non-deterministic one, unlike mMCS [10].²

Example 2. Consider the assisted living scenario from the Introduction in which we want to recognize potential threats caused, e.g., by overheating of the stove. We use the context C_{st} to monitor the stove. Its logic L_{st} is a storage logic taking $E = \{\text{pw}, \text{tm}(\text{cold}), \text{tm}(\text{hot})\}$ as the set of entries, representing the stove’s power status (on if pw is present, and off otherwise) and a qualitative value for its temperature (cold/hot). The current temperature and power state of the stove should be kept up to date in a knowledge base over L_{st} using the following operations:

$$OP_{st} = \{\text{setPower}(\text{off}), \text{setPower}(\text{on}), \text{setTemp}(\text{cold}), \text{setTemp}(\text{hot})\}$$

The semantics of the operations is given, for $OP' \subseteq OP_{st}$, by $\mathbf{mng}_{st}(OP', kb) =$

$$\begin{aligned} & \{\text{pw} \mid \text{setPower}(\text{on}) \in OP' \vee \\ & \quad (\text{pw} \in kb \wedge \text{setPower}(\text{off}) \notin OP')\} \cup \\ & \{\text{tm}(t) \mid \text{setTemp}(t) \in OP'\}. \end{aligned}$$

We assume a single constantly operating temperature sensor which triggers exactly one of the setTemp operations. Thus, the current value is simply inserted and we do not need to care about persistence of the temperature value or conflicting information. The power information, on the other hand, is based on someone toggling a switch, which requires persistence of the fluent pw. Note that \mathbf{mng}_{st} ensures that the stove is considered on whenever it is switched on, and also when it is not being switched off and already considered on in the given knowledge base kb. The second alternative implements persistence. Also note that whenever both conflicting setPower operations are in OP' , $\text{setPower}(\text{on})$ “wins”, that is, pw will be in the knowledge base. This is justified by the application: a potential overheating of the stove is considerably more harmful than unnecessarily turning off the electricity.

Assume we have a knowledge base $kb = \{\text{tm}(\text{cold})\}$. Then, an update with the set $OP = \{\text{setPower}(\text{on}), \text{setTemp}(\text{hot})\}$ of operations would result in the knowledge base $\mathbf{mng}_{st}(OP, kb) = \{\text{pw}, \text{tm}(\text{hot})\}$.

Contexts exchange information using *bridge rules* that are rules similar in spirit to those in logic programming and that determine which operations from OP_i to apply to

²Note that it is straightforward to adapt our definitions to use non-deterministic management functions. However, as they are not essential to our approach, we here refrain from doing so to keep notation simpler.

kb_i in a context C_i depending on beliefs held in other contexts, and, in our approach also on input from the outside. To keep the approach as abstract as possible, we only require that inputs be elements of some formal *input language* IL . Moreover, we allow for situations where input comes from different sources with potentially different input languages and thus consider tuples $\langle IL_1, \dots, IL_k \rangle$ of input languages.

Definition 2 (Bridge Rule). Let $C = \langle C_1, \dots, C_n \rangle$ be a tuple of contexts and $IL = \langle IL_1, \dots, IL_k \rangle$ a tuple of input languages. A bridge rule for C_i over C and IL , $i \in \{1, \dots, n\}$, is of the form

$$\mathbf{op} \leftarrow a_1, \dots, a_j, \mathbf{not} a_{j+1}, \dots, \mathbf{not} a_m \quad (1)$$

such that $\mathbf{op} = op$ or $\mathbf{op} = \mathbf{next}(op)$ for $op \in OP_i$, $j \in \{0, \dots, m\}$, and every atom a_ℓ , $\ell \in \{1, \dots, m\}$, is one of the following:

- a context atom $c:b$ with $c \in \{1, \dots, n\}$ and $b \in B$ for some $B \in BS_c$, or
- an input atom $s::b$ with $s \in \{1, \dots, k\}$ and $b \in IL_s$.

For a bridge rule r of the form (1) $\text{hd}(r)$ denotes \mathbf{op} , the head of r , while $\text{bd}(r) = \{a_1, \dots, a_j, \mathbf{not} a_{j+1}, \dots, \mathbf{not} a_m\}$ is the body of r . A literal is either an atom or an atom preceded by \mathbf{not} , and we differentiate between context literals and input literals.

Roughly, a set of bridge rules for C_i describes which operations to apply to its knowledge base kb_i , depending on whether currently available beliefs and external inputs match the literals in the body. Intuitively, rules with head op affect the computation of the semantics at the current time instant, while rules with head $\mathbf{next}(op)$ affect the computation of the knowledge base(s) at the next time instant. We define their precise semantics later in Section 2.2 and proceed by defining reactive multi-context systems.

Definition 3 (Reactive Multi-Context System). A reactive Multi-Context System (rMCS) is a tuple $M = \langle C, IL, BR \rangle$, where

- $C = \langle C_1, \dots, C_n \rangle$ is a tuple of contexts;
- $IL = \langle IL_1, \dots, IL_k \rangle$ is a tuple of input languages;
- $BR = \langle BR_1, \dots, BR_n \rangle$ is a tuple such that each BR_i , $i \in \{1, \dots, n\}$, is a set of bridge rules for C_i over C and IL .

Example 3. Consider rMCS $M_{ex3} = \langle \langle C_{st} \rangle, \langle IL_{ex3} \rangle, \langle BR_{ex3} \rangle \rangle$ with C_{st} as in Example 2.³ The input language $IL_{ex3} = \{\text{switch}\}$ is used to report whether the power switch of the stove has been turned. The bridge rules in BR_{ex3} are given by

$$\begin{aligned} \mathbf{next}(\text{setPower}(\text{on})) &\leftarrow ex3::\text{switch}, \mathbf{not} st:\text{pw}. \\ \mathbf{next}(\text{setPower}(\text{off})) &\leftarrow ex3::\text{switch}, st:\text{pw}. \end{aligned}$$

and react to switching the stove on or off: depending on the current power state of the stove that is stored in the knowledge base of C_{st} , whenever the switch is activated, the bridge rules derive an update of the knowledge base where the power state is reversed.

³Throughout the paper we use labels (such as st in C_{st}) instead of numerical indices in our examples.

2.2. Reacting to External Inputs - Semantics of rMCSs

To define the semantics of rMCSs, we first focus on the static case of a single time instant, and only subsequently introduce the corresponding dynamic notions for reacting to inputs changing over time.

We start with the evaluation of bridge rules, for which we need to know current beliefs and current external information. The former is captured by the notion of a *belief state* denoted by a tuple of belief sets – one for each context – similar as in previous work on multi-context systems.

Definition 4 (Belief State). Let $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$ be an rMCS. Then, a belief state for M is a tuple $\mathbf{B} = \langle B_1, \dots, B_n \rangle$ such that $B_i \in \text{BS}_i$, for each $i \in \{1, \dots, n\}$. We use Bel_M to denote the set of all belief states for M .

To also capture the current external information, we introduce the notion of an *input*.

Definition 5 (Input). Let $M = \langle \mathbf{C}, \langle \text{IL}_1, \dots, \text{IL}_k \rangle, \text{BR} \rangle$ be an rMCS. Then an input for M is a tuple $\mathbf{l} = \langle I_1, \dots, I_k \rangle$ such that $I_i \subseteq \text{IL}_i$, $i \in \{1, \dots, k\}$. The set of all inputs for M is denoted by Inp_M .

We are now ready to define when literals (in bridge rule bodies) are satisfied.

Definition 6 (Satisfaction of Literals). Let $M = \langle \mathbf{C}, \text{IL}, \text{BR} \rangle$ be an rMCS such that $\mathbf{C} = \langle C_1, \dots, C_n \rangle$ and $\text{IL} = \langle \text{IL}_1, \dots, \text{IL}_k \rangle$. Given an input $\mathbf{l} = \langle I_1, \dots, I_k \rangle$ for M and a belief state $\mathbf{B} = \langle B_1, \dots, B_n \rangle$ for M , we define the satisfaction of literals as:

- $\langle \mathbf{l}, \mathbf{B} \rangle \models a_\ell$ if a_ℓ is of the form $c:b$ and $b \in B_c$;
- $\langle \mathbf{l}, \mathbf{B} \rangle \models a_\ell$ if a_ℓ is of the form $s:b$ and $b \in I_s$;
- $\langle \mathbf{l}, \mathbf{B} \rangle \models \text{not } a_\ell$ if $\langle \mathbf{l}, \mathbf{B} \rangle \not\models a_\ell$.

Let r be a bridge rule for C_i over \mathbf{C} and IL . Then

- $\langle \mathbf{l}, \mathbf{B} \rangle \models \text{bd}(r)$ if $\langle \mathbf{l}, \mathbf{B} \rangle \models l$ for every $l \in \text{bd}(r)$.

If $\langle \mathbf{l}, \mathbf{B} \rangle \models \text{bd}(r)$, we say that r is *applicable under \mathbf{l} and \mathbf{B}* . The operations encoded in the heads of applicable bridge rules in an rMCS determine which knowledge base updates should take place. We collect them in two disjoint sets.

Definition 7 (Applicable Bridge Rules). Let $M = \langle \mathbf{C}, \text{IL}, \text{BR} \rangle$ be an rMCS such that $\mathbf{C} = \langle C_1, \dots, C_n \rangle$ and $\text{BR} = \langle \text{BR}_1, \dots, \text{BR}_n \rangle$. Given an input \mathbf{l} for M and a belief state \mathbf{B} for M , we define, for each $i \in \{1, \dots, n\}$, the sets

- $\text{app}_i^{\text{now}}(\mathbf{l}, \mathbf{B}) = \{\text{hd}(r) \mid r \in \text{BR}_i, \langle \mathbf{l}, \mathbf{B} \rangle \models \text{bd}(r), \text{hd}(r) \in \text{OP}_i\}$;
- $\text{app}_i^{\text{next}}(\mathbf{l}, \mathbf{B}) = \{\text{op} \mid r \in \text{BR}_i, \langle \mathbf{l}, \mathbf{B} \rangle \models \text{bd}(r), \text{hd}(r) = \text{next}(\text{op})\}$.

Intuitively, the operations in $\mathbf{app}_i^{now}(l, B)$ are used for non-persistent updates of the knowledge base that influence the semantics of an rMCS for a single point in time. The operations in $\mathbf{app}_i^{next}(l, B)$ on the other hand are used for changing knowledge bases over time. They are not used for computing the current semantics but are applied in the next point in time depending on the current semantics. This continuous change of knowledge bases over time is the reason why, unlike in previous work on MCSs, we do not consider knowledge bases as part of the contexts to which they are associated but store them in a separate configuration structure defined next.

Definition 8 (Configuration of Knowledge Bases). *Let $M = \langle C, lL, BR \rangle$ be an rMCS such that $C = \langle C_1, \dots, C_n \rangle$. A configuration of knowledge bases for M is a tuple $KB = \langle kb_1, \dots, kb_n \rangle$ such that $kb_i \in KB_i$, for each $i \in \{1, \dots, n\}$. We use Con_M to denote the set of all configurations of knowledge bases for M .*

The semantics of an rMCS for a single time instant is given by its *equilibria*.

Definition 9 (Equilibrium). *Let $M = \langle \langle C_1, \dots, C_n \rangle, lL, BR \rangle$ be an rMCS, $KB = \langle kb_1, \dots, kb_n \rangle$ a configuration of knowledge bases for M , and l an input for M . Then, a belief state $B = \langle B_1, \dots, B_n \rangle$ for M is an equilibrium of M given KB and l if, for each $i \in \{1, \dots, n\}$, we have that*

$$B_i \in \mathbf{acc}_i(kb'), \text{ where } kb' = \mathbf{mng}_i(\mathbf{app}_i^{now}(l, B), kb_i).$$

Example 4. *Consider rMCS M_{ex3} from Example 3 with the configuration of knowledge bases for M_{ex3} , $KB = \langle kb_{st} \rangle = \langle \emptyset \rangle$, representing that the stove is turned off, input $l = \langle \{\text{switch}\} \rangle$ for M_{ex3} and the belief state $B = \langle \emptyset \rangle$ for M_{ex3} . As both bridge rules in BR_{ex3} use the **next** operator, we have $\mathbf{app}_{st}^{now}(l, B) = \emptyset$ and consequently, following the definition of \mathbf{mng}_{st} in Example 2, kb_{st} remains unchanged, i.e., $\mathbf{mng}_{st}(\mathbf{app}_{st}^{now}(l, B), kb_{st}) = kb_{st}$. Thus, $\mathbf{acc}_{st}(\mathbf{mng}_{st}(\mathbf{app}_{st}^{now}(l, B), kb_{st})) = \{\emptyset\}$. Thus, B is an equilibrium of M_{ex3} given KB and l .*

Based on an equilibrium at the current time instant, we can compute an updated configuration of knowledge bases using the update function as introduced next.

Definition 10 (Update Function). *Let $M = \langle C, lL, BR \rangle$ be an rMCS such that $C = \langle C_1, \dots, C_n \rangle$, $KB = \langle kb_1, \dots, kb_n \rangle$ a configuration of knowledge bases for M , l an input for M , and B a belief state for M . Then, the update function for M is defined as $\mathbf{upd}_M(KB, l, B) = \langle kb'_1, \dots, kb'_n \rangle$, such that, for each $i \in \{1, \dots, n\}$, $kb'_i = \mathbf{mng}_i(\mathbf{app}_i^{next}(l, B), kb_i)$.*

We can finally show how an rMCS behaves in the presence of external information that changes over time. For this purpose, we assume that an rMCS receives data in a stream of inputs, i.e., an input for each time instant, and we represent individual time instants by natural numbers. These can be interpreted as logical time instants that do not necessarily represent specific physical time points, nor do we require that every pair of consecutive natural numbers represents equidistant physical time spans.

Definition 11 (Input Stream). *Let $M = \langle C, lL, BR \rangle$ be an rMCS such that $lL = \langle lL_1, \dots, lL_k \rangle$. An input stream for M (until τ) is a function $\mathcal{I} : [1.. \tau] \rightarrow \text{Inp}_M$ where $\tau \in \mathbb{N} \cup \{\infty\}$.*

We will omit the term “until τ ” whenever the limit of the stream is irrelevant. Clearly, an input stream for M until τ also fully determines an input stream for M until τ' for every $1 \leq \tau' < \tau$. For any input stream \mathcal{I} and $t \in [1..\tau]$, we will use \mathcal{I}^t to denote $\mathcal{I}(t)$, i.e., the input $\langle I_1, \dots, I_k \rangle$ for M at time t . We also term *stream* the restriction of an input stream \mathcal{I} to a single input language IL_i , i.e., a function $\mathcal{I}_i : [1..\tau] \rightarrow 2^{IL_i}$ that is fully determined by \mathcal{I} .

Note that \mathcal{I}^t encapsulates (input) data for every input language of M . Hence, we assume that, at every time instant, we receive input from every external source of the rMCS. This synchronicity is required since the evaluation of a bridge rule may depend on the availability of information from multiple streams. One possibility for modeling external sources that do not continuously provide information is setting \mathcal{I}_s^t to the empty set for representing a lack of input from the source with language IL_s at time t .

The semantics of an rMCS over time is given by its *equilibria streams* for a given initial configuration of knowledge bases and an input stream for the system.

Definition 12 (Equilibria Stream). *Let $M = \langle C, IL, BR \rangle$ be an rMCS, KB a configuration of knowledge bases for M , and \mathcal{I} an input stream for M until τ where $\tau \in \mathbb{N} \cup \{\infty\}$. Then, an equilibria stream of M given KB and \mathcal{I} is a function $\mathcal{B} : [1..\tau] \rightarrow \text{Bel}_M$ such that*

- \mathcal{B}^t is an equilibrium of M given \mathcal{KB}^t and \mathcal{I}^t , where \mathcal{KB}^t is inductively defined as

- $\mathcal{KB}^1 = \text{KB}$
- $\mathcal{KB}^{t+1} = \mathbf{upd}_M(\mathcal{KB}^t, \mathcal{I}^t, \mathcal{B}^t)$.

We will also refer to the function $\mathcal{KB} : [1..\tau] \rightarrow \text{Con}_M$ as the configurations stream of M given KB, \mathcal{I} , and \mathcal{B} .

Note that the limit τ of an equilibria stream is aligned with that of the given input stream. Following the definition, it is easy to see that if we have an equilibria stream \mathcal{B} of M given KB and \mathcal{I} , then the substream of \mathcal{B} of size τ' , with $\tau' \leq \tau$, is an equilibria stream of M given KB and \mathcal{I}' , where \mathcal{I}' is the substream of \mathcal{I} of size τ' . This implies that, conversely, each extension of the input stream can only lead to equilibria streams that extend those obtained given the original input stream.

Example 5. *Recall M_{ex3} , KB, and \mathcal{B} from Example 4, as well as an input stream \mathcal{I} until 3 with $\mathcal{I}^1 = \langle \{\text{switch}\} \rangle$, $\mathcal{I}^2 = \langle \emptyset \rangle$, and $\mathcal{I}^3 = \langle \{\text{switch}\} \rangle$. There is an equilibria stream \mathcal{B} of M_{ex3} given KB and \mathcal{I} . Note that the input \mathcal{I}^1 coincides with input 1 from Example 4. As \mathcal{B} is the only equilibrium of M_{ex3} given KB and 1, we have that $\mathcal{B}^1 = \mathcal{B}$.*

As we have $\mathbf{app}_{st}^{next}(1, \mathcal{B}) = \{\text{setPower}(\text{on})\}$, the update function provides the following configuration of knowledge bases for time instant 2 (with $\mathcal{KB}^1 = \text{KB}$):

$$\mathcal{KB}^2 = \mathbf{upd}_{M_{ex3}}(\mathcal{KB}^1, \mathcal{I}^1, \mathcal{B}^1) = \langle \mathbf{mng}_{st}(\mathbf{app}_{st}^{next}(1, \mathcal{B}), kb) \rangle = \langle \{\text{pw}\} \rangle.$$

Thus, switching the power state at time 1 leads to an updated knowledge base indicating that the stove is on at time 2. The table in Figure 1 summarizes the equilibria stream and the configurations stream given KB and \mathcal{I} .

t	\mathcal{KB}^t	\mathcal{I}^t	\mathcal{B}^t	$\mathbf{app}_{st}^{next}(\mathcal{I}^t, \mathcal{B}^t)$
1	$\langle \emptyset \rangle$	$\langle \{\text{switch}\} \rangle$	$\langle \emptyset \rangle$	$\{\text{setPower}(on)\}$
2	$\langle \{\text{pw}\} \rangle$	$\langle \emptyset \rangle$	$\langle \{\text{pw}\} \rangle$	\emptyset
3	$\langle \{\text{pw}\} \rangle$	$\langle \{\text{switch}\} \rangle$	$\langle \{\text{pw}\} \rangle$	$\{\text{setPower}(off)\}$

Figure 1: Streams and applicable operations for M_{ex3}

2.3. Rule Schemata

For convenience, we will represent sets of bridge rules using rule schemata. Intuitively, a rule schema is a parametrized bridge rule and each instance is a bridge rule. It will be convenient to use not only parameters, but also conditions to additionally constrain the set of instances. We will use them, e.g., for allowing arithmetic operations and comparison relations in bridge rules.

Given a tuple $\mathbf{C} = \langle C_1, \dots, C_n \rangle$ of contexts and a tuple \mathbf{IL} of input languages, let \mathcal{A} be the alphabet of symbols occurring in all possible bridge rules for all contexts $C_i \in \{C_1, \dots, C_n\}$ over \mathbf{C} and \mathbf{IL} . We call a string over \mathcal{A} an *instantiation term* for \mathbf{C} and \mathbf{IL} . Then, an *instantiation condition* for \mathbf{C} and \mathbf{IL} is a predicate $id(T_1, \dots, T_o)$, where T_1, \dots, T_o are strings over $P \cup \mathcal{A}$ and where $P = \{P_1, \dots, P_k\}$ is a set of symbols (called *parameters*) such that $P \cap \mathcal{A} = \emptyset$. A *rule schema* R for \mathbf{C} and \mathbf{IL} with parameters P is of the form

$$H \leftarrow A_1, \dots, A_j, \mathbf{not} A_{j+1}, \dots, \mathbf{not} A_m, D_1, \dots, D_s \quad (2)$$

such that H, A_1, \dots, A_m are strings over $P \cup \mathcal{A}$ and D_i is an instantiation condition for each $i \in \{1, \dots, s\}$.

A bridge rule $r = \mathbf{op} \leftarrow a_1, \dots, a_j, \mathbf{not} a_{j+1}, \dots, \mathbf{not} a_m$ for C_i over \mathbf{C} and \mathbf{IL} is an *instance* of a rule schema R for \mathbf{C} and \mathbf{IL} with parameters P of the form (2) if r results from $H \leftarrow A_1, \dots, A_j, \mathbf{not} A_{j+1}, \dots, \mathbf{not} A_m$ by a uniform substitution σ of parameters with instantiation terms, and for each $D_i = id(T_1, \dots, T_o)$, $i \in \{1, \dots, s\}$, the predicate $id(T_1\sigma, \dots, T_o\sigma)$ holds.

We adopt the convention from logic programming that parameters start with uppercase letters. Moreover, we will denote instantiation conditions representing comparison and arithmetic operations using standard infix notation. For example, the rule schema

$$\mathbf{add}(tmp(S, X)) \leftarrow S::temp = X, 45 \leq X \leq 65. \quad (3)$$

expresses that temperature values received on stream S should be added to the knowledge base (tagged with the stream index) if they lie between 45°C and 65°C.

Note that we assume that instantiation conditions only evaluate to true if their arguments are substituted by appropriate instantiation terms, for example when X is substituted by numeric values in the temperature example.

2.4. Example Scenario

So far, in our examples we have only shown rMCSs with a single context, to ease the presentation. We conclude this section with an extensive example of the assisted-living scenario thereby illustrating the usage of an rMCS with several contexts (and input languages) and the flow of information within these.

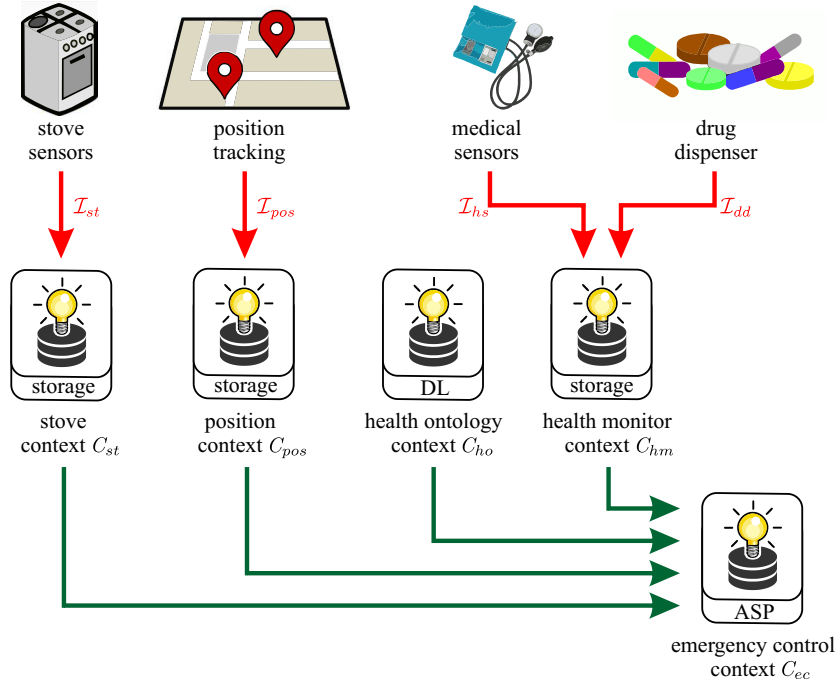


Figure 2: Structure of M_{al} .

The rMCS for our assisted living scenario, $M_{al} = \langle C, IL, BR \rangle$, as illustrated in Figure 2, with $C = \langle C_{st}, C_{pos}, C_{hm}, C_{ho}, C_{ec} \rangle$ has five contexts: C_{st} is the stove context, C_{pos} keeps track of the current position of Dave, C_{hm} monitors Dave's health status, C_{ho} is a medical ontology that the system may consult, and C_{ec} is the emergency control context that is responsible for detecting potential threats and taking actions, such as raising an alarm. The five contexts are connected by their corresponding sets of bridge rules $BR = \langle BR_{st}, BR_{pos}, BR_{hm}, BR_{ho}, BR_{ec} \rangle$, and four input languages $IL = \langle IL_{st}, IL_{pos}, IL_{hs}, IL_{dd} \rangle$ represent the sensors used in the assisted living scenario.

The stove context C_{st} has already been described in Examples 2 and 3. We adjust $IL_{st} = IL_{ex3} \cup \{tmp(T) \mid T \in \mathbb{Z}\}$ to allow for temperature sensor input data in IL_{st} and by adding further bridge rules for context C_{st} , i.e., BR_{st} extends BR_{ex3} (with st replacing $ex3$ in input literals) by the instances of the following bridge rule schemata

$$\begin{aligned} \text{setTemp}(cold) &\leftarrow st::tmp(T), T \leq 45. \\ \text{setTemp}(hot) &\leftarrow st::tmp(T), 45 < T. \end{aligned}$$

that classify the current temperature value as cold or hot. These bridge rules do not use the `next` operator and hence influence the computation of equilibria rather than the update of knowledge bases for the next step, i.e., temperature information does not persist in the knowledge base, but continuously computed from the input stream.

The input language IL_{pos} is given by

$$\{\text{enters}(\textit{kitchen}), \text{enters}(\textit{bathroom}), \text{enters}(\textit{bedroom})\}$$

and non-empty sensor (stream) input for I_{pos} signals when Dave changes rooms. This information is used in context C_{pos} to keep track of Dave's position. This context also uses a storage logic with $E = \{\text{pos}(P) \mid P \in \{\textit{kitchen}, \textit{bathroom}, \textit{bedroom}\}\}$, and its corresponding set of bridge rules BR_{pos} is given by the schema

$$\begin{aligned} \text{next}(\text{setPos}(P)) &\leftarrow \text{pos}::\text{enters}(P). \\ \text{setPos}(P) &\leftarrow \text{pos}::\text{enters}(P). \end{aligned}$$

The management function writes Dave's new position into the knowledge base whenever he changes rooms and keeps the previous position, otherwise, which is why we have a bridge rule with the operator **next**, to ensure that, as for $\text{setPower}(P)$ for C_{st} , the current position is stored. In addition, here we also have an identical copy of the bridge rule without **next** with the rationale that the information is available immediately. Intuitively, this will allow us to avoid situations in which the conditions for an alarm would be met including the absence of Dave, although he is just entering the room in question.

The input language IL_{hs} is used to provide stream data on the health sensor readings from Dave while IL_{dd} serves to provide data on the (usage of the) drug dispenser. The information from both sensors is used in context C_{hm} which serves as a health monitor and also builds on a storage logic with

$$E = \{\text{status}(\textit{asleep}), \text{m}(\textit{drugA}), \text{bp}(\textit{high}), \text{bp}(\textit{normal})\}.$$

Here, IL_{hs} contains $\{\textit{asleep}\} \cup \{\text{bpReading}(R) \mid R \text{ a blood pressure value}\}$ while IL_{dd} contains $\{\text{dispensed}(\textit{drugA})\}$ and the bridge rules in BR_{hm} are given by:

$$\begin{aligned} \text{setStatus}(\textit{asleep}) &\leftarrow \text{hs}::\textit{asleep}. \\ \text{next}(\text{setBP}(\textit{high})) &\leftarrow \text{hs}::\text{bpReading}(R), R \geq 140/90. \\ \text{next}(\text{setBP}(\textit{normal})) &\leftarrow \text{hs}::\text{bpReading}(R), R < 140/90. \\ \text{next}(\text{setMed}(\text{m}(\textit{drugA}))) &\leftarrow \text{dd}::\text{dispensed}(\textit{drugA}). \end{aligned}$$

While the sleep status can be monitored permanently and new data arrives at every time instant, blood pressure measurements and taking medications are only taken occasionally, so this data needs to be stored until new information arrives. Here, we abstract from the details on the duration of effects of medications. The management function $\text{mng}_{hm}(OP, kb)$ can then be defined quite similarly to $\text{mng}_{st}(OP, kb)$.

The context $C_{ho} = \langle L_{ho}, OP_{ho}, \text{mng}_{ho} \rangle$ represents a biomedical health ontology, which uses the DL logic described in Example 1. It contains information on diseases, treatments, and so on, and, it can, e.g., be used to consult for possible effects of medications by context C_{ec} as explained below. In our case, kb_{ho} is simplified to

$$\begin{aligned} \textit{drugA} &: \exists \text{contains.ephedrine} \\ \exists \text{contains.ephedrine} &\sqsubseteq \text{causesHighBP} \end{aligned}$$

which allows one to derive that *drugA* causes high blood pressure. In our simple setting, the management function \mathbf{mng}_{ho} leaves the knowledge base kb_{ho} unchanged.

$C_{ec} = \langle L_{ec}, OP_{ec}, \mathbf{mng}_{ec} \rangle$ is the context for detecting emergencies. It is implemented as an answer set program, hence the acceptable belief sets of L_{ec} are the answer sets of its knowledge bases as outlined in Example 1. The bridge rules in BR_{ec} do not refer to stream input but query other contexts:

$$\begin{aligned} \text{add}(\text{oven}(on, hot)) &\leftarrow st:pw, st:tm(hot). \\ \text{add}(\text{humanPos}(P)) &\leftarrow pos:pos(P). \\ \text{add}(\text{status}(asleep)) &\leftarrow hm:asleep. \\ \text{add}(\text{highBP}) &\leftarrow hm:bp(high). \\ \text{add}(\text{highBPMed}) &\leftarrow ho:causesHighBP(D), hm:m(D). \end{aligned}$$

Note that in the first bridge rule, we only import information for the case that will possibly be problematic, namely if the oven is turned on and hot, whereas, e.g., in the second, we use a schema to keep track of the position in all cases. The answer set program kb_{ec} is given by the rules:

$$\begin{aligned} \text{turnOff}(stove) &\leftarrow \text{oven}(on, hot), \text{status}(asleep). \\ \text{alert}(stove) &\leftarrow \text{oven}(on, hot), \mathbf{not} \text{humanPos}(kitchen), \mathbf{not} \text{status}(asleep). \\ \text{call}(medAssist) &\leftarrow \text{highBP}, \mathbf{not} \text{highBPMed}. \end{aligned}$$

The management function of C_{ec} , $\mathbf{mng}_{ec}(OP, kb)$, that adds information from the bridge rules temporarily as input facts to the context's knowledge base can be defined similar to the previous contexts.

An example run of this entire rMCS is given in Figure 3. Note that we have omitted from \mathcal{KB}^t and \mathcal{B}^t all information which is fixed to ease the reading, which is why, e.g., KB_{ho} and KB_{ec} in the figure are always empty. The presented situation begins with Dave in the kitchen, where he turns on the stove at $t = 1$. The effect of this is visible at $t = 2$, where he takes his medication (also in the kitchen). Note that blood pressure readings are not constantly done, and therefore at $t = 2$ no reading occurs. Subsequently, Dave leaves for the bedroom to rest (thus forgetting about the stove), and this change of position is available right away and stored. The next blood pressure reading at $t = 4$ is high, which is stored in the corresponding knowledge base from $t = 5$ onwards, at which time point a medical assistant would be called if the medication most likely causing this was not registered. Finally, at $t = 6$, the oven has become hot, but since Dave has fallen asleep, the context for detecting emergencies initiates turning off the stove to prevent possible problems. If Dave was not asleep yet, an alarm would have been sounded reminding him of the stove instead.

Finally, note that in our example scenario, currently the effects of taken medications do not wear off. Such extension can be handled by e.g., incorporating explicit time, one of the modeling features of rMCSs described in the next section.

3. Modeling with rMCSs

The running example demonstrates that rMCSs can jointly integrate different formalisms and deal with a dynamic environment. In this section, we want to focus on

t	$K\mathcal{B}^t$	\mathcal{I}^t	$\mathbf{app}_i^{now}(\mathcal{I}^t, \mathcal{B}^t)$	\mathcal{B}^t	$\mathbf{app}_i^{next}(\mathcal{I}^t, \mathcal{B}^t)$
1	$\langle \emptyset, \{\text{pos}(\text{kitchen})\}, \{\text{bp}(\text{normal})\}, \emptyset, \emptyset \rangle$	$\langle \{\text{tmp}(19), \text{switch}\}, \emptyset, \{\text{bpReading}(135/86)\}, \emptyset \rangle$	$\langle \{\text{setTemp}(\text{cold})\}, \emptyset, \emptyset, \emptyset, \{\text{add}(\text{humanPos}(\text{kitchen}))\} \rangle$	$\langle \{\text{tm}(\text{cold})\}, \{\text{pos}(\text{kitchen})\}, \{\text{bp}(\text{normal})\}, \emptyset, \{\text{humanPos}(\text{kitchen})\} \rangle$	$\langle \{\text{setPower}(\text{on})\}, \emptyset, \{\text{setBP}(\text{normal})\}, \emptyset, \emptyset \rangle$
2	$\langle \{\text{pw}\}, \{\text{pos}(\text{kitchen})\}, \{\text{bp}(\text{normal})\}, \emptyset, \emptyset \rangle$	$\langle \{\text{tmp}(21)\}, \emptyset, \emptyset, \{\text{dispensed}(\text{drugA})\} \rangle$	$\langle \{\text{setTemp}(\text{cold})\}, \emptyset, \emptyset, \emptyset, \{\text{add}(\text{humanPos}(\text{kitchen}))\} \rangle$	$\langle \{\text{tm}(\text{cold}), \text{pw}\}, \{\text{pos}(\text{kitchen})\}, \{\text{bp}(\text{normal})\}, \emptyset, \{\text{humanPos}(\text{kitchen})\} \rangle$	$\langle \emptyset, \emptyset, \{\text{setMed}(\text{m}(\text{drugA}))\}, \emptyset, \emptyset \rangle$
3	$\langle \{\text{pw}\}, \{\text{pos}(\text{kitchen})\}, \{\text{bp}(\text{normal}), \text{m}(\text{drugA})\}, \emptyset, \emptyset \rangle$	$\langle \{\text{tmp}(27)\}, \{\text{enters}(\text{bedroom})\}, \{\text{bpReading}(138/89)\}, \emptyset \rangle$	$\langle \{\text{setTemp}(\text{cold})\}, \{\text{setPos}(\text{bedroom})\}, \emptyset, \emptyset, \{\text{add}(\text{humanPos}(\text{bedroom})), \text{add}(\text{highBPMed})\} \rangle$	$\langle \{\text{tm}(\text{cold}), \text{pw}\}, \{\text{pos}(\text{bedroom})\}, \{\text{bp}(\text{normal})\}, \emptyset, \{\text{humanPos}(\text{bedroom}), \text{highBPMed}\} \rangle$	$\langle \emptyset, \{\text{setPos}(\text{bedroom})\}, \{\text{setBP}(\text{normal})\}, \emptyset, \emptyset \rangle$
4	$\langle \{\text{pw}\}, \{\text{pos}(\text{bedroom})\}, \{\text{bp}(\text{normal}), \text{m}(\text{drugA})\}, \emptyset, \emptyset \rangle$	$\langle \{\text{tmp}(36)\}, \emptyset, \{\text{bpReading}(148/97)\}, \emptyset \rangle$	$\langle \{\text{setTemp}(\text{cold})\}, \emptyset, \emptyset, \emptyset, \{\text{add}(\text{humanPos}(\text{bedroom})), \text{add}(\text{highBPMed})\} \rangle$	$\langle \{\text{tm}(\text{cold}), \text{pw}\}, \{\text{pos}(\text{bedroom})\}, \{\text{bp}(\text{normal})\}, \emptyset, \{\text{humanPos}(\text{bedroom}), \text{highBPMed}\} \rangle$	$\langle \emptyset, \emptyset, \{\text{setBP}(\text{high})\}, \emptyset, \emptyset \rangle$
5	$\langle \{\text{pw}\}, \{\text{pos}(\text{bedroom})\}, \{\text{bp}(\text{high}), \text{m}(\text{drugA})\}, \emptyset, \emptyset \rangle$	$\langle \{\text{tmp}(43)\}, \emptyset, \{\text{bpReading}(146/95)\}, \emptyset \rangle$	$\langle \{\text{setTemp}(\text{cold})\}, \emptyset, \emptyset, \emptyset, \{\text{add}(\text{humanPos}(\text{bedroom})), \text{add}(\text{highBPMed}), \text{add}(\text{highBP})\} \rangle$	$\langle \{\text{tm}(\text{cold}), \text{pw}\}, \{\text{pos}(\text{bedroom})\}, \{\text{bp}(\text{high})\}, \emptyset, \{\text{humanPos}(\text{bedroom}), \text{highBPMed}, \text{highBP}\} \rangle$	$\langle \emptyset, \emptyset, \{\text{setBP}(\text{high})\}, \emptyset, \emptyset \rangle$
6	$\langle \{\text{pw}\}, \{\text{pos}(\text{bedroom})\}, \{\text{bp}(\text{high}), \text{m}(\text{drugA})\}, \emptyset, \emptyset \rangle$	$\langle \{\text{tmp}(51)\}, \emptyset, \{\text{asleep}\}, \emptyset \rangle$	$\langle \{\text{setTemp}(\text{hot})\}, \emptyset, \{\text{setStatus}(\text{asleep})\}, \emptyset, \{\text{add}(\text{humanPos}(\text{bedroom})), \text{add}(\text{highBPMed}), \text{add}(\text{highBP}), \text{add}(\text{oven}(\text{on}, \text{hot})), \text{add}(\text{status}(\text{asleep}))\} \rangle$	$\langle \{\text{tm}(\text{hot}), \text{pw}\}, \{\text{pos}(\text{bedroom})\}, \{\text{bp}(\text{high})\}, \emptyset, \{\text{humanPos}(\text{bedroom}), \text{highBPMed}, \text{highBP}, \text{oven}(\text{on}, \text{hot}), \text{status}(\text{asleep}), \text{turnOff}(\text{stone})\} \rangle$	$\langle \emptyset, \emptyset, \emptyset, \emptyset, \emptyset \rangle$

Figure 3: Streams and applicable operations for the example scenario given by M_{ol}

aspects relevant to the knowledge engineer, namely how to model a certain scenario using an rMCS. To this end, we elaborate on different generic modeling techniques for rMCSs that we consider helpful in typical target applications. For concrete implementations, these techniques can still be refined and tailored towards the specific needs of the problem domain at hand. First, in Section 3.1, we discuss bridge rules as a device for translating stream data to a knowledge base language. When to use the `next` operator in the head of bridge rules is addressed in Section 3.2. In Section 3.3, we discuss how we deal with time on the object level, including the use of timestamps as well as external and logical clocks. Then, in Section 3.4, another technique is presented that allows for managing conflicting data in input streams on the modeling level, e.g., contradictory sensor measurements. An important issue for systems that are continuously online is when to keep data and when to remove it. In Section 3.5, we discuss how to do this and provide techniques for dynamically adjusting what information is kept in the system. These techniques allow us, e.g., to modify the size of sliding windows for stream data depending on the current situation. While the computation of acceptable belief sets may be easy for some contexts, it might be expensive for others that have to perform complex reasoning. In practice, it will therefore be wise to only evaluate these contexts if necessary.

3.1. Incorporating Stream Data

Bridge rules are responsible for providing a context with information from streams and other contexts. As we deal with heterogeneous context languages and since input languages may differ from context languages, one important aspect of rMCSs (and earlier types of MCSs) is that bridge rules can be seen as a translation device between the different languages: bridge rule bodies use the languages of the source contexts or streams whereas a bridge rule head is an operation that produces a knowledge base in the target language (via the management function). But bridge rules do not necessarily pass information the way it is used in the body. Rather, we can model bridge rules such that we they focus on information relevant to the target context and, e.g., translate sensor data into a form that is convenient for reasoning. Consider the bridge rule schemas shown in Section 2.4 for C_{st} :

$$\begin{aligned} \text{setTemp}(\text{cold}) &\leftarrow \text{ex3::tmp}(T), T \leq 45. \\ \text{setTemp}(\text{hot}) &\leftarrow \text{ex3::tmp}(T), 45 < T. \end{aligned}$$

In this case, the bridge rules only communicate whether the stove is hot or cold, abstracting away the exact temperature value coming from the sensor. This is in contrast to rule schema (3) discussed in the previous section where parameter X appears in the rule head whereby a concrete temperature value is added to the knowledge base.

3.2. Operational and Declarative Bridge Rules

For gaining a better understanding of when to use `next` in the head of a bridge rule, reconsider the example of turning a switch on and off as in Examples 2 and 3. Remember that the bridge rules

$$\begin{aligned} \text{next}(\text{setPower}(\text{on})) &\leftarrow \text{ex3::switch}, \text{not } \text{st:pw}. \\ \text{next}(\text{setPower}(\text{off})) &\leftarrow \text{ex3::switch}, \text{st:pw}. \end{aligned}$$

were used so that an occurrence of an input atom switch causes the knowledge base to contain an inverted power state at the next time point. In order to highlight the difference, assume we would instead use the bridge rules

$$\begin{aligned} \text{setPower}(on) &\leftarrow ex3::\text{switch}, \mathbf{not} \text{ } st:\text{pw}. \\ \text{setPower}(off) &\leftarrow ex3::\text{switch}, st:\text{pw}. \end{aligned}$$

without **next**. We refer to the version of M_{ex3} from Example 3 with these modified bridge rules by M'_{ex3} . Consider the configuration $\text{KB} = \langle kb_{st} \rangle$ of knowledge bases with $kb_{st} = \emptyset$ and the input $I = \langle \{\text{switch}\} \rangle$ as in Example 4. Indeed, M'_{ex3} has no equilibrium given KB and I : If we take belief state $B = \langle \emptyset \rangle$ as in Example 4, we have $\mathbf{app}_{st}^{now}(I, B) = \{\text{setPower}(on)\}$ (instead of \emptyset as in Example 4). Consequently, following the definition of \mathbf{mng}_{st} in Example 2, we get $\mathbf{mng}_{st}(\mathbf{app}_{st}^{now}(I, B), kb_{st}) = \{\text{pw}(on)\}$. But then B is not contained in $\mathbf{acc}_{st}(\mathbf{mng}_{st}(\mathbf{app}_{st}^{now}(I, B), kb_{st})) = \{\{\text{pw}(on)\}\}$, thus, B is not an equilibrium of M'_{ex3} given KB and I . This is in line with the intuition that without **next**, turning the switch should affect the current equilibrium. However, $B' = \langle \{\text{pw}(on)\} \rangle$ is also not an equilibrium of M'_{ex3} . Since $\mathbf{app}_{st}^{now}(I, B') = \{\text{setPower}(off)\}$, we get $\mathbf{mng}_{st}(\mathbf{app}_{st}^{now}(I, B'), kb_{st}) = \emptyset$. But then $\mathbf{acc}_{st}(\mathbf{mng}_{st}(\mathbf{app}_{st}^{now}(I, B'), kb_{st})) = \{\emptyset\}$ does not contain B' , consequently also B' is not an equilibrium of M'_{ex3} given KB and I . The two bridge rules without **next** prevent stability of a belief state required for an equilibrium: believing that the power is on (respectively off) causes an update of the knowledge base that the power is off (respectively on) which is in turn inconsistent to the belief.

The reason why the change considered here does not work is that the two types of bridge rule are meant to be used for different purposes. The bridge rules using **next** are responsible for changing knowledge bases over time. An rMCS without such bridge rules cannot alter its initial configuration of knowledge bases. Thus, these rules come with an *operational* flavor. Bridge rules without **next** on the other hand have a *pure declarative* nature and their purpose is to semantically integrate the contexts of an rMCS.

Still, as we have seen for M_{al} in Section 2.4, it sometimes makes sense to use essentially the same bridge rules with and without **next**: the bridge rule

$$\mathbf{next}(\text{setPos}(P)) \leftarrow pos::\text{enters}(P).$$

ensures that the information about the position of Dave persists in the knowledge base, whereas

$$\text{setPos}(P) \leftarrow pos::\text{enters}(P).$$

provides this information for computing equilibria in the current time instant.

3.3. Integration of Time

The structure of input streams and equilibria streams implicitly induces a discrete logical time for rMCSs. In order to operate in dynamic environments, in many cases it is necessary to deal with explicit physical time or logical time, and both can be achieved on the level of modeling. To this end, it is necessary to have access to explicit

time points on this level, i.e., in the bridge rules and knowledge bases. A natural way to make such time points explicitly available is the use of an external clock that provides the current time via an input stream \mathcal{I}_c . Thus, every input for the rMCS contains information about the current time which can then be queried in bridge rules.

Consider a context C_1 that requires information about the development of temperature values from an incoming sensor over time. The bridge rule schema

$$\text{next}(\text{add}(\text{tmpAtTime}(\text{Temp}, T))) \leftarrow \text{tmp}::\text{tmp}(\text{Temp}), c::\text{now}(T).$$

can be used to add atoms of form $\text{tmpAtTime}(\text{Temp}, T)$ to the context, where Temp is the current temperature and T stands for the time of the sensor reading. This setting also allows for querying, e.g., whether a temperature of 45°C was exceeded within the last 10 minutes, as expressed in the following bridge rule schema:

$$\begin{aligned} \text{add}(\text{recentlyHot}) &\leftarrow 1:\text{tmpAtTime}(\text{Temp}, T'), \\ &\text{Temp} > 45, c::\text{now}(T), T' \geq T - 10. \end{aligned}$$

Another possibility is to use logical time synchronized with the length of the equilibria stream, e.g., whenever the use of an external clock is not an option. We can obtain timestamps from the computation of knowledge base updates by using one context C_{clock} that keeps information about the current logical time that uses the following bridge rule schemas and whose initial knowledge base is assumed to be empty:

$$\begin{aligned} \text{setTime}(\text{now}(0)) &\leftarrow \text{not } \text{clock}::\text{timeAvailable}. \\ \text{next}(\text{add}(\text{timeAvailable})) &\leftarrow \text{clock}::\text{now}(0). \\ \text{next}(\text{setTime}(\text{now}(T + 1))) &\leftarrow \text{clock}::\text{now}(T). \end{aligned}$$

The first rule is used for initialization ensuring that if no time information is yet available, the logical time is set to the value 0. The third rule increments the current time by one and stores the updated value in the knowledge base of the next time instant. Finally, the second rule, ensures that once this value is greater than 0, the first rule can no longer be applied.

3.4. Handling Inconsistent Stream Data

In many situations, inconsistencies may occur when dealing with multiple external sources of data. Consider, e.g., an array of sensors that measure interdependent properties among which there may be a sensor that provides more fine-grained results, but which is also less reliable, so that it sometimes provides values that conflict with the data from the remaining sensors. In this case, we would like to use such a measure of reliability to consistently accommodate relevant sensor data. Next, we present a technique for integrating possibly inconsistent stream data into a context of an rMCS. Let $M = \langle C, \text{IL}, \text{BR} \rangle$ be an rMCS with $\text{IL} = \langle \text{IL}_1, \dots, \text{IL}_k \rangle$, and $C_i \in C$ a context whose aim is to receive and consistently accommodate (potentially conflicting) information

from the streams. To deal with possible inconsistencies, C_i has bridge rules of the form

$$\text{addC}(D, j) \leftarrow j::D.$$

for $j \in \{1, \dots, k\}$, where the operation addC is meant to consistently add the information of sensor j to the context. To address possible inconsistencies among sensors, we foresee a management function \mathbf{mng}_i that operates based on a total preference relation \prec over the available sensors. The second argument of the addC operation provides information about the source of a piece of information and thus a way of imposing preferences on the information to be added. Without loss of generality assume $IL_1 > \dots > IL_k$, that is, input language IL_1 has highest priority. Moreover, a notion of consistency needs to be specific to the context and so we assume a property $\text{cons}(kb)$ that holds if the knowledge base kb is consistent (according to such domain specific notion of consistency).

Given a set of operations OP , we define the sets of input data from each sensor $\text{inp}'_j = \{d \mid \text{addC}(d, j) \in OP\}$ for $j \in \{1, \dots, k\}$.

We then assume that $\text{inp}_0^c(OP) = \emptyset$ and let

$$\text{inp}_j^c(OP) = \begin{cases} \text{inp}_{j-1}^c(OP) \cup \text{inp}'_j & \text{if } \text{cons}(\text{inp}_{j-1}^c(OP) \cup \text{inp}'_j) \\ \text{inp}_{j-1}^c(OP) & \text{otherwise.} \end{cases}$$

Finally, we define $\mathbf{mng}_i(OP, kb_i) = kb_i \cup \text{inp}_k^c(OP)$.

The intuitive idea is that, starting with the input stream with highest priority, data from each sensor is only incorporated into the knowledge base if such data is consistent with the data already collected from the input streams with higher priority.

Note that by considering $\text{inp}_0^c(OP) = \emptyset$, the solution only considers inconsistency of data on the streams. For considering inconsistency between the streams and also the knowledge base kb_i of the context, we can set $\text{inp}_0^c(OP) = kb_i$.

We can also easily consider the incorporation of meta-information about sensors whose readings are considered inconsistent. This only requires a small change in the definition of $\text{inp}_j^c(OP)$ to $\text{inp}_{j-1}^c(OP) \cup \{\text{incons}(j)\}$ in case a conflict occurs. Such meta information can then be further leveraged by, e.g., initiating a control of the sensor if such measurements fall outside of the expected parameters of such sensor.

In all, this shows how the management function can solve conflicts due to inconsistent stream data based on preferences among the streams. Of course, many more strategies for integrating inconsistent stream data can be thought of. For example, in absence of a global ranking between streams, one way to ensure consistency is to select maximally consistent subsets of stream data. A corresponding management function could then be defined such that $\mathbf{mng}_i(OP, kb_i) = kb_i \cup \text{inp}_{mx}$, where inp_{mx} is a maximal set where $\text{inp}_{mx} \subseteq \{d \mid \text{addC}(d, j) \in OP, j \in \{1, \dots, m\}\}$ and $\text{cons}(\text{inp}_{mx})$ holds.

The strategies above show how we can deal with contradictory information in the processed data by means of modeling. In Section 4, we address inconsistency on the level of the formalism caused by nonexistence of equilibria or inconsistent belief states.

3.5. Selective Forgetting and Data Retention

As argued in Section 3.3 for the example where we were constantly adding sensor information to some knowledge base, sometimes it is necessary to forget (part of this)

knowledge again. In this section, we show how rMCSs can model scenarios where there is a need to dynamically adjust the size of the stored stream history. We do that by considering an extension of our running example. Recall from M_{al} in Section 2.4 that context C_{ec} is a context for detecting emergencies. In a situation where the stove is hot and Dave is asleep, $\text{turnOff}(\text{stove})$ is derived, signaling that the stove must be turned off. In case the stove is hot, but Dave is neither asleep nor in the kitchen, then an alarm is raised, signaling a potential emergency. In the latter case, we do not want to immediately turn the stove off, since it may well be the case that the absence of Dave from the kitchen is short. A situation is only considered a real emergency if Dave is absent from the kitchen for a (predefined) long period of time. To model such situation, we use a context C_{stE} , which collects timestamped alerts raised by context C_{ec} and uses these to check if a real stove-related emergency has occurred. Since we need to reason about time as in Section 3.3, we consider an input stream \mathcal{I}_c that provides the current time. The possible knowledge bases of C_{stE} contain elements of the form $\text{alert}(\text{stove}, t)$ where $t \in \mathbb{N}$, one element of the form $\text{winE}(t)$, which defines the limit size of the time window between two alerts above which an emergency should be raised, and possibly one element of the form $\text{emergency}(\text{stove})$ that signals the existence of a stove-related emergency. The set of bridge rules of C_{stE} is determined by the following bridge rule schemas:

$$\begin{aligned}
\text{next}(\text{add}(\text{alert}(\text{stove}, T))) &\leftarrow c::\text{now}(T), ec:\text{alert}(\text{stove}). \\
\text{next}(\text{del}(\text{alert}(\text{stove}, T))) &\leftarrow stE:\text{alert}(\text{stove}, T), \mathbf{not} \ ec:\text{alert}(\text{stove}). \\
\text{add}(\text{emergency}(\text{stove})) &\leftarrow c::\text{now}(T), ec:\text{alert}(\text{stove}), stE:\text{alert}(\text{stove}, T'), \\
&\quad stE:\text{winE}(Y), |T - T'| \geq Y.
\end{aligned}$$

The first rule adds a timestamped stove alert whenever such alert is active on context C_{ec} . The second rule removes all currently stored stove alerts whenever no such alert is coming from context C_{ec} . This guarantees that the knowledge base of C_{stE} does not accumulate unnecessary information. The last bridge rule triggers a real emergency alert whenever, in the history of alerts kept in the knowledge base of C_{stE} , there is an alert whose timestamp differs from the current time more than the acceptable emergency window.

Using context C_{stE} , we have shown how to model scenarios where tracking the stream history is triggered by alerts of possible emergencies. We now also consider the case where such alerts trigger the change of the window size of stream history to be kept. Consider a scenario with several potential emergencies, which can be just suspected or confirmed. Based on the status of the emergencies at each time point, we may need to adapt the size of the stream history that is kept. We generically model such scenario with an rMCS with a context C_d , which is used for emergency detection in such a dynamic environment, and an input language IL_s , which represents the possible observations. Assume there are m potential emergencies e_1, \dots, e_m we want the context to handle. The role of C_d is to check, based on the observations made, whether one or more of the emergencies e_i are suspected or confirmed. Based on information about potential emergencies, C_d adjusts the time window of the observations that are kept. This is the basis for intelligent forgetting based on dynamic windows.

The only assumptions we make about how C_d works internally are:

- C_d may signal that emergency e_i is suspected ($\text{susp}(e_i)$) or confirmed ($\text{conf}(e_i)$).
- C_d has information for each different observation p about default, respectively actual window sizes, $\text{defWin}(p, w)$, $\text{win}(p, w)$, and
- C_d has information about the window size for each observation relevant for a particular emergency, $\text{rel}(p, e, w)$.

The set of bridge rules for C_d includes the following rules.

$$\begin{aligned} \text{next}(\text{set}(\text{win}(P, X))) &\leftarrow d:\text{defWin}(P, X), \text{not } d:\text{susp}(E). \\ \text{next}(\text{set}(\text{win}(P, Y))) &\leftarrow d:\text{rel}(P, E, Y), d:\text{susp}(E). \\ \text{alarm}(E) &\leftarrow d:\text{conf}(E). \end{aligned}$$

The operation set sets the window size to a new value, deleting the old one, while alarm is an operation that adds information to the context knowledge base signaling an alarm. Since an observation can be relevant for more than one emergency, it may be the case that the management function has to deal with operations $\text{set}(\text{win}(p, w))$ with the same p but with different values of w . In that case, in order to avoid losing observations relevant for some suspected emergency, the management function takes the largest value of w as the window size for p .

Finally, the following bridge rule schemas define addition and deletion of observations from some stream s . The deletion of observations are performed in accordance with the determined window sizes:

$$\begin{aligned} \text{next}(\text{add}(P(T))) &\leftarrow t::\text{now}(T), s::P. \\ \text{next}(\text{del}(P(T'))) &\leftarrow d:P(T'), t::\text{now}(T), d:\text{win}(P, Z), T' < T - Z. \end{aligned}$$

The management function just performs additions and deletions on the context knowledge base. Since additions always include the (current) time of addition, deletions always refer to an earlier point in time, thus these two operators can never occur simultaneously.

4. Inconsistency Management

The occurrence of inconsistencies within frameworks that aim at integrating knowledge from different sources cannot be neglected, even more so in dynamic settings where knowledge changes over time. There are many reasons why rMCSs may fail to have an equilibria stream. These include the absence of an acceptable belief set for one of its contexts given its current knowledge base at some point in time, some occurring conflict between the operations in the heads of bridge rules, or simply because the input stream is such that the configuration of the flow of information within the rMCS, namely its bridge rules, prevent the existence of such an equilibria stream. In a real world situation, an rMCS without an equilibria stream is essentially useless. Not only can it not be used at the first time point equilibria ceased to exist, but it also cannot recover, even if what caused the problem was the particular input at that time point, which is bound to subsequently change into some other input that would no longer

cause any trouble. This is so because an equilibria stream requires the existence of an equilibrium at every time point.

In this section, we address the problem of inexistent equilibria streams, also known as *global inconsistency*. We begin by defining a notion of coherence associated with individual contexts which allows us to first establish sufficient conditions for the existence of equilibria streams, and then abstract away from problems due to specific incoherent contexts and focus on those problems essentially caused by the way the flow of information in rMCSs is organized through its bridge rules. We introduce the notion of a *repair*, which modifies an rMCS by changing its bridge rules at some particular point in time in order to obtain some equilibria stream, which we dub *repaired equilibria stream*. We establish sufficient conditions for the existence of repaired equilibria streams and briefly discuss different possible strategies to define such repairs. However, repaired equilibria streams may not always exist either, e.g., because some particular context is incoherent. To deal with such situations, we relax the concept of equilibria stream and introduce the notion of *partial equilibria stream*, which essentially allows the non-existences of equilibria at some time points. It turns out that *partial equilibria streams* always exist thus solving the problem of global inconsistency for rMCSs.

In related work, the problem of global inconsistency has been addressed in the context of mMCSs [15] by establishing sufficient conditions for the existence of equilibria. We also follow that idea, but among the two notions established in [15], *diagnosis* and *explanation*, the former corresponding to rules that need to be altered to restore consistency, and the latter corresponding to combinations of rules that cause inconsistency, we focus on the former. This is justified by the fact that the two notions turn out to be dual of each other, and somehow correspond to our notion of repair. The main difference here is that we opt to not allow the (non-standard) strengthening of bridge-rule to restore consistency, and, of course, the fact that our repairs need to take into account the dynamic nature of rMCSs.

We start by introducing two notions of global consistency differing only on whether we consider a particular input stream or all possible input streams.

Definition 13. *Let M be an rMCS, KB a configuration of knowledge bases for M , and \mathcal{I} an input stream for M . Then, M is consistent with respect to KB and \mathcal{I} if there exists an equilibria stream of M given KB and \mathcal{I} . M is strongly consistent with respect to KB if, for every input stream \mathcal{I} for M , M is consistent with respect to KB and \mathcal{I} .*

Obviously, for a fixed configuration of knowledge bases, strong consistency implies consistency w.r.t. any input stream, but not vice-versa.

Unfortunately, verifying strong consistency is in general highly complex since it requires checking all possible equilibria streams. Nevertheless, we can establish conditions that ensure that an rMCS M is strongly consistent with respect to a given configuration of knowledge bases KB , hence guaranteeing the existence of an equilibria stream independently of the input. It is based on two notions – *totally coherent contexts* and *acyclic rMCSs* – that together are sufficient to ensure (strong) consistency.

Total coherence imposes that each knowledge base of a context always has at least one acceptable belief set.

Definition 14. *A context C_i is totally coherent if $\text{acc}_i(kb) \neq \emptyset$, for every $kb \in KB_i$.*

The second notion describes cycles between contexts which may be a cause of inconsistency. Acyclic rMCSs are those whose bridge rules have no cycles.

Definition 15. Given an rMCS $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$, \triangleleft_M is the binary relation over contexts of M such that $(C_i, C_j) \in \triangleleft_M$ if there is a bridge rule $r \in \text{BR}_i$ and $j: b \in \text{bd}(r)$ for some b . If $(C_i, C_j) \in \triangleleft_M$, also denoted by $C_i \triangleleft_M C_j$, we say that C_i depends on C_j in M , dropping the reference to M whenever unambiguous.

Definition 16. An rMCS M is acyclic if the transitive closure of \triangleleft_M is irreflexive.

We can show that acyclicity and total coherence together are indeed sufficient to ensure strong consistency.

Proposition 1. Let $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$ be an acyclic rMCS such that every C_i , $1 \leq i \leq n$, is totally coherent, and KB a configuration of knowledge bases for M . Then, M is strongly consistent with respect to KB.

Nevertheless, these conditions are rather restrictive since there are many useful cyclic rMCSs which only under some particular configurations of knowledge bases and input streams may have no equilibria streams.

To deal with these, and recover an equilibria stream, one possibility is to repair the rMCSs by locally, and selectively, eliminating some of its bridge rules. Towards introducing the notion of *repair*, given an rMCS $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$, we denote by br_M the set of all bridge rules of M , i.e., $br_M = \bigcup_{1 \leq i \leq n} \text{BR}_i$. Moreover, given a set $R \subseteq br_M$, denote by $M[R]$ the rMCS obtained from M by restricting the bridge rules to those not in R .

Definition 17 (Repair). Let $M = \langle C, \text{IL}, \text{BR} \rangle$ be an rMCS, KB a configuration of knowledge bases for M , and \mathcal{I} an input stream for M until τ where $\tau \in \mathbb{N} \cup \{\infty\}$. Then, a repair for M given KB and \mathcal{I} is a function $\mathcal{R} : [1..\tau] \rightarrow 2^{br_M}$ such that there exists a function $\mathcal{B} : [1..\tau] \rightarrow \text{Bel}_M$ such that

- \mathcal{B}^t is an equilibrium of $M[\mathcal{R}^t]$ given \mathcal{KB}^t and \mathcal{I}^t , where \mathcal{KB}^t is inductively defined as
 - $\mathcal{KB}^1 = \text{KB}$
 - $\mathcal{KB}^{t+1} = \text{upd}_{M[\mathcal{R}^t]}(\mathcal{KB}^t, \mathcal{I}^t, \mathcal{B}^t)$.

We refer to \mathcal{B} as a repaired equilibria stream of M given KB, \mathcal{I} and \mathcal{R} .

Note the generality of this notion, which considers to be a *repair* essentially any sequence of bridge rules (defined by the repair function \mathcal{R}) that, if removed from the rMCS at their corresponding time point, will allow for an equilibrium at that time point. This may include repairs that unnecessarily eliminate some bridge rules, and even the *empty repair* i.e. the repair \mathcal{R}_\emptyset such that $\mathcal{R}_\emptyset^t = \emptyset$ for every t , whenever M already has an equilibria stream given KB and \mathcal{I} . This ensures that the set of repaired equilibria streams properly extends the set of equilibria streams, since equilibria streams coincide with repaired equilibria streams given the empty repair.

Proposition 2. *Every equilibria stream of M given KB and \mathcal{I} is a repaired equilibria stream of M given KB , \mathcal{I} and the empty repair \mathcal{R}_\emptyset .*

It turns out that for rMCSs composed of totally coherent contexts, repaired equilibria streams always exist.

Proposition 3. *Let $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$ be an rMCS such that each C_i , $i \in \{1, \dots, n\}$, is totally coherent, KB a configuration of knowledge bases for M , and \mathcal{I} an input stream for M until τ . Then, there exists $\mathcal{R} : [1.. \tau] \rightarrow 2^{\text{br}_M}$ and $\mathcal{B} : [1.. \tau] \rightarrow \text{Bel}_M$ such that \mathcal{B} is a repaired equilibria stream given KB , \mathcal{I} and \mathcal{R} .*

Whenever repair operations are considered in the literature, e.g., in the context of databases [4], there is a special emphasis on seeking repairs that are somehow minimal, the rational being that we want to change things as little as possible to regain consistency. In the case of repairs of rMCS, it is easy to establish an order relation between them, based on a comparison of the bridge rules to be deleted at each time point.

Definition 18. *Let \mathcal{R}_a and \mathcal{R}_b be two repairs for some rMCS M given a configuration of knowledge bases for M , KB and \mathcal{I} , an input stream for M until τ . We say that $\mathcal{R}_a \leq \mathcal{R}_b$ if $\mathcal{R}_a^i \subseteq \mathcal{R}_b^i$ for every $i \leq \tau$, and that $\mathcal{R}_a < \mathcal{R}_b$ if $\mathcal{R}_a \leq \mathcal{R}_b$ and $\mathcal{R}_a^i \subset \mathcal{R}_b^i$ for some $i \leq \tau$.*

This relation can be straightforwardly used to check whether a repair is minimal, and we can restrict ourselves to adopting minimal repairs. However, there may be good reasons to adopt non-minimal repairs, e.g., so that they can be determined *as we go*, or so that *deleted* bridge rules are not reinstated, etc. Even though investigating specific types of repairs falls outside the scope of this paper, we nevertheless present and briefly discuss some possibilities.

Definition 19 (Types of Repairs). *Let \mathcal{R} be a repair for some rMCS M given KB and \mathcal{I} . We say that \mathcal{R} is a:*

Minimal Repair *if there is no repair \mathcal{R}_a for M given KB and \mathcal{I} such that $\mathcal{R}_a < \mathcal{R}$.*

Global Repair *if $\mathcal{R}^i = \mathcal{R}^j$ for every $i, j \leq \tau$.*

Minimal Global Repair *if \mathcal{R} is global and there is no global repair \mathcal{R}_a for M given KB and \mathcal{I} such that $\mathcal{R}_a < \mathcal{R}$.*

Incremental Repair *if $\mathcal{R}^i \subseteq \mathcal{R}^j$ for every $i \leq j \leq \tau$.*

Minimally Incremental Repair *if \mathcal{R} is incremental and there is no incremental repair \mathcal{R}_a and $j \leq \tau$ such that $\mathcal{R}_a^i \subset \mathcal{R}^i$ for every $i \leq j$.*

Minimal repairs perhaps correspond to the ideal situation in the sense that they never unnecessarily remove bridge rules. In some circumstances, it may be the case that if a bridge rule is somehow involved in some inconsistency, it should not be used at any time point, leading to the notion of *global repair*. Given the set of all repairs, checking which are global is also obviously less complex than checking which are

minimal. A further refinement – *minimal global repairs* – would be to only consider repairs that are minimal among the global ones, which would be much simpler to check than checking whether it is simply minimal. Note that a minimal global repair is not necessarily a minimal repair. One of the problems with these types of repairs is that we can only check whether they are of that type globally, i.e., we can only check once we know the entire input stream \mathcal{I} . This was not the case with *plain* repairs, as defined in Definition 17, which could be checked *as we go*, i.e., we can determine what bridge rules to include in the repair at a particular time point by having access to the input stream \mathcal{I} up to that time point only. This is important so that rMCSs can be used to effectively react to their environment. The last two types of repairs defined above allow for just that. *Incremental repairs* essentially impose that removed bridge rules cannot be reused in the future, i.e., that the set of removed bridge rules monotonically grows with time, while *minimally incremental repairs* further impose that only minimal sets of bridge rules can be added at each time point. Other types of repairs could be defined, e.g., by defining some priority relation between bridge rules, some distance measure between subsets of bridge rules and minimize it when considering the repair at consecutive time points, among many other options, whose investigation we leave for future work. Repairs could also be extended to allow for the strengthening of bridge rules, besides their elimination, generalizing ideas from [15] and [10] where the extreme case of eliminating the entire body of bridge rules as part of a repair is considered.

Despite the existence of repaired equilibria streams for large classes of systems, two problems remain: first, computing a repair may be excessively complex, and second, there remain situations where no repaired equilibria stream exists, namely when the rMCS contains contexts that are not totally coherent. The second issue could be dealt with by ensuring that for each non-totally coherent context there would be some bridge rule with a management operation in its head that would always restore consistency of the context, and that such rule could always be *activated* through a repair (for example, by adding a negated reserved atom to its body, and another bridge rule with that atom in its head and an empty body, so that removing this latter rule through a repair would activate the management function and restore consistency of the context). But this would require special care in the way the system is specified, and its analysis would require a very complex analysis of the entire system including the specific behavior of management functions. In practice, it would be quite hard – close to impossible in general – to ensure the existence of repaired equilibria streams, and we would still be faced with the first problem, that of the complexity of determining the repairs.

A solution to this problem is to relax the notion of equilibria stream so that it does not require an equilibrium at every time point. This way, if no equilibrium exists at some time point, the equilibria stream would be undefined at that point, but possibly defined again in subsequent time points. This leads to the following notion of *partial equilibria stream*.⁴

⁴The notion of partial equilibria has also been used in the static case in [14], but in a quite different sense, namely to focus only on those contexts that are relevant for determining the semantics of a given context C .

Definition 20 (Partial Equilibria Stream). Let $M = \langle C, IL, BR \rangle$ be an rMCS, $KB = \langle kb_1, \dots, kb_n \rangle$ a configuration of knowledge bases for M , and \mathcal{I} an input stream for M until τ where $\tau \in \mathbb{N} \cup \{\infty\}$. Then, a partial equilibria stream of M given KB and \mathcal{I} is a partial function $\mathcal{B} : [1..\tau] \rightarrow \text{Bel}_M$ such that

- \mathcal{B}^t is an equilibrium of M given \mathcal{KB}^t and \mathcal{I}^t , where \mathcal{KB}^t is inductively defined as

$$\begin{aligned} & - \mathcal{KB}^1 = KB \\ & - \mathcal{KB}^{t+1} = \begin{cases} \text{upd}_M(\mathcal{KB}^t, \mathcal{I}^t, \mathcal{B}^t), & \text{if } \mathcal{B}^t \text{ is not undefined.} \\ \mathcal{KB}^t, & \text{otherwise.} \end{cases} \end{aligned}$$

- or \mathcal{B}^t is undefined.

As expected, this is a proper generalization of the notion of equilibria stream:

Proposition 4. Every equilibria stream of M given KB and \mathcal{I} is a partial equilibria stream of M given KB and \mathcal{I} .

And it turns out that partial equilibria streams always exist.

Proposition 5. Let M be an rMCS, KB a configuration of knowledge bases for M , and \mathcal{I} an input stream for M until τ . Then, there exists $\mathcal{B} : [1..\tau] \rightarrow \text{Bel}_M$ such that \mathcal{B} is a partial equilibria stream given KB and \mathcal{I} .

One final word to note is that partial equilibria streams not only allow us to deal with situations where equilibria do not exist at some time instants, but they also open the ground to consider other kinds of situations where we do not wish to consider equilibria at some time point, for example because we were not able to compute them on time, or simply because we do not wish to process the input at every time point, e.g., whenever we just wish to sample the input with a lower frequency than it is generated. If we wish to restrict that partial equilibria streams only relax equilibria streams when necessary, i.e., when equilibria do not exist at some time point, we can further impose the following condition on Definition 20:

$$\mathcal{B}^t \text{ is undefined} \implies \text{there is no equilibrium of } M \text{ given } \mathcal{KB}^t \text{ and } \mathcal{I}^t.$$

5. Non-Determinism and Well-Founded Semantics

Reactive MCSs as considered so far are non-deterministic for two reasons. On the one hand, we allow for contexts whose semantics may return multiple belief sets for the same knowledge base, and on the other hand, the flow of information between contexts established by bridge rules may be the source of non-determinism. As this leads to multiple equilibria and thus to exponentially many equilibria streams, in practice, this may be undesired, which is why it is important to study under which conditions non-determinism can be avoided.

The first source of non-determinism solely depends on the choice of the contexts occurring in the rMCS, i.e., it can be avoided when determining/designing the rMCS in question. The second source, however, requires to consider the entire rMCS when eliminating the non-determinism. This can be achieved by introducing preferences on equilibria using, e.g., preference functions as proposed by Ellmauthaler [17]. One might also adopt language constructs for expressing preferences in ASP such as optimization statements [24] or weak constraints [13], which essentially assign a quality measure to equilibria, or, more recently, add a pre-order on the contexts [36]. Alternatively, and inspired by notions developed for MCSs [9], we may consider restrictions on rMCSs such that these non-determinisms do not occur in the first place. In the remainder of this section, we will focus on such restrictions leading to an alternative well-founded semantics for rMCSs.

As a first step towards this objective, only equilibria that are subset-minimal will be considered, also with the aim of avoiding unnecessary self-justifications of information resulting from the interaction of various contexts. Then, grounded equilibria as a special case for so-called reducible MCSs will be presented for which the existence of minimal equilibria can be effectively checked. Subsequently, a well-founded semantics for such reducible MCSs will be defined under which an approximation of all grounded equilibria can be computed deterministically. In the following, we transfer these notions from static MCSs in [9] to dynamic rMCSs and discuss under which (non-trivial) conditions they actually can be applied.

We start with the notion of minimal equilibria. Formally, given an rMCS $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$, a configuration of knowledge bases KB for M and an input l for M , an equilibrium $B = \langle B_1, \dots, B_n \rangle$ of M given KB and l is *minimal* if there is no equilibrium $B' = \langle B'_1, \dots, B'_n \rangle$ of M given KB and l such that $B'_i \subseteq B_i$ for all i with $1 \leq i \leq n$ and $B'_j \subsetneq B_j$ for some j with $j \in \{1 \dots n\}$.

This notion of minimality avoids to some extent unnecessary self-justifications in equilibria (streams). A simple example of such self-justification is shown in Example 2 in [9], which can be transcribed to our setting using a simple rMCS without input languages, a single propositional context C_1 with empty kb and BR_1 only containing $\text{add}(a) \leftarrow 1:a$. Then both $\langle \emptyset \rangle$ and $\langle \{a\} \rangle$ are equilibria. The latter is considered dubious as a is justified by itself and minimality would avoid such self-justification. Note that this does not always work.

Example 6. *A single storage context with empty kb and bridge rules $\text{add}(b) \leftarrow \text{not } 1:a$. and $\text{add}(a) \leftarrow 1:a$. has two minimal equilibria $\langle \{a\} \rangle$ and $\langle \{b\} \rangle$, the former being self-justified.*

Still, avoiding non-minimality whenever possible certainly is a decent guiding principle, and, as we will see later, it can be avoided altogether under certain conditions.

The problem is that checking for minimality commonly raises the computational complexity of determining equilibria. To avoid that, we now formalize conditions under which minimal equilibria can be effectively checked.

For that purpose, we start by introducing the notion of monotonic logics. Namely, a logic $L = \langle \text{KB}, \text{BS}, \text{acc} \rangle$ is *monotonic* if

1. $\text{acc}(kb)$ is a singleton set for each $kb \in \text{KB}$, and

2. $B \subseteq B'$ whenever $kb \subseteq kb'$, $\mathbf{acc}(kb) = \{B\}$, and $\mathbf{acc}(kb') = \{B'\}$.

In other words, L is monotonic if \mathbf{acc} is deterministic and monotonic.

Maybe not surprisingly, the logics of non-monotonic formalisms such as L_a for the answer set semantics in Example 1 do not satisfy this definition. We therefore proceed with introducing the notion of a reducible logic, which covers monotonic logics, but also includes logics that can be reduced to monotonic ones given a belief set of the logic.

Logic $L = \langle KB, BS, \mathbf{acc} \rangle$ is *reducible* iff, for some $KB^* \subseteq KB$ and some reduction function $\mathbf{red} : KB \times BS \rightarrow KB^*$,

1. $\langle KB^*, BS, \mathbf{acc} \rangle$ is monotonic,
2. for each $kb \in KB$, and all $B, B' \in BS$:
 - $\mathbf{red}(kb, B) = kb$ whenever $kb \in KB^*$,
 - $\mathbf{red}(kb, B) \subseteq \mathbf{red}(kb, B')$ whenever $B' \subseteq B$, and
 - $\mathbf{acc}(\mathbf{red}(kb, B)) = \{B\}$ iff $B \in \mathbf{acc}(kb)$.

Intuitively, L is reducible, if a) the restriction of L to KB^* is monotonic, and b) there is a reduction function which should not reduce kb further if it already is in KB^* , which is antitonic, and by means of which acceptability of a belief set can be checked by the reduction (see also [9]).

In a reducible rMCS, the logics of all contexts have to be reducible. Additionally, we require \mathbf{red} and \mathbf{mng} to be applicable in arbitrary order for all contexts, and that sequences of increasing sets of operations are monotonic in the following sense. Given rMCS $M = \langle \langle C_1, \dots, C_n \rangle, \mathbf{IL}, \mathbf{BR} \rangle$ and, for some $i \in \{1, \dots, n\}$, $kb \in KB_i$, we say that $OP \subseteq OP_i$ in context C_i is *monotonic w.r.t. kb*, if $kb \subseteq \mathbf{mng}(OP, kb)$. Intuitively, a set of monotonic operations will not remove content from the knowledge base to which it is applied.

Definition 21. Let $M = \langle \langle C_1, \dots, C_n \rangle, \mathbf{IL}, \mathbf{BR} \rangle$ be an rMCS, KB a configuration of knowledge bases for M and $S_i = \{\mathbf{hd}(r) \in OP_i \mid r \in BR_i\}$ for all $i \in \{1, \dots, n\}$. Then, M is reducible given KB iff, for all $i \in \{1, \dots, n\}$,

1. L_i is reducible;
2. for all $OP \subseteq S_i$, all $kb \in KB_i$, and all $B_i \in BS_i$:

$$\mathbf{red}_i(\mathbf{mng}_i(OP, kb), B_i) = \mathbf{mng}_i(OP, \mathbf{red}_i(kb, B_i));$$

3. for all $OP^1 \subseteq \dots \subseteq OP^m \subseteq S_i$ and all $j \in \{1, \dots, m\}$: OP^j is monotonic w.r.t. kb_i^j where
 - (a) $kb_i^1 = kb_i$;
 - (b) $kb_i^h = \mathbf{mng}_i(OP^{h-1}, kb_i^{h-1})$ for $h \in \{2, \dots, m\}$.

The second condition essentially requires that no \mathbf{mng}_i introduces formulas that are affected by \mathbf{red}_i . The third condition in addition checks whether any sequence of increasing sets of (applicable) operations is monotonic, i.e., whether some kb_i^m can be obtained in a step-wise, monotonic iteration.

Example 7. Recall M_{al} from Section 2.4, for which we want to check whether it is reducible. First, all contexts but C_{ec} build on a monotonic logic, so their logics are automatically reducible. This also means that the second condition is trivially satisfied for these contexts using the identity function for \mathbf{red}_i . At the same time, L_{ec} is reducible using the well-known Gelfond-Lifschitz reduct [25], and clearly none of the bridge rules in BR_{ec} can introduce any knowledge base formula (for any $kb \in KB_{ec}$) which is affected by such \mathbf{red}_{ec} . More care needs to be taken regarding the third condition: in fact, e.g., applying $\mathbf{setTemp}(\mathit{hot})$ to $\{\mathit{tm}(\mathit{cold})\}$ yields $\{\mathit{tm}(\mathit{hot})\}$, which is clearly not monotonic. This can however be avoided if it is ensured that kb_{st} does not contain information on the temperature. Then, any increasing sequence of operations is monotonic w.r.t. the corresponding kb_i^j (cf. Example 2). The same is not true for BR_{pos} which is why we have to omit the bridge rules $\mathbf{setPos}(P) \leftarrow \mathit{pos}::\mathit{enters}(P)$. in BR_{pos} , thus only preventing the information of Dave changing the room from being available right away. With this slight adjustment, the rMCS M_{al} is in fact reducible.

Following notions from logic programming, we introduce definite rMCSs.

Definition 22. Let $M = \langle \langle C_1, \dots, C_n \rangle, \mathbb{L}, \mathbb{BR} \rangle$ be a reducible rMCS given KB. Then, M is definite given KB iff, for all $i \in \{1, \dots, n\}$,

1. no $r \in BR_i$ contains **not**;
2. for all $B \in BS_i$: $kb_i = \mathbf{red}(kb_i, B)$.

Thus, in a definite rMCS (given KB), the bridge rules are monotonic and all knowledge bases are already reduced. This suggests the following iteration.

Definition 23. Let M be a definite rMCS given KB, and l an input for M . For all $i \in \{1 \dots n\}$, let $kb_i^0 = kb_i$ and define, for each successor ordinal $\alpha + 1$,

$$kb_i^{\alpha+1} = \mathbf{mng}(\mathbf{app}_i^{\mathit{now}}(l, B^\alpha), kb_i^\alpha),$$

where $B^\alpha = \langle B_1^\alpha, \dots, B_n^\alpha \rangle$ and $\mathbf{acc}_i(kb_i^\alpha) = \{B_i^\alpha\}$ for any ordinal α . Furthermore, for each limit ordinal α , define $kb_i^\alpha = \bigcup_{\beta \leq \alpha} kb_i^\beta$, and let $kb_i^\infty = \bigcup_{\alpha > 0} kb_i^\alpha$.

We next show two properties of the iteration defined in Definition 23 that will prove useful subsequently.

Lemma 1. Let M be a definite rMCS given KB, and l an input for M . The following holds for the iteration in Definition 23 for all ordinals α :

1. M is definite given KB^α ;
2. for any ordinal β with $\beta \leq \alpha$ we have $kb_i^\beta \subseteq kb_i^\alpha$ for $i \in \{1, \dots, n\}$.

Lemma 1 guarantees that the iteration is well-defined and monotonic and we can show that it yields an equilibrium. In fact, it is the unique minimal equilibrium of M given KB and l .

Proposition 6. Let M be a definite rMCS given KB, and l an input for M . Then $\langle B_1^\infty, \dots, B_n^\infty \rangle$ is the unique minimal equilibrium of M given KB and l . We call it grounded equilibrium of M given KB and l , denoted by $\mathbf{GE}(M, \mathbb{KB}, l)$.

As pointed out in [9], for many logics, $kb_i^\infty = kb_i^\omega$ holds and the iteration even stops after finitely many steps. This is also the case for the slightly revised scenario in Example 7. Moreover, as a consequence of Proposition 6, no self-justification can occur in the grounded equilibrium of definite rMCSs.

This fixpoint iteration cannot be applied to arbitrary reducible rMCSs right away as, e.g., **not** in the bridge rule bodies and non-reduced knowledge bases may allow the removal of already derived information in the iteration. To counter that, we introduce the notion of a reduct for rMCSs where, for a bridge rule of the form (1), $\text{bd}(r)^+ = \{a_1, \dots, a_j\}$ and $\text{bd}(r)^- = \{a_{j+1}, \dots, a_m\}$.

Definition 24. Let $M = \langle C, \text{IL}, \langle BR_1, \dots, BR_n \rangle \rangle$ be a reducible rMCS given KB, I an input for M , and $\text{B} = \langle B_1, \dots, B_n \rangle$ a belief state of M . The (I, B) -reduct of M and KB is obtained as $M^{(\text{I}, \text{B})} = \langle \text{IL}, C, \langle BR_1^{(\text{I}, \text{B})}, \dots, BR_n^{(\text{I}, \text{B})} \rangle \rangle$ where $BR_i^{(\text{I}, \text{B})} = \{\text{hd}(r) \leftarrow \text{bd}(r)^+ \mid r \in BR_i, \text{hd}(r) = \text{op}$ with $\text{op} \in \text{OP}_i, \langle \text{I}, \text{B} \rangle \not\models a_\ell$ for all $a_\ell \in \text{bd}(r)^-\}$ and as $\text{KB}^{(\text{I}, \text{B})} = \langle kb_1^{B_1}, \dots, kb_n^{B_n} \rangle$ where $kb_i^{B_i} = \text{red}_i(kb_i, B_i)$.

Note that bridge rules with an operation under next in its head are ignored here, as these do not affect the computation of the equilibria anyway.

For all reducible rMCSs M given KB, all inputs I for M , and all belief states B for M , the (I, B) -reduct of M and KB is definite. In this case, we can check whether B is a grounded equilibrium for M given KB and I in the usual manner.

Definition 25. Let $M = \langle C, \text{IL}, \text{BR} \rangle$ be a reducible rMCS given KB, and I an input for M . A belief state B of M is a grounded equilibrium of M given KB and I iff $\text{B} = \mathbf{GE}(M^{(\text{I}, \text{B})}, \text{KB}^{(\text{I}, \text{B})}, \text{I})$.

Grounded equilibria of reducible rMCSs given some KB are also minimal.

Proposition 7. Every grounded equilibrium of a reducible rMCS M given KB and an input I is a minimal equilibrium of M given KB and I .

Grounded equilibria of reducible rMCSs are not unique equilibria in general. Consider, e.g., again Example 6. Then $\langle \{b\} \rangle$ is a grounded equilibrium, while $\langle \{a\} \rangle$ is not. Still, as checking for grounded equilibria relies on the unique grounded equilibrium of the reduct, we know that no self-justification can occur in grounded equilibria, which is also why $\langle \{a\} \rangle$ is filtered.

We can now introduce grounded equilibria streams provided that rMCS M given KB is reducible for each KB occurring in \mathcal{KB} . Unfortunately, checking for reducibility cannot be done a priori as the KB will only be known at runtime (due to 3. of Definition 21). To complete the picture, we also spell out under which conditions an initially reducible rMCS given KB remains reducible.

Definition 26. Let $M = \langle C, \text{IL}, \text{BR} \rangle$ be a reducible rMCS given KB. Then M is persistently reducible given KB iff for all sequences $\text{OP}^1, \dots, \text{OP}^m$ with $m > 0$ and $\text{OP}^j \subseteq \{\text{op} \mid \text{next}(\text{op}) \in \text{hd}(r), r \in BR_i\}$, M is reducible given KB^j where $\text{KB}^1 = \text{KB}$ and $\text{KB}^j = \langle kb_1^j, \dots, kb_n^j \rangle$ for $j > 1$ with $kb_i^j = \mathbf{mng}(\text{OP}^{j-1}, kb_i^{j-1})$.

Thus, any reducible rMCS (given some concrete KB) is persistently reducible if applying any arbitrary sequence of the operations under `next` occurring in the bridge rules yields again a reducible rMCS.

It should be noted that checking for reducible rMCSs is not trivial in general and even less so for persistently reducible rMCSs. Still, certain kinds of contexts clearly satisfy the restrictions including those occurring in the example scenario in Example 7. Namely, they either store certain pieces of information, but only do change them using `next` and not otherwise, or their *kb* is initially empty and this is never changed using `next`, so information is only stored temporarily for the computation of the current equilibrium. The scenario is thus persistently reducible for any knowledge base configuration which does not store pieces of information in the initial kb_i for contexts of the latter kind.

This allows us to introduce grounded equilibria streams.

Definition 27. Let $M = \langle C, \text{IL}, \text{BR} \rangle$ be a persistently reducible rMCS given KB, \mathcal{I} an input stream for M until s , $\mathcal{B} = \langle \mathcal{B}^1, \dots, \mathcal{B}^s \rangle$ an equilibria stream of M given KB and \mathcal{I} , and \mathcal{KB} the configurations stream of M given KB, \mathcal{I} , and \mathcal{B} . Then, \mathcal{B} is a grounded equilibria stream of M given KB and \mathcal{I} iff, for each $t \in \{1 \dots s\}$, $\mathcal{B}^t = \mathbf{GE}(M^{(l, \mathcal{B})}, (\mathcal{KB}^t)^{(l, \mathcal{B})}, \mathcal{I}^t)$.

A grounded equilibria stream of some M given KB and \mathcal{I} can thus be understood as a stream of grounded equilibria. It is thus straightforward to see that each equilibrium in a grounded equilibria stream is minimal w.r.t. the knowledge base configuration and input of its time point.

We remark that the notion of persistently reducible rMCSs substantially differs from reducible MCSs on which the former are founded. On the one hand, permitting operations in `mng` beyond simple addition, as used already for mMCSs, requires the non-trivial tests for reducible rMCSs, and in addition, persistent reducibility needs to be verified for the dynamic aspect of rMCSs. Thus, these new notions also extend the work in [9] from MCSs to mMCSs.

For reducible rMCSs in general, determining grounded equilibria is not deterministic yet, since we have to guess and check the grounded equilibrium in each step. This is why we now introduce the well-founded semantics for reducible rMCSs M following the ideas in [9]. Its definition is based on the operator $\gamma_{M, \text{KB}, l}(\mathcal{B}) = \mathbf{GE}(M^{(l, \mathcal{B})}, \text{KB}^{(l, \mathcal{B})}, l)$, provided BS_i for each logic L_i in all the contexts of M has a least element B_i^* (w.r.t. subset inclusion). Such rMCSs are called *normal*.

It can be shown that $\gamma_{M, \text{KB}, l}$ is antitonic which means that applying $\gamma_{M, \text{KB}, l}$ twice yields a monotonic operator. Hence, by the Knaster-Tarski theorem, $(\gamma_{M, \text{KB}, l})^2$ has a least fixpoint which determines the well-founded semantics.

Definition 28. Let M be a normal, reducible rMCS given KB, and l an input for M . The well-founded model of M given KB and l , denoted $\text{WF}(M, \text{KB}, l)$, is the least fixpoint of $(\gamma_{M, \text{KB}, l})^2$.

Starting with the belief state $\mathcal{B}^* = \langle B_1^*, \dots, B_n^* \rangle$, this fixpoint can be iterated, establishing the relation between $\text{WF}(M, \text{KB}, l)$ and grounded equilibria of M .

Proposition 8. *Let $M = \langle C, IL, BR \rangle$ be a normal, reducible rMCS given KB, I an input for M , $WF(M, KB, I) = \langle W_1, \dots, W_n \rangle$, and $B = \langle B_1, \dots, B_n \rangle$ a grounded equilibrium of M given KB and I . Then $W_i \subseteq B_i$ for $i \in \{1 \dots n\}$.*

The well-founded model of M can thus be viewed as the belief state representing what is accepted in all grounded equilibria, even though $WF(M, KB, I)$ may itself not necessarily be a grounded equilibrium. Yet, if the rMCS is acyclic (i.e., no cyclic dependencies over bridge rules exist between beliefs in the rMCS, see Sect. 4), then the grounded equilibrium of M given KB and I is unique and identical to the well-founded model. This is indeed the case for the scenario in Example 7.

The well-founded semantics can be generalized to streams as follows.

Definition 29. *Let $M = \langle C, IL, BR \rangle$ be a normal, persistently reducible rMCS given KB, and \mathcal{I} an input stream for M until τ . The well-founded stream of M given KB and \mathcal{I} is a function $\mathcal{WF} : [1.. \tau] \rightarrow \text{Bel}_M$ such that*

- \mathcal{WF}^t is the well-founded model of M given \mathcal{KB}^t and \mathcal{I}^t , where \mathcal{KB}^t is defined as
 - $\mathcal{KB}^1 = \text{KB}$
 - $\mathcal{KB}^t = \text{upd}_M(\mathcal{KB}^{t-1}, \mathcal{I}^{t-1}, \mathcal{WF}^{t-1})$, for $t > 1$.

Clearly, Proposition 8 also generalizes to the well-founded stream (of M given KB and I) including all the made observations. That is, the well-founded stream may not coincide with any (grounded) equilibria stream, unless the rMCS in question is acyclic, in which case the grounded equilibria stream is unique and does coincide with the well-founded stream. Again, this can be verified for the scenario in Example 7.

6. Complexity

In this section, we want to analyze the complexity of answering queries over equilibria streams of rMCSs. As our framework has many degrees of freedom, we need to impose some restrictions in order to get meaningful complexity results. In particular, we are only interested in decision problems with input of finite size, since otherwise the decision problem would unfortunately be undecidable right away. Thus, in the following, we only consider *finite* rMCSs, i.e., we do not consider rule schemas (which stand potentially for an infinite amount of bridge rules), and assume that all knowledge bases in the given rMCS are finite. Also, we restrict our considerations to finite input streams.

We start by introducing the two reasoning problems we consider.

Definition 30. *The problem Q^\exists , respectively Q^\forall , is deciding whether for a given finite rMCS M , a belief b for the k -th context of M , a configuration of knowledge bases KB for M , and an input stream \mathcal{I} until τ , it holds that $b \in B_k$ for some $\mathcal{B}^t = \langle B_1, \dots, B_n \rangle$, ($1 \leq t \leq \tau$), for some, respectively all, equilibria stream(s) \mathcal{B} given KB and \mathcal{I} .*

$\mathcal{CC}(M, k:b)$	Q^\exists	Q^\forall
P	NP	coNP
$\Delta_i^P (i \geq 2)$	Σ_i^P	Π_i^P
$\Sigma_i^P (i \geq 1)$	Σ_i^P	Π_i^P
PSPACE	PSPACE	PSPACE
EXPTIME	EXPTIME	EXPTIME

Table 1: Complexity results of checking Q^\exists and Q^\forall

As the complexity of an rMCS depends on that of its individual contexts, we introduce the notion of *context complexity* (cf. [15]). To do so, we need to focus on relevant parts of belief sets by means of projection. Intuitively, among all beliefs, we only need to consider belief b that we want to query and the beliefs that contribute to the application of bridge rules for deciding Q^\exists or Q^\forall . Given M , b , k , and \mathcal{I} as in Definition 30, the set of relevant beliefs for a context C_i of M is given by

$$RB_i(M, k:b) = \{b' \mid r \in BR_h, i:b' \in \text{bd}(r) \vee \text{not } i:b' \in \text{bd}(r), h \in \{1 \dots, n\}\} \cup \{b \mid k = i\}.$$

Then, a *projected belief state* for M and $k:b$ is a tuple

$$B_{|M}^{k:b} = \langle B_1 \cap RB_1(M, k:b), \dots, B_n \cap RB_n(M, k:b) \rangle$$

where $B = \langle B_1, \dots, B_n \rangle$ is a belief state for M . If B is an equilibrium, then we call this tuple *projected equilibrium*. The *context complexity* of C_i in M with respect to $k:b$ for a fixed input l is the complexity of deciding the *context problem* of C_i , that is, whether for a given projected belief state $B = \langle B_1, \dots, B_n \rangle$ for M and $k:b$, there is some belief set B'_i for C_i with $B_i = B'_i \cap RB_i(M, k:b)$ and $B'_i \in \text{acc}_i(\text{mng}_i(\text{app}_i(l, B), kb_i))$. The *context complexity* $\mathcal{CC}(M, k:b)$ of an entire rMCS is a (smallest) upper bound for the context complexity classes of its contexts.

Theorem 1. *Table 1 summarizes the complexities of membership of problems Q^\exists and Q^\forall for finite input streams (until some $\tau \in \mathbb{N}$) depending on the context complexity. Hardness also holds if it holds for the context complexity.*

These results can also be used to show the complexity results of a strongly related problem, namely ensuring that a certain belief does always hold over all time instants for some or all equilibria streams. This can be achieved by simply considering the co-problem of ensuring that the negated belief does not hold at any time instant for some or all equilibria streams. In case the considered belief set itself does not allow us expressing this negated belief, an appropriate additional auxiliary context with adequate bridge rules can be used to represent it and query for it.

Note that allowing for input streams of infinite length leads to undecidability.

Proposition 9. *Given a finite rMCS M , the problems Q^\exists and Q^\forall are undecidable for infinite input streams (when $\tau = \infty$).*

The reason is that rMCSs are expressive enough (even with very simple context logics) to simulate a Turing machine such that deciding Q^\exists or Q^\forall for infinite runs solves the halting problem. We provide an rMCS that implements a Turing machine service in Appendix C: a fixed rMCS can read the configuration of a Turing machine TM as well as the corresponding input and then simulate a run of TM on that input.

For persistently reducible rMCSs, it turns out that the complexity results for the analog problems Q_g^\exists and Q_g^\forall on grounded equilibria streams are identical to Q^\exists and Q^\forall on equilibria streams.

Theorem 2. *Let $M = \langle C, IL, BR \rangle$ be a finite, persistently reducible rMCS given KB. Then membership and hardness results for Q_g^\exists and Q_g^\forall on grounded equilibria streams for finite input streams coincide with those for Q^\exists and Q^\forall in Theorem 1.*

Note that this only holds if checking whether M is persistently reducible is already known/can be neglected. While such assumption is a strong one in general, in Section 5 it is argued that this can be done easily for contexts of certain kinds, including the ones used in the example scenario in Section 2.4 (and in Example 7).

Regarding the well-founded stream, we can restrict our attention to the problem $Q_{w_f}^\exists$ for the well-founded stream, since it coincides with $Q_{w_f}^\forall$ for the unique well-founded stream, and to polynomial contexts given the motivation for the well-founded stream.

Theorem 3. *Let $M = \langle C, IL, BR \rangle$ be a finite, normal, persistently reducible rMCS given KB such that $CC(M, k:b) = \mathbf{P}$ for $Q_{w_f}^\exists$. Then, $Q_{w_f}^\exists$ is in \mathbf{P} . In addition, hardness holds provided $CC(M, k:b) = \mathbf{P}$ is hard.*

This result, together with the observation that the well-founded stream coincides with the unique grounded equilibrium stream until s , allows us to verify that computing the results in our use case scenario in Section 2.4 can be done in polynomial time.

7. Related Work

Several systems and frameworks for modeling the dynamics of knowledge and the flow of information have been developed. These systems are obviously related to reactive multi-context systems. In this section, we present those approaches we consider most relevant, stressing differences as well as commonalities with rMCSs. Note that we focus entirely on dynamic approaches and do not include systems which handle heterogeneous information in a static setting. An interesting example of the latter type are Lierler and Truszczyński's abstract modular systems [33]. In a nutshell, modular systems realize the communication between different reasoning modules through joint vocabularies rather than bridge rules. These systems are best viewed as alternatives to classical MCSs. It is an open question how to adapt them to a dynamic setting.

Since rMCSs have been designed for applications involving stream reasoning, we will first consider two recent approaches for stream reasoning, namely LARS [7] and

STARQL [37]. Then, we consider EVOLP [2], a framework that focuses on dynamics in the form of updates in the restricted setting of generalized logic programs building on similar notions as the operator `next` used for rMCSs. Finally, we consider asynchronous multi-context systems [19], and also reactive ASP [23, 8].⁵

7.1. Reactive Multi-Context Systems and Stream Reasoning

Reasoning over streaming data is a topic of increasing interest. Key driver is the central role that streams play in current endeavors towards the Internet of Things, the Internet of Services, as well as the vision of a fourth industrial revolution. Stream reasoning has been a prominent issue in the Semantic Web community for several years,⁶ and has also received substantial attention by researchers from areas such as Knowledge Representation and Reasoning and Databases lately [18].

We consider rMCSs to be a well-suited formalism for stream reasoning that addresses important challenges that naturally arise in this context. One important benefit of rMCSs is that their managing capabilities provide a dynamic way to decide which data to keep for future reasoning and which data to drop. In this respect, rMCSs offer greater flexibility than sliding-window approaches. Nevertheless, as also demonstrated in Section 3, rMCSs can be used to implement windowing techniques.

In the following, we consider two recently proposed frameworks for stream reasoning, LARS [7] and STARQL [37].

LARS

The Logic-based framework for Analyzing Reasoning over Streams (LARS) [7] aims at providing a formal declarative logical language for reasoning with streams. LARS is a rule-based formalism, whose language features not only provide different means to refer to or abstract from time, but also a novel window operator, thus providing a flexible mechanism to represent and reason with views on streaming data. The semantics of LARS is based on the FLP semantics of logic programs [20].

Since the semantics of a LARS program is defined for a fixed input data stream and for a particular time point, it is in fact mainly static.

Given their generality, rMCSs can be used to add a dynamic layer to LARS. In Appendix D.1 we show how that can be done by defining an rMCS with a LARS context. The system continuously feeds this context with the relevant (recent) substream of the input stream such that the LARS semantics can be computed for each time point.

STARQL

STARQL (pronounced Star-Q-L) [37], a framework for ontology-based stream reasoning, is developed within the Optique Project [38] that comes with a stream query language inspired by the RDF query language SPARQL [30]. Streams in this framework come in the form of timestamped Description Logic assertions. Both input as

⁵In Appendix D, we provide detailed comparisons between rMCSs and the stream reasoning approaches LARS and STARQL (Appendix D.1) as well as EVOLP (Appendix D.2).

⁶cf. <http://streamreasoning.org>

well as answers of STARQL queries are streams of this kind. Unlike the related language continuous SPARQL (C-SPARQL) [6], STARQL goes beyond RDF semantics and allows for DL reasoning.

Due to the abstract nature of rMCSs, there are many ways to realize STARQL queries as rMCSs. One such way would be to assume a correspondence of one equilibrium computation per STARQL query evaluation. In that case, STARQL input stream and time stamps can be represented using rMCSs input streams, and rMCSs contexts can be devised to handle the different components of a STARQL query. We illustrate the realization in detail in Appendix D.1.

7.2. EVOLP

The framework of evolving logic programs EVOLP [2] is a powerful approach for modeling updates of (propositional) generalized logic programs. Evolving logic programs are defined as general logic programs built over a special language which allows them to express self-evolution. For that, the language includes a reserved unary predicate, `assert`, whose argument may itself be a full-blown rule, thus making arbitrary nesting possible. The idea of EVOLP is that programs can update their own rules thus describing their possible self-evolution. Moreover, besides self-evolution, evolving logic programs also consider evolution caused by the addition of external rules.

The semantics of evolving logic programs is based on sequences of interpretations, called *evolution stable models*.

Given their general nature, rMCSs can capture EVOLP in such a way that equilibria streams of an rMCS correspond to the evolution stable models of evolving logic programs as we demonstrate in Appendix D.2. At the heart of this correspondence between evolution stable models in EVOLP and equilibria streams for rMCSs is the fact that, conceptionally, the operators `next` and `assert` are rather similar.

7.3. Reactive ASP

Closely related to EVOLP are the two frameworks of Reactive ASP, one implemented as a solver *oclingo* [23] and one described in [8]. The system *oclingo* extends an ASP solver for handling external modules provided at runtime by a controller. The output of these external modules can be seen as the observations of EVOLP. Unlike the observations in EVOLP, which can be rules, external modules in *oclingo* are restricted to produce atoms so the evolving capabilities are very restricted. On the other hand, *oclingo* permits committing to a specific answer set at each state for the sake of efficiency, a feature that is not part of EVOLP. Reactive ASP as described in [8] can be seen as a more straightforward generalization of EVOLP where operations other than `assert` for self-updating a program are permitted. Since EVOLP can be captured by rMCSs, and since rMCSs permit several (evolution) operations in the head of bridge rules, it should be clear that Reactive ASP as described in [8] can be captured by rMCSs.

7.4. Asynchronous Multi-Context Systems

Asynchronous multi-context systems (aMCSs) are a framework for loosely coupling knowledge representation formalisms and services [19]. Like rMCSs, they consist of heterogeneous contexts and are aware of continuous streams of information.

However, the semantics of aMCSs is not defined in terms of equilibria but every context delivers output whenever available. Intuitively, equilibria realize a tight integration approach in which the semantics of the individual contexts are interdependent at each time point, while aMCSs do not reach this high level of integration. Rather, contexts communicate with each other by means of input and output streams over time. Consequently, instead of bridge rules that depend on system-wide equilibria, aMCSs use output rules that define which information should be sent to another context or an output stream of the overall system based on a result of a single context.

A further difference is the role of non-determinism in the semantics of aMCSs and rMCSs. While multiple equilibria give rise to non-determinism at each step in a run, for aMCSs, all accepted belief sets of a context are computed in a consecutive way. Nevertheless, there is also a source of non-determinism in the case of aMCSs. The durations of computations and communication are taken into consideration in aMCS but their lengths are left open. Thus, the order in which contexts finish their computation can influence the system and resulting knowledge bases.

Finally, both aMCSs and rMCSs are very general frameworks that allow for simulating Turing machines (cf. Appendix C) and thus for performing multi-purpose computations. A setup to simulate an rMCS by an aMCS has been presented in [19].

8. Conclusions

In this paper, we have introduced *reactive Multi-Context Systems (rMCSs)*, an adaptation of multi-context systems suitable for continuous reasoning in dynamic environments, with the objective to achieve two goals at the same time: integrating heterogeneous knowledge sources and opening knowledge-based systems for dynamic scenarios in the presence of incoming information. For addressing the first goal, we have built rMCSs upon managed multi-context systems, inheriting their functionality for integrating heterogeneous knowledge sources, admitting also relevant operations on knowledge bases. To accommodate the dynamic aspects, rMCSs are equipped with several extensions. For one, bridge rules in our framework allow for input atoms to incorporate the information of multiple external streams of data. Moreover, contrary to standard MCSs which possess only one type of bridge rules modeling the information flow which needs to be taken into account when equilibria are computed (or the operations that need to be applied in case of mMCSs), rMCS have an additional, different type of bridge rules, distinguished by the occurrence of the operator **next** in the head. These rules are used to specify how the configuration of knowledge bases evolves whenever an equilibrium was computed providing the definition of equilibria streams which define the semantics of rMCSs over time.

The resulting formalism is indeed very expressive offering the capabilities to model the integration of heterogeneous knowledge in the presence of incoming information in different dynamic scenarios. Based on a running example dealing with assisted living, we have demonstrated how to model such different dynamic scenarios using rMCSs and addressed several temporal aspects of modeling such as incorporating time on the object level and forgetting.

Other real world use cases of an even larger scale can be handled as well, such as the one described in [40] where the customs service needs to assess imported cargo

for a variety of risk factors including terrorism, narcotics, food and consumer safety, pest infestation, tariff violations, and intellectual property rights. Assessing this risk, even at a preliminary level, involves extensive knowledge about commodities, business entities, trade patterns, government policies and trade agreements. Some of this knowledge may be external to a given customs agency: for instance the broad classification of commodities according to the international Harmonized Tariff System (HTS), or international trade agreements. Other knowledge may be internal to a customs agency, such as lists of suspected violators or of importers who have a history of good compliance with regulations. In [40], all this extensive knowledge is encoded in ontologies based on description logics and logic programming rules under the answer set semantics, and rMCSs naturally cover these formalisms as shown, e.g., in the running example. In addition, they easily allow for the direct modular integration of information given in databases or for example business rules without the need for any prior conversion (which may not be readily available for arbitrary formalisms). While some of the knowledge in this risk assessment scenario is relatively stable, much of it changes rapidly. Changes are made not only at a specific level, such as knowledge about the expected arrival date of a shipment; but at a more general level as well. For instance, while the broad HTS code for tomatoes (0702) does not change, the full classification and tariffs for cherry tomatoes for import into the US changes seasonally. Here again, rMCSs provide mechanisms to easily make changes no matter if they are of a mere temporary nature or more persistent by using sensor data and the incorporation of knowledge via **next** respectively. And, unlike [40], this flexibility is achieved without having to ensure or to test whether the integrated heterogeneous knowledge is organized in so-called layers.

Naturally, dealing with inconsistency is an important issue in dynamic settings where knowledge changes over time. Indeed, we may face different kinds of inconsistencies when using rMCSs, e.g., in the form of inconsistent sensor input for which we discuss modeling-based solutions. Another type of inconsistency is the absence of equilibria at certain time points which could potentially render an entire system useless. We have addressed this problem first by showing sufficient conditions on the contexts and the bridge rules that ensure the existence of an equilibria stream. In the cases where these conditions are not met, we have presented two possible solutions, one following an approach based on repairs – essentially the selective removal of bridge rules to regain an equilibria stream – and a second by relaxing the notion of equilibria stream to ensure that intermediate inconsistent states can be recovered. Thus, rMCSs remain usable in the face of possible inconsistencies which may always occur in real use cases where knowledge and sensor data from different sources is integrated.

We have also addressed the non-determinism of rMCSs and discussed possibilities for avoiding it in situations where it is not desired. To this end, we have also provided a well-founded stream semantics for rMCSs. The study of computational complexity we conducted confirms worst case complexities similar to managed multi-context systems. Hence, the additional integration of sensor data does not raise the worst case complexity of reasoning with heterogeneous knowledge integrated in multi-context systems. In addition, the results confirm that if very efficient reasoning with such rich knowledge is essential, we can restrict to the well-founded stream. This implies some restrictions on the permitted expressiveness, but we have argued that often these restrictions can be

accommodated.

Finally, we compared our approach to related work. Most importantly, we have shown that rMCSs can capture two relevant approaches in stream reasoning, namely LARS [7] and STARQL [37], thus showcasing that rMCSs are indeed “defined in the right way” for the intended integration of constant streams of data.

Regarding future work, several interesting lines of research can be considered. First, we may want to extend rMCSs to enhance their applicability in an even wider set of scenarios. Namely, one might adopt language constructs for expressing preferences in ASP such as optimization statements [24] or weak constraints [13], which assign a quality measure to equilibria. This would allow us, e.g., to improve on the potentially high non-determinism when faced with several equilibria at one time point and thus avoid having to consider a possibly huge number of different equilibria. We would also like to explore giving bridge rules access to the entire belief set of another context or an input and not just to one element. A potential application would be counting, e.g., to ask if there was no sensor reading.

An alternative to deal with inconsistent states is to follow a paraconsistent approach, as proposed for hybrid knowledge bases in [21, 31]. Also, following the idea of EVOLP [2] as explored in [28], we could allow the bridge rules to change with time, strengthening the evolving and adaptation capabilities of rMCSs. We would also like to establish bridges to asynchronous MCSs [19], a framework for loosely coupling knowledge representation formalisms whose semantics assumes that every context delivers output whenever available. Finally, we may build on existing implementations of distributed (static) MCSs [14] for providing an implementation of rMCSs that handles the integration and evolution of distributed, heterogeneous knowledge with incoming streams of data and reasoning over such integration.

The framework introduced in this paper is highly abstract. Needless to say this was our intention. We wanted to capture a wide range of situations without imposing any restrictions on the KR formalisms used. It is a matter of the application to choose the best suited formalisms to solve a specific task. The downside of this generality, of course, is that the framework as such needs to be instantiated in various ways before it is ready-to-use. In particular, we need to select the context logics and define the ways in which contexts and dynamic information interact by specifying adequate bridge rules. We still believe our approach provides a valuable - and to the best of our knowledge unique - solution to the problems outlined in the introduction, that is, problems which originate in the heterogeneous and dynamic nature of the information available in many challenging applications.

Acknowledgements

We would like to thank K. Schekotihin and the anonymous reviewers for their comments, which helped improving this paper. G. Brewka, S. Ellmauthaler, and J. Pührer were partially supported by the German Research Foundation (DFG) under grants BR-1817/7-1/2 and FOR 1513. R. Gonçalves, M. Knorr and J. Leite were partially supported by Fundação para a Ciência e a Tecnologia (FCT) under project NOVA LINES (UID/CEC/04516/2013) and project RIVER (PTDC/CCI-COM/30952/2017)

Moreover, R. Gonçalves was partially supported by FCT grant SFRH/BPD/100906/2014 and M. Knorr by FCT grant SFRH/BPD/86970/2012.

References

- [1] Alferes, J., Leite, J., Pereira, L., Przymusinska, H., Przymusinski, T., 2000. Dynamic updates of non-monotonic knowledge bases. *J. Log. Program.* 45 (1-3), 43–70.
- [2] Alferes, J. J., Brogi, A., Leite, J. A., Pereira, L. M., 2002. Evolving logic programs. In: Flesca, S., Greco, S., Leone, N., Ianni, G. (Eds.), *Logics in Artificial Intelligence, European Conference, JELIA 2002, Cosenza, Italy, September, 23-26, Proceedings*. Vol. 2424 of LNCS. Springer, pp. 50–61.
- [3] Anicic, D., Rudolph, S., Fodor, P., Stojanovic, N., 2012. Stream reasoning and complex event processing in ETALIS. *Semantic Web* 3 (4), 397–407.
- [4] Arenas, M., Bertossi, L. E., Chomicki, J., 1999. Consistent query answers in inconsistent databases. In: Vianu, V., Papadimitriou, C. H. (Eds.), *Proceedings of the Eighteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, May 31 - June 2, 1999, Philadelphia, Pennsylvania, USA*. ACM Press, pp. 68–79.
- [5] Baader, F., Calvanese, D., McGuinness, D. L., Nardi, D., Patel-Schneider, P. F. (Eds.), 2007. *The Description Logic Handbook: Theory, Implementation, and Applications, 2nd Edition*. Cambridge University Press.
- [6] Barbieri, D., Braga, D., Ceri, S., Valle, E. D., Grossniklaus, M., 2010. C-SPARQL: a continuous query language for RDF data streams. *Int. J. Semantic Computing* 4 (1), 3–25.
- [7] Beck, H., Dao-Tran, M., Eiter, T., Fink, M., 2015. LARS: A logic-based framework for analyzing reasoning over streams. In: Bonet, B., Koenig, S. (Eds.), *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*. AAAI Press, pp. 1431–1438.
- [8] Brewka, G., 2013. Towards reactive multi-context systems. In: Cabalar, P., Son, T. C. (Eds.), *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'13)*. Vol. 8148 of LNCS. Springer, pp. 1–10.
- [9] Brewka, G., Eiter, T., 2007. Equilibria in heterogeneous nonmonotonic multi-context systems. In: *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22-26, 2007, Vancouver, British Columbia, Canada*. AAAI Press, pp. 385–390.
- [10] Brewka, G., Eiter, T., Fink, M., Weinzierl, A., 2011. Managed multi-context systems. In: Walsh, T. (Ed.), *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*. IJCAI/AAAI, pp. 786–791.

- [11] Brewka, G., Ellmauthaler, S., Pührer, J., 2014. Multi-context systems for reactive reasoning in dynamic environments. In: ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014). pp. 159–164.
- [12] Brewka, G., Roelofsen, F., Serafini, L., 2007. Contextual default reasoning. In: Veloso, M. M. (Ed.), IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007. pp. 268–273.
- [13] Buccafurri, F., Leone, N., Rullo, P., 1997. Strong and weak constraints in disjunctive datalog. In: Dix, J., Furbach, U., Nerode, A. (Eds.), Logic Programming and Nonmonotonic Reasoning, 4th International Conference, LPNMR'97, Dagstuhl Castle, Germany, July 28-31, 1997, Proceedings. Vol. 1265 of LNCS. Springer, pp. 2–17.
- [14] Dao-Tran, M., Eiter, T., Fink, M., Krennwallner, T., 2015. Distributed evaluation of nonmonotonic multi-context systems. *J. Artif. Intell. Res.* 52, 543–600.
- [15] Eiter, T., Fink, M., Schüller, P., Weinzierl, A., 2014. Finding explanations of inconsistency in multi-context systems. *Artif. Intell.* 216, 233–274.
- [16] Eiter, T., Ianni, G., Lukasiewicz, T., Schindlauer, R., Tompits, H., 2008. Combining answer set programming with description logics for the semantic web. *Artif. Intell.* 172 (12-13), 1495–1539.
- [17] Ellmauthaler, S., 2013. Generalizing multi-context systems for reactive stream reasoning applications. In: Jones, A. V., Ng, N. (Eds.), 2013 Imperial College Computing Student Workshop, ICCSW 2013, September 26/27, 2013, London, United Kingdom. Vol. 35 of OASICS. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, pp. 19–26.
- [18] Ellmauthaler, S., Pührer, J. (Eds.), 2014. Proceedings of the International Workshop on Reactive Concepts in Knowledge Representation (ReactKnow 2014). Leipzig University.
- [19] Ellmauthaler, S., Pührer, J., 2015. Asynchronous multi-context systems. In: Eiter, T., Strass, H., Truszczynski, M., Woltran, S. (Eds.), Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday. Vol. 9060 of LNCS. Springer, pp. 141–156.
- [20] Faber, W., Leone, N., Pfeifer, G., 2004. Recursive aggregates in disjunctive logic programs: Semantics and complexity. In: Alferes, J. J., Leite, J. A. (Eds.), Logics in Artificial Intelligence, 9th European Conference, JELIA 2004, Lisbon, Portugal, September 27-30, 2004, Proceedings. Vol. 3229 of LNCS. Springer, pp. 200–212.

- [21] Fink, M., 2012. Paraconsistent hybrid theories. In: Brewka, G., Eiter, T., McIlraith, S. A. (Eds.), Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10-14, 2012. AAAI Press, pp. 391–401.
- [22] Gebser, M., Grote, T., Kaminski, R., Obermeier, P., Sabuncu, O., Schaub, T., 2012. Stream reasoning with answer set programming: Preliminary report. In: Brewka, G., Eiter, T., McIlraith, S. A. (Eds.), Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10-14, 2012. AAAI Press, pp. 613–617.
- [23] Gebser, M., Grote, T., Kaminski, R., Schaub, T., 2011. Reactive answer set programming. In: Delgrande, J. P., Faber, W. (Eds.), Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings. Vol. 6645 of LNCS. Springer, pp. 54–66.
- [24] Gebser, M., Kaminski, R., Kaufmann, B., Ostrowski, M., Schaub, T., Thiele, S., 2010. A users guide to gringo, clasp, clingo, and iclingo. Potassco Team.
- [25] Gelfond, M., Lifschitz, V., 1991. Classical negation in logic programs and disjunctive databases. *New Generation Comput.* 9 (3/4), 365–386.
- [26] Giunchiglia, F., 1993. Contextual reasoning. *Epistemologia XVI*, 345–364.
- [27] Giunchiglia, F., Serafini, L., 1994. Multilanguage hierarchical logics or: How we can do without modal logics. *Artif. Intell.* 65 (1), 29–70.
- [28] Gonçalves, R., Knorr, M., Leite, J., 2014. Evolving bridge rules in evolving multi-context systems. In: Bulling, N., van der Torre, L. W. N., Villata, S., Jamroga, W., Vasconcelos, W. W. (Eds.), Computational Logic in Multi-Agent Systems - 15th International Workshop, CLIMA XV, Prague, Czech Republic, August 18-19, 2014. Proceedings. Vol. 8624 of LNCS. Springer, pp. 52–69.
- [29] Gonçalves, R., Knorr, M., Leite, J., 2014. Evolving multi-context systems. In: ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014). pp. 375–380.
- [30] Harris, S., Seaborne, A., 2013. SPARQL 1.1 query language for RDF. W3C Recommendation.
- [31] Kaminski, T., Knorr, M., Leite, J., 2015. Efficient paraconsistent reasoning with ontologies and rules. In: Yang, Q., Wooldridge, M. (Eds.), Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015. AAAI Press, pp. 3098–3105.
- [32] Lécué, F., Pan, J. Z., 2013. Predicting knowledge in an ontology stream. In: Rossi, F. (Ed.), IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013. IJCAI/AAAI, pp. 2662–2669.

- [33] Lierler, Y., Truszczyński, M., 2016. On abstract modular inference systems and solvers. *Artif. Intell.* 236, 65–89.
- [34] McCarthy, J., 1987. Generality in artificial intelligence. *Commun. ACM* 30 (12), 1029–1035.
- [35] Motik, B., Rosati, R., 2010. Reconciling description logics and rules. *Journal of the ACM* 57 (5), 93–154.
- [36] Mu, K., Wang, K., Wen, L., 2016. Preferential multi-context systems. *Int. J. Approx. Reasoning* 75, 39–56.
- [37] Özçep, Özgür. L., Möller, R., Neuenstadt, C., Zheleznayakow, D., Kharlamov, E., Horrocks, I., Hubauer, T., Roshchin, M., 2013. D5.1 Executive Summary: A semantics for temporal and stream-based query answering in an OBDA context. Tech. rep., Optique FP7-ICT-2011-8-318338 Project.
- [38] Rodriguez-Muro, M., Jupp, S., Srinivas, K. (Eds.), 2013. The Optique Project: Towards OBDA Systems for Industry (Short Paper). Vol. 1080 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- [39] Roelofsen, F., Serafini, L., 2005. Minimal and absent information in contexts. In: Kaelbling, L. P., Saffiotti, A. (Eds.), *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, Edinburgh, Scotland, UK, July 30-August 5, 2005. Professional Book Center, pp. 558–563.
- [40] Slota, M., Leite, J., Swift, T., 2015. On updates of hybrid knowledge bases composed of ontologies and rules. *Artif. Intell.* 229, 33–104.

Appendix A. Proofs Inconsistency Management

Proposition 1. *Let $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$ be an acyclic rMCS such that every C_i , $1 \leq i \leq n$, is totally coherent, and KB a configuration of knowledge bases for M . Then, M is strongly consistent with respect to KB.*

PROOF. Let $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \langle BR_1, \dots, BR_n \rangle \rangle$ be an acyclic rMCS with totally coherent contexts. We first prove that M has an equilibrium given KB and l , for any knowledge base configuration $\text{KB} = \langle kb_1, \dots, kb_n \rangle$ for M and input l for M .

We prove this by induction on the number of contexts of M , making use of the following simple observation: if M does not have cycles, then there exists some $i \in \{1, \dots, n\}$ such that $\text{ref}_r(j, i)$ does not hold for any $j \in \{1, \dots, n\}$ and $r \in BR_j$, where $\text{ref}_r(j, i)$ holds precisely when r is a bridge rule of context C_j and $i:b$ occurs in the body of r . It is quite easy to see that if this condition is violated then a cycle necessarily exists.

Let $n = 1$. Then, since there are no cycles, no bridge rule in BR_1 contains atoms of the form $1:b$ in its body. Thus, $\text{app}_i^{\text{now}}(l, B)$ does not depend on B . Total coherence then immediately implies that M has an equilibrium given KB and l .

Let $n = m + 1$. We use the above observation, and assume, w.l.o.g., that C_1 is a context for which $\text{ref}_r(j, 1)$ does not hold for any $j \in \{1, \dots, m + 1\}$ and $r \in BR_j$. Then, the rMCS $M^* = \langle \langle C_2, \dots, C_{m+1} \rangle, \text{IL}, \langle BR_2, \dots, BR_{m+1} \rangle \rangle$ has m contexts and it is still acyclic. By induction hypothesis, we can conclude that M^* has an equilibrium given $\text{KB}^* = \langle kb_2, \dots, kb_{m+1} \rangle$ and l . Let $B^* = \langle B_2, \dots, B_{m+1} \rangle$ be such equilibrium. Then, since C_1 is assumed to be a totally coherent context, there exists $B_1 \in BS_1$ such that $B = \langle B_1, B_2, \dots, B_n \rangle$ is an equilibrium of M given KB and l . This follows easily from the fact that no set $\text{app}_i^{\text{now}}(l, B)$ depends on the choice of B_1 .

We have shown that the existence of an equilibrium for M is independent of the given knowledge base configuration KB for M and input l for M . This immediately implies that for any input stream \mathcal{I} for M (until τ), and any knowledge base configuration KB for M , there exists an equilibria stream of M given KB and \mathcal{I} .

Proposition 2. *Every equilibria stream of M given KB and \mathcal{I} is a repaired equilibria stream of M given KB, \mathcal{I} and the empty repair \mathcal{R}_\emptyset .*

PROOF. This result follows easily from the observation that $M[\mathcal{R}_\emptyset] = M$. In this case the conditions in the definition of an equilibria stream of M coincide with those in the definition of a repaired equilibria stream of M .

Proposition 3. *Let $M = \langle \langle C_1, \dots, C_n \rangle, \text{IL}, \text{BR} \rangle$ be an rMCS such that each C_i , $i \in \{1, \dots, n\}$, is totally coherent, KB a configuration of knowledge bases for M , and \mathcal{I} an input stream for M until τ . Then, there exists $\mathcal{R} : [1.. \tau] \rightarrow 2^{\text{br}_M}$ and $\mathcal{B} : [1.. \tau] \rightarrow \text{Bel}_M$ such that \mathcal{B} is a repaired equilibria stream given KB, \mathcal{I} and \mathcal{R} .*

PROOF. Since each context of M is totally coherent, Proposition 1 guarantees the existence of an equilibrium if M is acyclic. Now just note that if we take $\mathcal{R} : [1.. \tau] \rightarrow 2^{\text{br}_M}$ such that $\mathcal{R}^t = \text{br}_M$ for every t , then each $M[\mathcal{R}^t]$ does not have bridge rules and it is therefore acyclic. Then, for every t , $M[\mathcal{R}^t]$ is strongly consistent. Therefore we

can easily inductively construct $\mathcal{B} : [1..\tau] \rightarrow \text{Bel}_M$ such that \mathcal{B} is a repaired equilibria stream given KB , \mathcal{I} and \mathcal{R} .

Proposition 4. *Every equilibria stream of M given KB and \mathcal{I} is a partial equilibria stream of M given KB and \mathcal{I} .*

PROOF. This result follows easily from the observation that for an equilibria stream \mathcal{B} of M given KB and \mathcal{I} , and every t , \mathcal{B}^t is never undefined. Therefore, in this case the conditions in the definition of partial equilibria stream coincide with those for equilibria stream.

Proposition 5. *Let M be an rMCS, KB a configuration of knowledge bases for M , and \mathcal{I} an input stream for M until τ . Then, there exists $\mathcal{B} : [1..\tau] \rightarrow \text{Bel}_M$ such that \mathcal{B} is a partial equilibria stream given KB and \mathcal{I} .*

PROOF. We just need to note that if we take $\mathcal{B} : [1..\tau] \rightarrow \text{Bel}_M$ such that, for every t , \mathcal{B}^t is undefined, then \mathcal{B} is trivially a partial equilibria stream given KB and \mathcal{I} .

Appendix B. Proofs Non-Determinism and Well-Founded Semantics

We first show some properties of the iteration defined in Def. 23 that will prove useful subsequently.

Lemma 1. *Let M be a definite rMCS given KB , and l an input for M . The following holds for the iteration in Definition 23 for all ordinals α :*

1. M is definite given KB^α ;
2. for any ordinal β with $\beta \leq \alpha$ we have $kb_i^\beta \subseteq kb_i^\alpha$ for $i \in \{1, \dots, n\}$.

PROOF. We have to show that the lemma holds for all ordinals α . For the (initial) limit ordinal $\alpha = 0$, Claims 1 and 2 hold trivially.

Now, suppose we have shown that 1 and 2 hold for all ordinals $\alpha_1 \leq \alpha$. First, consider a successor ordinal $\alpha + 1$. Regarding Claim 1, i.e., M being definite given $\text{KB}^{\alpha+1}$, Condition 1 of Definition 22 holds trivially as the bridge rules are fix, and Condition 2 follows from the fact that all kb_i^α are already in reduced form, and from Condition 2 of Definition 21, which prevents the introduction of reducible content by any mng_i . Regarding Claim 2, we already know by the hypothesis that this holds for α , so we only have to show that $kb_i^\alpha \subseteq kb_i^{\alpha+1}$ holds for all i as well. As all acc_i for the definite rMCS are monotonic, we also know that $\text{B}^\beta \subseteq \text{B}^\alpha$ holds for all $\beta \leq \alpha$. In fact, this holds for any of the β just as well, i.e., there is an increasing sequence of B^β associated to the iteration. Now, since the rMCS is definite, hence no negation occurs in bridge rules, and l is fix, the sequence of all $\text{app}_i^{\text{now}}(l, \text{KB}^\beta)$ for each β is also increasing. But then, by Condition 3 of Definition 21, $kb_i^\alpha \subseteq kb_i^{\alpha+1}$ holds for all i .

Now consider a limit ordinal α' (with $\alpha \leq \alpha'$). Since in this case, all $kb_i^{\alpha'}$ are just the union of all kb_i^β with $\beta \leq \alpha'$, the Claims 1 and 2 trivially follow.

With the help of Lemma 1, we can show that the iteration indeed yields a unique minimal equilibrium (of M given KB and l).

Proposition 6. *Let M be a definite rMCS given KB, and l an input for M . Then $\langle B_1^\infty, \dots, B_n^\infty \rangle$ is the unique minimal equilibrium of M given KB and l . We call it grounded equilibrium of M given KB and l , denoted by $\mathbf{GE}(M, \text{KB}, l)$.*

PROOF. We first note that $\langle B_1^\infty, \dots, B_n^\infty \rangle$ is indeed an equilibrium. This follows directly from the definition of equilibria (Definition 9) and the fact that the fixpoint in Definition 23 precisely matches it.

Concerning minimality, suppose it is not a minimal equilibrium. Then, there is some $\langle B', \dots, B' \rangle$ which is also an equilibrium such that $B'_i \subseteq B_i^\infty$ for all $i \in \{1, \dots, n\}$ and $B'_i \subset B_i^\infty$ for at least one i . Therefore some belief $b \in B_i^\infty \setminus B'_i$ and a least ordinal α exist such that $b \in B_i^\alpha$. We consider a b where α is least for all such beliefs. If $\alpha = 0$, then $\mathbf{acc}_i(kb_i^0) = B_i^0$ and since, by Lemma 1, M is definite given KB^0 , this belief set is unique, hence B' cannot be an equilibrium (following from monotonicity of the iteration). If $\alpha > 0$, then necessarily some operation in a bridge rule head, applied when creating kb_i^α , triggered the occurrence of b in $\mathbf{acc}_i(kb_i^\alpha) = B_i^\alpha$. Now, as M is definite given all KB^β in the iteration of KB^∞ and since α is least, this head is also applicable by monotonicity w.r.t. B' . Hence, we obtain a contradiction to B' being an equilibrium.

Regarding uniqueness, suppose it is a minimal equilibrium, but not unique. Then, there is some $\langle B', \dots, B' \rangle$ which is also a minimal equilibrium such that at least one of the following holds:

- neither $B'_i \subseteq B_i^\infty$ nor $B_i^\infty \subseteq B'_i$ for at least one $i \in \{1, \dots, n\}$;
- $B'_i \subset B_i^\infty$ and $B'_j \supset B_j^\infty$ for $i \neq j$ with $i, j \in \{1, \dots, n\}$.

In both cases, consider some $b \in B_i^\infty \setminus B'_i$. Again, there is a least α such that $b \in B_i^\alpha$ and α is least among all these b . We can apply the same argument as used for proving minimality and obtain a contradiction to B' being an equilibrium.

Proposition 7. *Every grounded equilibrium of a reducible rMCS M given KB and an input l is a minimal equilibrium of M given KB and l .*

PROOF. We first show that a grounded equilibrium B is indeed an equilibrium. For this, note that B is the unique minimal equilibrium of the definite $M^{(l, B)}$ given $\text{KB}^{(l, B)}$ and l by Definition 25 and Proposition 6. Thus, by Definition 9 and since $M^{(l, B)}$ is definite, we know for all $i \in \{1, \dots, n\}$:

$$B_i = \mathbf{acc}_i(kb'), \text{ where } kb' = \mathbf{mng}_i(\mathbf{app}_i^{\text{now}}(l, B), kb_i^{B_i}). \quad (\text{B.1})$$

with $kb_i^{B_i} = \mathbf{red}_i(kb_i, B_i)$ by Definition 24. Now, according to Definition 24, $M^{(l, B)}$ only differs from M on the sets of bridge rules $BR_i^{(l, B)}$, and the way these are obtained from the BR_i ensures that $\mathbf{app}_i^{\text{now}}(l, B)$ is identical for both M and $M^{(l, B)}$. Moreover, by 2. of Definition 21, (B.1) can be rewritten to (for all i):

$$B_i = \mathbf{acc}_i(kb'), \text{ where } kb' = \mathbf{red}_i(\mathbf{mng}_i(\mathbf{app}_i^{\text{now}}(l, B), kb_i), B_i). \quad (\text{B.2})$$

In this case, by the definition of \mathbf{red} , we know that, (for all i),

$$B_i \in \mathbf{acc}_i(\mathbf{mng}_i(\mathbf{app}_i^{\text{now}}(l, B), kb_i)) \quad (\text{B.3})$$

and this shows, by Definition 9, that B is an equilibrium.

Now suppose that B is not minimal. Then there is an equilibrium $B' = \langle B'_1, \dots, B'_n \rangle$ of M given KB and I such that $B'_i \subseteq B_i$ for all i with $1 \leq i \leq n$ and $B'_j \subsetneq B_j$ for some j with $j \in \{1, \dots, n\}$. Since \mathbf{red} is antitonic by definition, we know that $\mathbf{red}_i(kb_i, B_i) \subseteq \mathbf{red}_i(kb_i, B'_i)$ holds for all $i \in \{1, \dots, n\}$. Also, by Definition 24, $BR_i^{(I, B^\alpha)} \subseteq BR_i^{(I, B'^\alpha)}$ for all i . It can thus be shown by induction on α for the monotonic iteration in Definition 23 that $B_i^\alpha \subseteq B'^\alpha_i$ holds for all $i \in \{1, \dots, n\}$:

- $\alpha = 0$: this holds right away by monotonicity of \mathbf{acc}_i for reduced knowledge bases.
- Suppose the claim holds for all ordinals $\alpha_1 \leq \alpha$.
- Consider a successor ordinal $\alpha + 1$. As $BR_i^{(I, B^\alpha)} \subseteq BR_i^{(I, B'^\alpha)}$ holds, and no bridge rule contains **not** due to the reduction, we have, by the induction hypothesis that $\mathbf{app}_i^{now}(I, B^\alpha) \subseteq \mathbf{app}_i^{now}(I, B'^\alpha)$. Then, since \mathbf{mng}_i in the iteration cannot remove beliefs (see 3. of Definition 21), we obtain $kb_i^{\alpha+1} \subseteq kb_i'^{\alpha+1}$ and thus $B_i^{\alpha+1} \subseteq B_i'^{\alpha+1}$.
- Consider a limit ordinal α' with $\alpha \leq \alpha'$. As the corresponding knowledge bases are simply the union of those of all smaller ordinals, we necessarily have $kb_i^{\alpha'} \subseteq kb_i'^{\alpha'}$, and thus $B_i^{\alpha'} \subseteq B_i'^{\alpha'}$.

Since we already know that B is a grounded equilibrium, we have that $B_i = B_i^\infty$ for all i . We conclude that $B'_j \subset B_j^\infty$ holds for at least one $j \in \{1, \dots, n\}$. Thus, B' itself cannot be a grounded equilibrium of M given KB and I by Definition 25.

So suppose that B' is an equilibrium, but not grounded. Then we know that, by Definition 9, for all $i \in \{1, \dots, n\}$:

$$B'_i \in \mathbf{acc}_i(kb'), \text{ where } kb' = \mathbf{mng}_i(\mathbf{app}_i^{now}(I, B'), kb_i). \quad (\text{B.4})$$

In this case, by the definition of \mathbf{red} , we have for all i :

$$\mathbf{acc}_i(\mathbf{red}_i(kb', B'_i)) = \{B'_i\} \quad (\text{B.5})$$

Therefore, by 2. of Definition 21:

$$\mathbf{acc}_i(\mathbf{mng}_i(\mathbf{app}_i^{now}(I, B'), \mathbf{red}_i(kb_i, B'_i))) = \{B'_i\} \quad (\text{B.6})$$

Now, the kb_i^0 used in the iteration of B_i^α are precisely the $\mathbf{red}_i(kb_i, B'_i)$ for all i . As argued in the proof of Lemma 1, the sequence of operations in the iteration is monotonically increasing. We consider two cases. First, the sequence reaches $\mathbf{app}_i^{now}(I, B')$. Then, for all $OP^j \subseteq \mathbf{app}_i^{now}(I, B')$, $\mathbf{acc}_i(\mathbf{mng}_i(OP^j, kb_i^\alpha)) \subseteq B'_i$. As $B'_j \subset B_j^\infty$ holds for at least one $j \in \{1, \dots, n\}$, the iteration has to continue beyond the step involving $\mathbf{app}_i^{now}(I, B')$, but this contradicts B' being an equilibrium. Alternatively, the sequence does not reach a step with precisely the operations in $\mathbf{app}_i^{now}(I, B')$. In this case, there is a least ordinal such that an operation not contained in $\mathbf{app}_i^{now}(I, B')$ occurs in the iteration, but this again contradicts that B' is an equilibrium following the argument applied in the proof of Proposition 6.

Proposition 8. *Let $M = \langle C, IL, BR \rangle$ be a normal, reducible rMCS given KB, I an input for M , $WF(M, KB, I) = \langle W_1, \dots, W_n \rangle$, and $B = \langle B_1, \dots, B_n \rangle$ a grounded equilibrium of M given KB and I . Then $W_i \subseteq B_i$ for $i \in \{1 \dots n\}$.*

PROOF. We show the claim by proving that $(\gamma_{M,KB,I})^2 \uparrow \alpha)_i \subseteq B_i$ holds for all grounded equilibria B of M given KB and I . For the case $\alpha = 0$, this holds trivially, as the initial belief state is B^* .

Suppose the claim holds for all ordinals $\alpha_1 \leq \alpha$. Consider a successor ordinal $\alpha + 1$. If $(\gamma_{M,KB,I})^2 \uparrow \alpha = (\gamma_{M,KB,I})^2 \uparrow (\alpha + 1)$ then the claim trivially holds by the induction hypothesis. Thus, suppose that $(\gamma_{M,KB,I})^2 \uparrow \alpha \subset (\gamma_{M,KB,I})^2 \uparrow (\alpha + 1)$, i.e., new beliefs are added in the step $\alpha + 1$ for (at least) some i . Using Definition 23 and the induction hypothesis, it can be shown that these new beliefs are indeed contained in each grounded equilibrium, while all beliefs that do not occur in the intermediate fixpoint of the squared operator are not true in any grounded equilibrium of M given KB and I , in the very same way as usual for the alternating fixpoint, as e.g., for logic programs, on which this construction is based.

Finally, regarding a limit ordinal α' with $\alpha \leq \alpha'$, the claim holds trivially by the induction hypothesis.

Appendix C. Proofs Complexity

Theorem 1. *Table 1 summarizes the complexities of membership of problems Q^\exists and Q^\forall for finite input streams (until some $\tau \in \mathbb{N}$) depending on the context complexity. Hardness also holds if it holds for the context complexity.*

PROOF. The membership results for the Q^\exists cases (with the exception of $\mathcal{CC}(M, k:b) = \text{EXPTIME}$) can be argued for as follows: a non-deterministic Turing machine can be used to guess a projected belief state $\mathcal{B}^t = \langle B_1, \dots, B_n \rangle$ for all τ inputs in \mathcal{I} in polynomial time. Then, iteratively for each of the consecutive inputs \mathcal{I}^t , first the context problems can be solved either polynomially or using an oracle for the context complexity (the guess of \mathcal{B}^t and the oracle guess can be combined which explains why we stay on the same complexity level for higher context complexity). If the answer is 'yes', \mathcal{B}^t is the projected equilibrium. We can check whether $b \in B_i$, compute the updated knowledge bases and continue the iteration until reaching the last input. For **PSPACE** the same line of argumentation holds as **PSPACE** = **NPSPACE**. In the case of $\mathcal{CC}(M, k:b) = \text{EXPTIME}$, we iterate through the exponentially many projected belief states for which we solve the context problem in exponential time and proceed as before. The argument is similar for the co-problem of Q^\forall . Hardness holds because being able to solve Q^\exists , respectively the co-problem of Q^\forall , one can decide equilibrium existence for managed MCSs which is hard for the same complexity classes [10] given hardness for the context complexity of the managed MCS.

Theorem 2. *Let $M = \langle C, IL, BR \rangle$ be a finite, persistently reducible rMCS given KB. Then membership and hardness results for Q_g^\exists and Q_g^\forall on grounded equilibria streams for finite input streams coincide with those for Q^\exists and Q^\forall in Theorem 1.*

PROOF. As M is assumed to be finite and persistently reducible, the argument is exactly identical to that of Theorem 1. The only difference is here that we also have to guess the intermediate belief states in the iteration, but this does not raise the (worst-case) complexity itself.

Theorem 3. *Let $M = \langle C, \text{IL}, \text{BR} \rangle$ be a finite, normal, persistently reducible rMCS given KB such that $\text{CC}(M, k:b) = \mathbf{P}$ for Q_{wf}^\exists . Then, Q_{wf}^\exists is in \mathbf{P} . In addition, hardness holds provided $\text{CC}(M, k:b) = \mathbf{P}$ is hard.*

PROOF. If $\text{CC}(M, k:b) = \mathbf{P}$, then deciding all context problems is polynomial. Also, no guessing of equilibria is required, as each equilibrium can be computed iteratively in polynomial time starting from the belief state containing all least elements of each belief set. Since the context complexity also determines the complexity of performing the update function to obtain the new configuration of knowledge bases, we conclude that the result holds. The result for hardness follows in the same manner as in the proofs of Theorems 1 and 2.

Finally, to demonstrate the expressiveness of rMCSs and prove Proposition 9 we provide an rMCS M_{TM} that implements a service that reads the configuration of a Turing machine (TM) from external sources and simulates a run of the TM on request. The fixed rMCS M_{TM} comes with inexpressive context logics and management functions computable in linear time, thus showing that Turing-completeness arises from the interplay between the contexts over time. We use the following variant of TMs, where \mathbb{S} is the set of all tape symbols and \mathbb{Q} is the set of all states:

Definition 31. *A TM is a tuple $\langle Q, \Gamma, \sqcup, \Sigma, \delta, q_0, F \rangle$, where*

- $Q \subseteq \mathbb{Q}$ is a finite, non-empty set of states,
- $\Gamma \subseteq \mathbb{S}$ is a finite, non-empty set of tape symbols,
- $\sqcup \in \Gamma$ is the blank symbol,
- $\Sigma \subseteq \Gamma$ is the set of input symbols,
- q_0 is the initial state,
- $F \subseteq Q$ is the set of final states, and
- $\delta : Q \setminus F \times \Gamma \rightarrow Q \times \Gamma \times \{\leftarrow, \rightarrow\}$ is the (partial) transition function.

We assume familiarity with the computation model of TMs, in particular what it means that a TM halts and accepts an input word $w \in \Sigma^*$. M_{TM} is an rMCS with four contexts. Context C_t simulates a tape of a TM, C_q contains information about TM states, C_f encodes a transition function, and C_c is a control context for operating the TM simulation and presenting results. All contexts use a storage logic as in Example 1, where the set of entries is respectively given by

- $E_t = \bigcup_{p \in \mathbb{Z}, s \in \mathbb{S}} \{t(p, s), \text{curP}(p)\},$

- $E_q = \bigcup_{q \in \mathbb{Q}} \{\text{final}(q), \text{curQ}(q)\}$,
- $E_f = \{f(q, s, q', s', m) \mid q, q' \in \mathbb{Q}, s, s' \in \mathbb{S}, m \in \{L, R\}\}$, and
- $E_c = \{\text{computing}, \text{answer}(yes), \text{answer}(no)\}$.

The input of M_{TM} is provided by four streams over the languages

- $IL_t = E_t$,
- $IL_q = E_q$,
- $IL_f = E_f$, and
- $IL_c = \{\text{start}, \text{reset}\}$,

where IL_t allows for setting an initial configuration for the tape, IL_q the allowed and final states of the TM to simulate, IL_f the transition function, and input over IL_c is used to start and reset the simulation.

The bridge rule schemata of the tape context C_t are given by:

$$\begin{aligned}
\text{next}(\text{add}(t(P, S'))) &\leftarrow q:f(Q, S, Q', S', D), t:\text{curP}(P), q:\text{curQ}(Q), \\
&\quad t:t(P, S), S \neq S', c:\text{computing}. \\
\text{next}(\text{rm}(t(P, S))) &\leftarrow q:f(Q, S, Q', S', D), t:\text{curP}(P), q:\text{curQ}(Q), \\
&\quad t:t(P, S), S \neq S', c:\text{computing}. \\
\text{add}(\text{nextP}(P - 1)) &\leftarrow q:f(Q, S, Q', S', \leftarrow), q:\text{curQ}(Q), \\
&\quad t:t(P, S), t:\text{curP}(P), c:\text{computing}. \\
\text{add}(\text{nextP}(P + 1)) &\leftarrow q:f(Q, S, Q', S', \rightarrow), q:\text{curQ}(Q), \\
&\quad t:t(P, S), t:\text{curP}(P), c:\text{computing}. \\
\text{add}(\text{nextPdefined}) &\leftarrow t:\text{nextP}(P), c:\text{computing}. \\
\text{next}(\text{add}(\text{curP}(P))) &\leftarrow t:\text{nextP}(P), c:\text{computing}. \\
\text{next}(\text{rm}(\text{curP}(P))) &\leftarrow t:\text{curP}(P), t:\text{nextPdefined}, c:\text{computing}. \\
\text{next}(\text{add}(X)) &\leftarrow t::X, \text{not } c:\text{computing}. \\
\text{next}(\text{clear}) &\leftarrow c::\text{reset}.
\end{aligned}$$

The bridge rule schemata for the state context C_q are the following:

$$\begin{aligned}
\text{next}(\text{add}(\text{curQ}(Q'))) &\leftarrow q:f(Q, S, Q', S', D), t:\text{curP}(P), q:\text{curQ}(Q), \\
&\quad t:t(P, S), Q \neq Q', c:\text{computing}. \\
\text{next}(\text{rm}(\text{curQ}(Q))) &\leftarrow q:f(Q, S, Q', S', D), t:\text{curP}(P), q:\text{curQ}(Q), \\
&\quad t:t(P, S), Q \neq Q', c:\text{computing}. \\
\text{next}(\text{add}(X)) &\leftarrow q::X, \text{not } c:\text{computing}. \\
\text{next}(\text{clear}) &\leftarrow c::\text{reset}.
\end{aligned}$$

The state C_f for the transition function has the bridge rules schemata given next:

$$\begin{aligned}
\text{next}(\text{add}(X)) &\leftarrow f::X, \text{not } c:\text{computing}. \\
\text{next}(\text{clear}) &\leftarrow c::\text{reset}.
\end{aligned}$$

Finally, the schemata for C_c are:

$$\begin{aligned}
\text{add}(\text{answer}('Y')) &\leftarrow q:\text{curQ}(Q), q:\text{finalQ}(Q), c:\text{computing}. \\
\text{add}(\text{answer}('N')) &\leftarrow \mathbf{not} t:\text{nextPdefined}(Q), q:\text{curQ}(Q), \\
&\quad \mathbf{not} q:\text{finalQ}(Q), c:\text{computing}. \\
\mathbf{next}(\text{add}(\text{answer}(X))) &\leftarrow c:\text{answer}(X), c:\text{computing}. \\
\mathbf{next}(\text{rm}(\text{computing})) &\leftarrow c:\text{answer}(X), c:\text{computing}. \\
\mathbf{next}(\text{clear}) &\leftarrow c::\text{reset}. \\
\mathbf{next}(\text{add}(\text{computing})) &\leftarrow c::\text{start}.
\end{aligned}$$

All contexts use the following management function:

$$\mathbf{mng}(OP, kb) = \begin{cases} \emptyset & \text{if clear} \in OP \\ kb \setminus \{X \mid \text{rm}(X) \in OP\} \cup \{X \mid \text{add}(X) \in OP\} & \text{else} \end{cases}$$

Let $T = \langle Q, \Gamma, \vdash, \Sigma, \delta, q_0, F \rangle$ be a TM and $w \in \Sigma^*$ an input word for T . We want to use M_{TM} with input stream \mathcal{I} to simulate T . Assume we start at time t . We first make sure that all knowledge bases are empty by setting $\mathcal{I}_c^t = \{\text{reset}\}$. This activates the bridge rules in all contexts of M_{TM} that derive $\mathbf{next}(\text{clear})$. As a consequence, at time $t + 1$ the contents of all knowledge bases are deleted. Next, we feed T and w to M_{TM} by sending

- $\text{final}(q)$ for all $q \in F$ and $\text{curQ}(q_0)$ on the input stream q ,
- $f(q, s, q', s')$ iff $\delta(q, s) = \langle q', s' \rangle$ on stream f , and
- $\text{curP}(\theta)$ and $t(p, s)$ iff $s = s_p$ for $w = s_0, s_1, s_2, \dots$ on the tape stream t .

Note that it does not matter whether we do this all at once at time $t + 1$ or scattered over multiple time points greater than t . Assume that we finished to incorporate all this information to the knowledge bases at time t' . Then, we set $\mathcal{I}_c^j = \{\text{start}\}$ to initiate the simulation of T . At time $t' + 1$ the entry computing is contained in the knowledge base of context C_c , activating the bridge rules in all contexts that are responsible for the simulation. From now on, depending on the current state $\text{curQ}(q)$ and the transition function, the bridge rules of tape context C_t always change the content of the tape on the current tape position indicated by $\text{curP}(p)$. A new position p' of the tape head indicated by the transition function is reflected by deriving $\text{nextP}(p')$. If such a belief is in the equilibrium so is nextPdefined and $\text{curP}(p')$ is added at the next time point. For context C_q the current state is updated according to the transition function. Note that the auxiliary belief nextPdefined is also used in the bridge rules of context C_c for indicating that if the current state is not final and the transition function is undefined for the current state and input symbol, then the answer of the TM is no, indicated by $\text{answer}('N')$. Conversely, if we arrive at a final state then $\text{answer}('Y')$ is derived. If T does not halt on input w , then also the simulation in M_{TM} will continue forever, unless we stop the computation by sending reset on input stream \mathcal{I}_c once more.

Appendix D. Comparative Studies

In this part, selected related approaches are compared to rMCS in more detail by demonstrating how they can be implemented in our framework.

Appendix D.1. Reactive Multi-Context Systems and Stream Reasoning

In the following, we relate rMCSs to two the frameworks for stream reasoning discussed in Section 7.1 by demonstrating how tasks suited for these approaches can be solved by an rMCS. Here, we focus at presenting the intuition on how problems can be modeled in the different approaches rather than presenting formal translations between them. As the involved frameworks are quite general, it would only make sense to look at technical mappings for well-defined restrictions of the general settings, which would limit their value for a general comparison. Moreover, such formal translations would have to mainly deal with bridging technical differences in semantics (e.g., between the equilibrium semantics of rMCS and the FLP semantics used in [7]) rather than giving insight into how we can model stream reasoning tasks in the respective approaches.

LARS

The LARS language [7] is built over a set of atoms \mathcal{A} defined over disjoint sets of predicates \mathcal{P} and constants \mathcal{C} as usual, where \mathcal{P} is divided into two disjoint subsets, the *extensional predicates* $\mathcal{P}^{\mathcal{E}}$, and the *intensional predicates* $\mathcal{P}^{\mathcal{I}}$. The former is intended to be used for input streams, and the latter for intermediate and output streams.

Given $i, j \in \mathbb{N}$, an *interval* is a set of the form $[i, j] = \{k \in \mathbb{N} \mid i \leq k \leq j\}$. An *evaluation function* over an interval T is a function $v : \mathbb{N} \rightarrow 2^{\mathcal{A}}$ such that $v(t) = \emptyset$ if $t \notin T$. Then, a *stream* in LARS is a tuple $s = \langle T, v \rangle$, where T is an interval and v is an evaluation function over T . A stream is a *data stream* if it contains only extensional atoms. A stream $S = \langle T, v \rangle$ is a *substream* of a stream $S' = \langle T', v' \rangle$, denoted by $S \subseteq S'$, if $T \subseteq T'$ and $v(t) \subseteq v'(t)$ for all $t \in T$.

The formulas of LARS are defined using the following grammar:

$$\alpha := a \mid \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \alpha \rightarrow \alpha \mid \diamond\alpha \mid \square\alpha \mid @_t\alpha \mid \boxplus_{\iota}^{\mathbf{x}}\alpha$$

where $a \in \mathcal{A}$ and $t \in \mathbb{N}$. The connectives \neg, \wedge, \vee and \rightarrow are the usual classical connectives. Similarly to modal logics, the operators \diamond and \square are used to represent that a formula holds for some and for each time instant within some interval for a given stream, respectively. The exact operator $@_t$ represents the reference to some specific time $t \in \mathbb{N}$, and window operators of the form $\boxplus_{\iota}^{\mathbf{x}}\alpha$ allow focusing on more recent substreams, where ι represents the type of window and \mathbf{x} the tuple of parameters for the respective type. Among the operators presented in [7] are, for example, time-based operators \boxplus_r^n , which allow focusing on the last n time instants of a given stream, or the partition-based operators $\boxplus_p^{\text{id}, n}$, which first split the stream into substreams and then allow focusing on the last n tuples of a particular substream. LARS programs are based on rules composed of such LARS formulas in a way similar to those in logic programming. More precisely, a LARS rule is of the form

$$\alpha \leftarrow \beta_1, \dots, \beta_j, \text{not } \beta_{j+1}, \dots, \text{not } \beta_n$$

where $\alpha, \beta_1, \dots, \beta_n$ are formulas and α contains only intentional predicates.

Consider a simplified version of the example in [7] that models a scenario where we want to reason about a tram network, including, for example, the prediction of the expected arrival time of a tram at some stop. The key idea is that such information should not only depend on a fixed timetable, but also dynamically on real-time information about the arrival time at a previous stop and the expected travel time between stations (which heavily depends on real-time traffic jam information). LARS allows such combination of static knowledge with a stream of information. Here, static knowledge is composed of atoms of the form $\text{plan}(L, X, Y, T)$, where L is the line identifier, X and Y are consecutive tram stops on line L , and T is the expected travel time between X and Y , and atoms of the form $\text{line}(Id, L)$ mean that the tram Id operates on line L . The stream of information contains atoms of the form $\text{tram}(Id, X)$ meaning that tram Id is at stop X , and atoms of the form $\text{jam}(X)$ meaning that there is a traffic jam near station X . Consider the following example of a LARS rule:

$$\begin{aligned} @_T \text{exp}(Id, Y) \leftarrow \boxplus_p^{\text{idx},1} @_{T_1} \text{tram}(Id, X), \text{line}(Id, L), \mathbf{not} \boxplus_r^{20} \diamond \text{jam}(X), \\ \text{plan}(L, X, Y, Z), T = T_1 + Z. \end{aligned}$$

The intuitive idea of the above rule is that the tram is expected to arrive at stop Y at time T whenever five conditions are met: *i*) the fact that tram Id stopped at X at time T_1 is the last information about stops of Id in the stream; *ii*) X and Y are consecutive stops of line L ; *iii*) $T - T_1$ is the travel time between X and Y ; *iv*) the tram Id is operating on line L ; and *v*) there is no information about traffic jams near station X within the last 20 minutes. Note that the use of default negation \mathbf{not} allows for the representation of exceptions, whereas the partition-based operator $\boxplus_p^{\text{idx},1}$ allows for the focus on the last item of the form $\text{tram}(X)$ in the stream, and the time-based operator \boxplus_r^{20} allows for the focus on the last 20 time instants of the stream.

The semantics of LARS is based on the FLP semantics of logic programs [20]. Given an input stream $D = \langle T, v \rangle$ and time point $t \in T$, each LARS program P is associated with a set of streams, the *answer streams* of P for D at t . Since the semantics of a LARS program is defined for a fixed input data stream and for a particular time point, it is in fact mainly static.

We now describe how rMCSs can be used to run a LARS program over a (possibly infinite) input stream. The idea is to define an rMCS with a LARS context and a recent substream of a (possibly infinite) input stream (in the sense of rMCSs). At each time point, the knowledge base of the LARS context contains the LARS program and the relevant substream of the input stream. Then, the answer streams of the program given the available data stream and the current time point can be computed.

More precisely, we assume fixed sets of predicates \mathcal{P} and constants \mathcal{C} , a fixed window size $w \in \mathbb{N}$ and a LARS program P over \mathcal{A} , the set of atoms obtained from \mathcal{P} and \mathcal{C} . Let $\mathcal{A}^\mathcal{E}$ be the set of atoms that include only extensional predicates from $\mathcal{P}^\mathcal{E}$, and \mathcal{A}^T be the set of time-tagged atoms, i.e., $\mathcal{A}^T = \{\langle a, t \rangle \mid a \in \mathcal{A} \text{ and } t \in \mathbb{N}\}$.

Consider the rMCS $M = \langle \langle C_{\text{LARS}} \rangle, \langle IL_1, \text{Clock} \rangle, \langle BR_{C_{\text{LARS}}} \rangle \rangle$ obtained from P and w in the following way:

- $C_{\text{LARS}} = \langle L, OP, \mathbf{mng} \rangle$ where

- $L = \langle KB, BS, \mathbf{acc} \rangle$ is such that
- $KB = \{P \cup A \cup \{\text{now}(t)\} \mid A \subseteq \mathcal{A}^T \text{ and } t \in \mathbb{N}\}$
- $BS = \{S \mid S \text{ is a stream for } \mathcal{A}\}$
- $\mathbf{acc}(kb) = \{S \mid S \text{ is an answer stream of } P \text{ for } D^{kb} \text{ at time } t^{kb}\}$
 - $t^{kb} = t$ such that $\text{now}(t) \in kb$
 - $T^{kb} = [t^{kb} - w, t^{kb}]$
 - $v^{kb}(t) = \{a \mid \langle a, t \rangle \in kb\}$, for each $t \in \mathbb{N}$
 - $D^{kb} = \langle T^{kb}, v^{kb} \rangle$
- $OP = \{\text{add}(\delta) \mid \delta \in \mathcal{A}^T\} \cup \{\text{del}(\delta) \mid \delta \in \mathcal{A}^T\} \cup \{\text{add}(\text{now}(t)) \mid t \in \mathbb{N}\}$
- $\mathbf{mng}(kb, op) = (kb \cup \{\delta \mid \text{add}(\delta) \in op\}) \setminus \{\delta \mid \text{del}(\delta) \in op\}$
- $IL_1 = \mathcal{A}^\mathcal{E}$
- $Clock = \mathbb{N}$
- BR_{CLARS} contains the following rules for managing history:
 - $\text{add}(\text{now}(T)) \leftarrow \text{clock}::T$
 - $\text{add}(\langle A, T \rangle) \leftarrow 1::A, \text{clock}::T$
 - $\text{del}(\langle A, T \rangle) \leftarrow \text{clock}::T', T < T' - w$
 - $\mathbf{next}(\text{add}(\langle A, T \rangle)) \leftarrow 1::A, \text{clock}::T$
 - $\mathbf{next}(\text{del}(\langle A, T \rangle)) \leftarrow \text{clock}::T', T \leq T' - w$

Given an input stream \mathcal{I} for M and a time point $t \in \mathbb{N}$, we consider $t_t^\mathcal{I}$, the unique element of stream $Clock$ at step t , which represents the current time at step t . We also consider the LARS input data stream at time t , $D_t^\mathcal{I} = \langle T, v \rangle$, such that $T = [t_t^\mathcal{I} - w, t_t^\mathcal{I}]$ and $v(t') = \{a \in \mathcal{A}^\mathcal{E} \mid \text{there exists } t'' \leq t \text{ such that } t' = t_{t''}^\mathcal{I}, \text{ and } a \in \mathcal{I}_1^{t''}\}$ for $t' \in T$, and $v(t') = \emptyset$ otherwise. Then, given an input stream \mathcal{I} for M , at each time point $t \in \mathbb{N}$, each equilibria stream \mathcal{B} for M given $KB = \langle \{P\} \rangle$ and \mathcal{I} is composed of an answer stream of P for $D_t^\mathcal{I}$ at time $t_t^\mathcal{I}$.

Note that at each time instant the knowledge base contains only the relevant part of the (possibly infinite) input stream, meaning that information no longer valid is really discarded, and that the current time, given by the stream $Clock$, is decoupled from the time steps at which equilibria are evaluated. For the sake of presentation, we have assumed a fixed time window w , yet an extension in the spirit of what we presented in Section 3.5 can easily be imagined.

STARQL

Streams in the STARQL framework [37] come in the form of timestamped Description Logic assertions (called ABox assertions). Both, input as well as answers of STARQL queries are streams of this kind. A STARQL select expression is structured as follows.

```
SELECT selectClause( $\vec{x}, \vec{y}$ )
FROM listOfWindowedStreamExpressions
USING listOfResources
WHERE  $\Psi(\vec{x})$ 
SEQUENCE BY seqMethod
HAVING  $\Phi(\vec{x}, \vec{y})$ 
```

For a comprehensive description of syntax and semantics of STARQL, we refer to Optique Deliverable 5.1 [37]. We make use of an example (partly taken from [37]) to explain components of STARQL queries and to describe the core aspects that are relevant for us. It is based on the following STARQL query that creates an output stream `S_out` that indicates temperature sensors whose readings were monotonically increasing for the past two seconds.

```
CREATE STREAM S_out AS

SELECT { ?sens rdf:type MonIncTemp } <NOW>
FROM S 0s <-[NOW-2s, NOW]->1s
USING  STATIC ABOX <http://example.org/Astatic>,
        TBOX <http://example.org/TBox>
WHERE { ?sens rdf:type TempSensor }
SEQUENCE BY StdSeq AS SEQ1
HAVING FORALL i <= j in SEQ1 , x, y:
        IF ( { ?sens rd ?x } <i> AND { ?sens rd ?y } <j> )
        THEN ?x <= ?y
```

The considered input stream contains readings of sensors s_0, s_1 :

$$S = \{ rd(s_0, 90) \langle 0s \rangle, \\ rd(s_1, 30) \langle 0s \rangle, \\ rd(s_0, 93) \langle 1s \rangle, \\ rd(s_1, 32) \langle 1s \rangle, \\ rd(s_0, 94) \langle 2s \rangle, \\ rd(s_0, 91) \langle 3s \rangle, \\ rd(s_0, 93) \langle 4s \rangle, \\ rd(s_0, 95) \langle 5s \rangle \}$$

In addition, the TBox at `http://example.org/TBox` contains the axiom

$$BurnerTipTempSensor \sqsubseteq TempSensor$$

stating that every *BurnerTipTempSensor* is a temperature sensor, and the ABox at `http://example.org/Astatic` contains the assertion

$$BurnerTipTempSensor(s_0)$$

stating that s_0 is of type *BurnerTipTempSensor*. The other sensor s_1 is thus not (derivable to be) a temperature sensor.

Taking S as input, the query returns the output stream, represented here as a sequence of timestamped ABox assertions:

$$S_{out} = \{ \text{MonIncTemp}(s_0)\langle 0s \rangle \\ \text{MonIncTemp}(s_0)\langle 1s \rangle \\ \text{MonIncTemp}(s_0)\langle 2s \rangle \\ \text{MonIncTemp}(s_0)\langle 5s \rangle \}$$

We observe that temperature sensor s_0 had two seconds of monotonic readings at time points $0s$, $1s$, $2s$, and $5s$. The FROM part of the query specifies that we consider S as input where the window expression $0s < - [\text{NOW} - 2s, \text{NOW}] -> 1s$ states that, at time point t of the evaluation, we are interested in assertions with a time stamp between $t - 2s$ and t . The slide parameter (specified by $-> 1s$) expresses that every second the window moves forward in time. We also assume a pulse of one second, i.e., one query evaluation per second. In STARQL, such a pulse can be defined by a pulse expression declared separately from the SELECT query. The so-called slack parameter for handling out-of-order stream assertions is assumed to be zero ($0s < -$) and not used in the initial version of STARQL.

The input stream data that falls into the window is gathered in a *temporal ABox*, i.e., a set of timestamped ABox assertions:

Time	Temporal ABox
$0s$	$\{rd(s_0, 90)\langle 0s \rangle, rd(s_1, 30)\langle 0s \rangle\}$
$1s$	$\{rd(s_0, 90)\langle 0s \rangle, rd(s_1, 30)\langle 0s \rangle, rd(s_0, 93)\langle 1s \rangle, rd(s_1, 32)\langle 1s \rangle\}$
$2s$	$\{rd(s_0, 90)\langle 0s \rangle, rd(s_1, 30)\langle 0s \rangle, rd(s_0, 93)\langle 1s \rangle, rd(s_1, 32)\langle 1s \rangle, rd(s_0, 94)\langle 2s \rangle\}$
$3s$	$\{rd(s_0, 93)\langle 1s \rangle, rd(s_1, 32)\langle 1s \rangle, rd(s_0, 94)\langle 2s \rangle, rd(s_0, 91)\langle 3s \rangle\}$
$4s$	$\{rd(s_0, 94)\langle 2s \rangle, rd(s_0, 91)\langle 3s \rangle, rd(s_0, 93)\langle 4s \rangle\}$
$5s$	$\{rd(s_0, 91)\langle 3s \rangle, rd(s_0, 93)\langle 4s \rangle, rd(s_0, 95)\langle 5s \rangle\}$

Variable bindings for \vec{x} in the condition $\Psi(\vec{x})$ of the WHERE clause of a query are determined by finding the certain answers for \vec{x} with respect to the ontology formed by the static ABoxes and TBoxes provided by the USING clause. In our example, that means that the variable $?sens$ is bound to the sensor named s_0 , because the TBox at <http://example.org/TBox> together with ABox <http://example.org/Astatic> identify only this sensor to be a temperature sensor.

The SEQUENCE directive of the query defines how to cluster the information in the temporal ABox obtained as above into a sequence of ABoxes. In the basic STARQL version, the only method for doing so is `StdSeq` which separates the assertions based on their timestamp. For example, at time $2s$ we get a sequence of three ABoxes: $\{rd(s_0, 90), rd(s_1, 30)\}\langle 1 \rangle$, $\{rd(s_0, 93), rd(s_1, 32)\}\langle 2 \rangle$, and $\{rd(s_0, 94)\}\langle 3 \rangle$. Note that the absolute timestamps are replaced by natural numbers marking the ordinal position of the ABox in the sequence.

The HAVING clause of the query specifies a condition $\Phi(\vec{x}, \vec{y})$ that has to hold for this sequence of ABoxes, taking also the ontology parts from the USING clause into consideration. The variables in \vec{x} have already been bound by the WHERE clause,

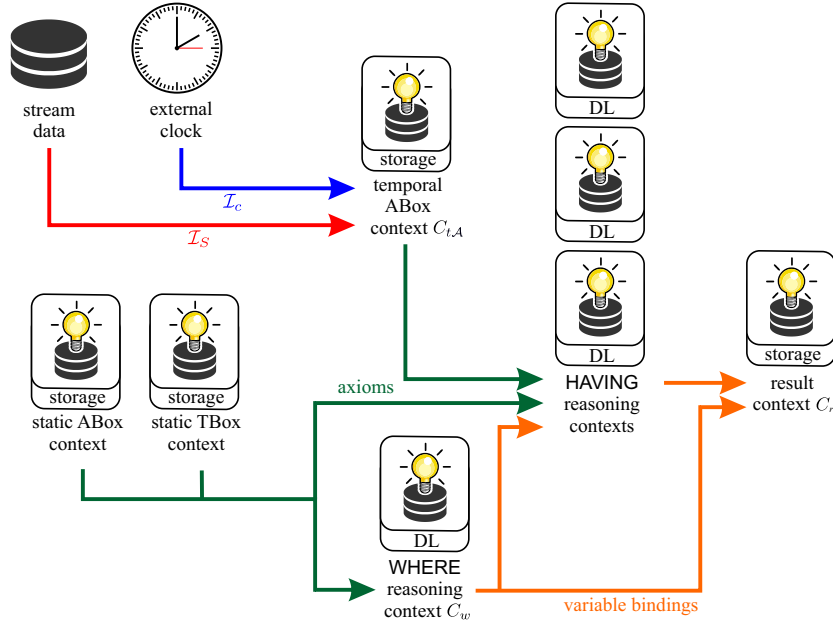


Figure D.4: Structure of an rMCS implementing a STARQL query with an external clock.

so here $?sens$ is being bound to s_0 . The HAVING condition can either bind the remaining open variables in \vec{y} or, as in the example, act as a boolean condition when there are no further open variables. Here, the condition is true if, for all numbers i, j of the sequence (that is from $\{1, 2, 3\}$) with $i \leq j$, the condition holds that if there are sensor readings of s_0 in the ABoxes with sequence positions i and j , then the reading in ABox i is smaller or equal than that of ABox j .

Finally, the expression following the SELECT keyword determines the form of the output stream. For every evaluation time in which the HAVING clause holds, an ABox assertion in the form of an RDF triple $\{?sens \text{ rdf:type MonIncTemp} \langle \text{NOW} \rangle\}$ augmented with a timestamp is returned where $?sens$ is replaced by s_0 and NOW by the evaluation time.

We want to illustrate next how STARQL queries can be realized as rMCSs (Fig. D.4). Certainly, due to the abstract nature of rMCSs, there are many different ways to do so. Here, we aim at a setting where we assume a correspondence of one equilibrium computation per STARQL query evaluation, i.e., in our example there is one equilibrium evaluation per second. We assume one rMCS input stream \mathcal{I}_S that corresponds to the STARQL input stream and one rMCS input stream \mathcal{I}_c that provides the current time as in Section 3.3. The data in \mathcal{I}_S can either be timestamped assertions such as in the original STARQL approach, or unstamped raw data (such as pure sensor values) that will be automatically time-stamped by bridge rules using the current time provided by \mathcal{I}_c . While in the former approach it is possible to have timestamps in the current input that differs from the current time, such a situation is avoided with unstamped data. Both variants have their place and purpose. Here, we assume unstamped data.

We reserve one rMCS context $C_{t,A}$ for the temporal ABox containing the data currently in the window. Its bridge rules add new data from \mathcal{I}_S and timestamp it similar as in Section 3.3. Moreover, other bridge rules ensure that data not “visible” in the window is deleted from $C_{t,A}$.

There is one context C_w that computes the variable bindings of the WHERE clause. Its knowledge base is a description logic ontology and its bridge rules import assertions and axioms according to the USING clause of the STARQL query. Thus, we assume that each static ABox and TBox mentioned in the query is available in a corresponding storage context of the rMCS. These axioms and assertions are also imported to DL contexts for evaluating the HAVING clause: we assume as many such contexts as there can be ABoxes in the sequence of ABoxes created by the SEQUENCE statement. In our example, we have three contexts of this type because there are assertions with at most three different timestamps in the temporal ABox representing the current window. Thus, the sequence operator `StdSeq` generates ABox sequences of at most size three. The reasoning contexts evaluate the different timepoints in the window. For example, the i -th context computes assignments for variable $?x$ of the expression $\{ ?sens \ rd \ ?x \} <i>$. We still need to reason on top of the results for the individual DL contexts for fully evaluating the HAVING clause: In the example, this is the evaluation of the FORALL condition that ranges over different time points. This can be done with another context C_r that also combines the results of the WHERE and HAVING clauses and produces output as specified in the SELECT statement. Thus, in an equilibrium for time t , the belief set for C_r contains timestamped RDF triples of form $\{s_0 \ rd \ ?type \ MonIncTemp\} <t>$ whenever s_0 had monotonically increasing readings during the previous two seconds.

Appendix D.2. EVOLP

We show that rMCSs are expressive enough to capture EVOLP [2] (discussed in Section 7.2) as a particular case. We assume a basic familiarity with logic programs as such, and only point out that a general logic program is a set of logic program rules that allow default negation in the head. *Evolving logic programs* are then defined as general logic programs built over a special language which allows them to express self-evolution. The language includes a reserved unary predicate, `assert`, whose argument may itself be a full-blown rule, thus making arbitrary nesting possible. Formally, given a propositional language \mathcal{L} , the extended language $\mathcal{L}_{\text{assert}}$ over \mathcal{L} is defined inductively as follows:

1. All propositional atoms in \mathcal{L} are atoms in $\mathcal{L}_{\text{assert}}$;
2. if R is a rule over $\mathcal{L}_{\text{assert}}$ then `assert(R)` is an atom in $\mathcal{L}_{\text{assert}}$;
3. nothing else is an atom in $\mathcal{L}_{\text{assert}}$.

An *evolving logic program* over \mathcal{L} is a generalized logic program over $\mathcal{L}_{\text{assert}}$. We denote by $\mathcal{R}_{\mathcal{L}}$ the set of all rules over $\mathcal{L}_{\text{assert}}$.

The idea of EVOLP is that programs can update their own rules thus describing their possible self-evolution. Each self-evolution can be represented by a sequence of programs, each program corresponding to a state, and these sequences of programs can be treated as in Dynamic Logic Programs (DLPs) [1]. Dynamic logic programs are

sequences $P_1 \oplus \dots \oplus P_n$ of generalized logic programs, whose semantics is based on the causal rejection principle. The idea is that the most recent rules are put in force, (where P_n is to be seen as the most recent set of rules), and the previous rules are valid as far as possible, i.e., they are only kept if they do not conflict with more recent rules. Here, these intuitions about DLPs are sufficient, and we point to [1] for more details.

The semantics of evolving logic programs is based on sequences of interpretations. More precisely, an *evolution interpretation* of length m of an evolving logic program P over a propositional language \mathcal{L} is a finite sequence $I = \langle I^1, \dots, I^m \rangle$ of sets of propositional atoms of $\mathcal{L}_{\text{assert}}$. The evolution trace associated with an evolution interpretation I is the sequence of programs $P_I = \langle P^1, \dots, P^m \rangle$ such that $P^1 = P$ and, for each $2 \leq j \leq m$, $P^j = \{r \mid \text{assert}(r) \in I^{j-1}\}$.

Given an evolving logic program, the main intuition for the construction of a sequence of programs that corresponds to a possible evolution of P is that whenever an atom $\text{assert}(r)$ belongs to an interpretation in a sequence, then the rule r must belong to the program in the next state.

Besides self-evolution, evolving logic programs also consider evolution caused by the addition of external rules. These rules, called events, are represented as a sequence of evolving logic programs. Given an evolving logic program P over \mathcal{L} , such an *event sequence over P* is a sequence of evolving logic programs over \mathcal{L} .

This leads to the central model notion of evolving logic programs that also takes into account an incoming event sequence. Let P be an evolving logic program over \mathcal{L} , $I = \langle I^1, \dots, I^m \rangle$ an evolution interpretation of length m of P with evolution trace $P_I = \langle P^1, \dots, P^m \rangle$, and $E = \langle E^1, \dots, E^\ell \rangle$ an event sequence over P such that $\ell \geq m$. Then, I is an *evolution stable model of P* given E iff for every $1 \leq j \leq m$, we have that I^j is a stable model of $P^1 \oplus P^2 \oplus \dots \oplus (P^j \cup E^j)$.

We now show how EVOLP can be captured in the framework of rMCSs. For that, given an evolving logic program P we aim to construct an rMCS M_P whose equilibria streams corresponds to the evolution stable models of P . First of all, note that the events in EVOLP can be modeled by the input streams of rMCSs.

The intuitive idea for the construction of M_P is that, at each instant, the incoming events are temporarily added to a knowledge base. To keep track of these, we consider a context C , whose possible knowledge bases are pairs. The first component of such is to keep track of the sequence of programs $P^1 \oplus \dots \oplus P^j$, which corresponds to the trace of an interpretation until the current instant, and the second component is reserved for the current event E^j . Given this intuition, it is clear that $\text{acc}(\langle P^1 \oplus \dots \oplus P^j, E^j \rangle)$ should be defined as the set of stable models of $P^1 \oplus \dots \oplus P^{j-1} \oplus (P^j \cup E^j)$. To incorporate events, we consider an input language IL of M_P defined precisely as the language of events, i.e., the set of evolving logic programs. Moreover, context C contains bridge rules to access the current events, which are then used to update the event's component of the knowledge base. Also, context C has bridge rules designated to guarantee that the formulas that occur under the predicate assert at the current instant are used to update the knowledge base in the next instant. Finally, the management function mng of C is such that $\text{mng}(op, \langle P, E \rangle)$ updates the program sequence component or the event component depending on whether op is a set of asserts, or a set of events, respectively.

Formally, let P be an evolving logic program, and consider the following rMCS $M_P = \langle \langle C \rangle, \langle IL \rangle, \langle BR_C \rangle \rangle$ where

- $C = \langle L, OP, \mathbf{mng} \rangle$ such that
- $L = \langle KB, BS, \mathbf{acc} \rangle$ is a logic such that
- KB is the set of pairs $\langle D, E \rangle$ where D is a dynamic logic program over \mathcal{L} , and E is an evolving logic program over \mathcal{L}
- BS is the set of all subsets of $\mathcal{L}_{\text{assert}}$
- $\mathbf{acc}(\langle P^1 \oplus \dots \oplus P^j, E \rangle)$ is the set of stable models of $P^1 \oplus \dots \oplus P^{j-1} \oplus (P^j \cup E)$
- $OP = \{\text{as}(r) \mid r \in \mathcal{R}_{\mathcal{L}}\} \cup \{\text{ob}(r) \mid r \in \mathcal{R}_{\mathcal{L}}\}$
- $\mathbf{mng}(op, \langle D, E \rangle) = \{\langle D \oplus U, E' \rangle\}$ where $U = \{r \in \mathcal{R}_{\mathcal{L}} \mid \text{as}(r) \in op\}$ and $E' = \{r \in \mathcal{R}_{\mathcal{L}} \mid \text{ob}(r) \in op\}$
- $IL = \mathcal{R}_{\mathcal{L}}$
- $BR_C = \{\mathbf{next}(\text{as}(s)) \leftarrow 1:\text{assert}(s) \mid \text{assert}(s) \in \mathcal{L}_{\text{assert}}\} \cup \{\text{ob}(s) \leftarrow 1::s \mid s \in IL\}$

Evolution stable models of P are strongly related to the equilibria streams of M_P . Namely, given an event sequence $E = \langle E_1, \dots, E_\ell \rangle$ over P consider its associated input stream $\mathcal{I}_E = \langle \langle E_1 \rangle, \dots, \langle E_\ell \rangle \rangle$. Then, $I = \langle I^1, \dots, I^\ell \rangle$ is an evolution stable model of P given E iff I is an equilibria stream for M_P given $\text{KB} = \langle \langle P, \emptyset \rangle \rangle$ and \mathcal{I}_E .

This coincidence between evolution stable models in EVOLP and equilibria streams for rMCSs as defined above can be explained by the fact that, conceptually, the operators **next** and **assert** are rather similar.