TPLP **20** (1): 1–43, 2020. © Cambridge University Press 2019
doi:10.1017/S1471068419000073 First published online 19 June 2019

A Logic Framework for P2P Deductive Databases

LUCIANO CAROPRESE and ESTER ZUMPANO
Department of Computer, Modelling, Electronics and Systems Engineering, University of Calabria, Rende (CS), Italy
(*e-mails:* l.caroprese@dimes.unical.it, e.zumpano@dimes.unical.it)

submitted 10 May 2017; revised 8 May 2019; accepted 9 May 2019

Abstract

This paper presents a logic framework for modeling the interaction among deductive databases in a peer-to-peer (P2P) environment. Each peer joining a P2P system *provides or imports data* from its neighbors by using a set of *mapping rules*, that is, a set of semantic correspondences to a set of peers belonging to the same environment. By using mapping rules, as soon as it enters the system, a peer can participate and access all data available in its neighborhood, and through its neighborhood it becomes accessible to all the other peers in the system. A query can be posed to any peer in the system and the answer is computed by using locally stored data and all the information that can be consistently imported from the neighborhood. Two different types of mapping rules are defined: mapping rules allowing to import a maximal set of atoms not leading to inconsistency (called *maximal mapping rules*) and mapping rules allowing to import a minimal set of atoms needed to restore consistency (called *minimal mapping rules*). Implicitly, the use of maximal mapping rules states *it is preferable to import as long as no inconsistencies arise*; whereas the use of minimal mapping rules states that *it is preferable not to import unless a inconsistency exists*. The paper presents three different declarative semantics of a P2P system: (i) the *Max Weak Model Semantics*, in which mapping rules are used to import *as much knowledge as possible* from a peer's neighborhood without violating local integrity constraints; (ii) the *Min Weak Model Semantics*, in which the P2P system can be locally inconsistent and the information provided by the neighbors is used to restore consistency, that is, to only integrate the missing portion of a correct, but incomplete database; (iii) the *Max-Min Weak Model Semantics* that unifies the previous two different perspectives captured by the Max Weak Model Semantics and Min Weak Model Semantics. This last semantics allows to characterize each peer in the neighborhood as a resource used either to enrich (integrate) or to fix (repair) the knowledge, so as to define a kind of *integrate–repair* strategy for each peer. For each semantics, the paper also introduces an equivalent and alternative characterization, obtained by rewriting each mapping rule into prioritized rules so as to model a P2P system as a prioritized logic program. Finally, results about the computational complexity of P2P logic queries are investigated by considering *brave* and *cautious* reasoning.

KEYWORDS: peer data exchange, incompleteness, inconsistency, integrity constraints, relational databases, prioritized logic program

1 Introduction

Data exchange consists in sharing data from a source schema to a target schema according to specifications fixed by source-to-target constraints (Fagin *et al.* 2005; Fagin *et al.* 2005; Fuxman *et al.* 2006). This challenging topic is closely related to data integration and consistent query answering (Lenzerini 2002; Greco *et al.* 2003; Leone *et al.* 2005; Arenas *et al.* 1999; Calì *et al.* 2003; Calì *et al.* 2004). Data integration is one of the most fundamental processes in intelligent systems, from individuals to societies. At the present, the most

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

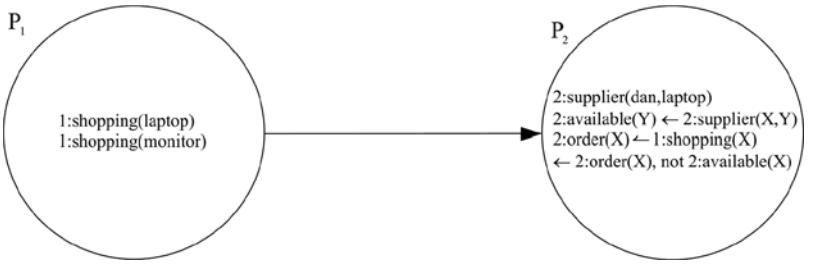


Fig. 1. A P2P system with maximal mapping rules.

important application of data integration is any form of peer-to-peer (P2P) interaction and cooperation. Ideally, in P2P systems there is no selection, but integration of the valuable contributions of every participant.

In a *Peer Data Managment System*, PDMS, a number of peers interact and exchange data. More specifically, each peer joining a P2P system uses a set of *mapping rules*; that is, a set of semantic correspondences to a set of peers belonging to the same environment, to both *provide or import data* from its neighbors. Therefore, in a P2P system the entry of a new source, *peer*, is extremely simple as it just requires the definition of the mapping rules. By using mapping rules, as soon as it enters the system, a peer can participate and access all data available in its neighborhood, and through its neighborhood it becomes accessible to all other peers in the system.

The possibility for the users of sharing knowledge from a large number of informative sources has enabled the development of new methods for data integration easily usable for processing distributed and autonomous data.

Due to this, there have been several proposals which consider the integration of information and the computation of queries in an open-ended network of distributed peers (Bernstein *et al.* 2002; Bertossi and Bravo 2004; Calvanese *et al.* 2004; Calvanese *et al.* 2003; Franconi *et al.* 2003) as well as the problem of schema mediation (Halevy *et al.* 2003; Madhavan and Halevy 2003; Halevy *et al.* 2005), query answering, and query optimization in P2P environments (Abiteboul and Duschka 1998; Tatarinov and Halevy 2004; Gribble *et al.* 2001; Fagin *et al.* 2005).

Previously proposed approaches investigate the data integration problem in a P2P system by considering each peer as locally consistent. Therefore, the introduction of inconsistency is only caused by the operation of importing data from other peers. These approaches assume that for each peer it is preferable to import as much knowledge as possible.

Our previous works, in the context of P2P data integration, follow this direction.

In Caroprese *et al.* (2006a); Caroprese *et al.* (2006b); Caroprese and Zumpano (2007); Caroprese and Zumpano (2008); Caroprese and Zumpano (2012a), it is adopted the classical idea that a peer imports maximal sets of atoms. More specifically, the interaction among deductive databases in a P2P system has been modeled by importing maximal sets of atoms not violating integrity constraints, that is, *maximal sets of atoms that allow the peer to enrich its knowledge while preventing inconsistency anomalies.* The following examples will clarify the perspective used by *maximal mapping rules* to import in each peer maximal sets of atoms not violating integrity constraints.

Example 1.
Consider the P2P system depicted in Figure 1.

- Peer P_1 stores information about products that should be ordered. It contains the facts: 1:*shopping*(*laptop*) and 1:*shopping*(*monitor*). The special syntax used for a fact – its first part is the *peer identifier* – will be formally presented in Section 3.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

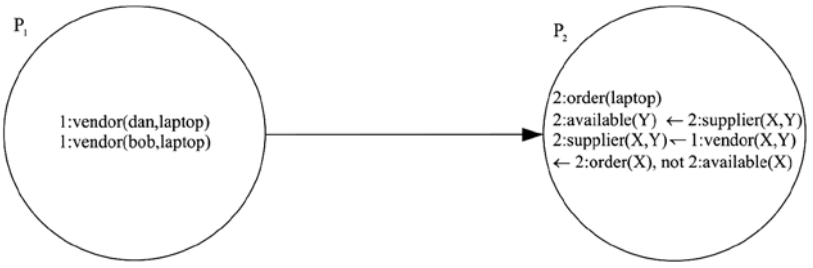


Fig. 2. A P2P system with minimal mapping rules.

- Peer P_2 contains
  - the fact 2:$supplier(dan, laptop)$, whose meaning is "*Dan is a supplier of laptops*";
  - the maximal mapping rule, 2:$order(X) \leftarrow$ 1:$shopping(X)$, whose precise syntax and semantics will be formally defined in Section 3. Intuitively, this rule allows to import as many orders as possible from the relation shopping of P_1 into the relation order of P_2. In fact, it states that if 1:$shopping(X)$ is *true* in the source peer P_1, the atom 2:$order(X)$ can be imported in the target peer P_2 (that is 2:$order(X)$ is *true* in the target peer) only if it does not imply the violation of some integrity constraints;
  - the rule 2:$available(Y) \leftarrow$ 2:$supplier(X, Y)$ stating that a product Y is available if there is a supplier X of Y;
  - the integrity constraint $\leftarrow$ 2:$order(X)$, *not* 2:$available(X)$, stating that the order of a device cannot exist if it is not available.

Intuitively, peer P_1 provides two facts, but the maximal set of them that P_2 can import, using the mapping rule, is {2:$order(laptop)$}. The fact 2:$order(monitor)$ cannot be imported as it would violate the integrity constraint; in fact, no supplier of the device *monitor* exists.

Besides the basic classical idea followed in the previous example, a different perspective could be argued. Often, in real-world P2P systems, peers use the available import mechanisms to extract knowledge from the rest of the system only if this knowledge is strictly needed to repair an inconsistent local database. The work in Caroprese and Zumpano (2012b) stems from this different perspective. A peer can be locally inconsistent and it can use the information provided by its neighbors in order to restore consistency, that is, to only integrate the missing portion of a correct, but incomplete database. Then, an inconsistent peer, in the interaction with different peers, just imports the information allowing to restore consistency, that is, *minimal sets of atoms allowing the peer to enrich its knowledge so as to restore inconsistency anomalies.*

The following example will intuitively clarify this perspective.

Example 2.
Consider the P2P system depicted in Figure 2. It consists of the following two peers:

- Peer P_1 stores information about vendors of devices and contains the following facts: 1:$vendor$ $(dan, laptop)$, whose meaning is *"Dan is a vendor of laptops,"* and 1:$vendor(bob, laptop)$, whose meaning is *"Bob is a vendor of laptops."*
- Peer P_2 contains:
  - the fact 2:$order(laptop)$, stating that the order of a laptop exists;
  - the minimal mapping rule 2:$supplier(X, Y) \leftarrow$ 1:$vendor(X, Y)$, whose precise syntax and semantics will be formally defined in Section 3. The rule is used to

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

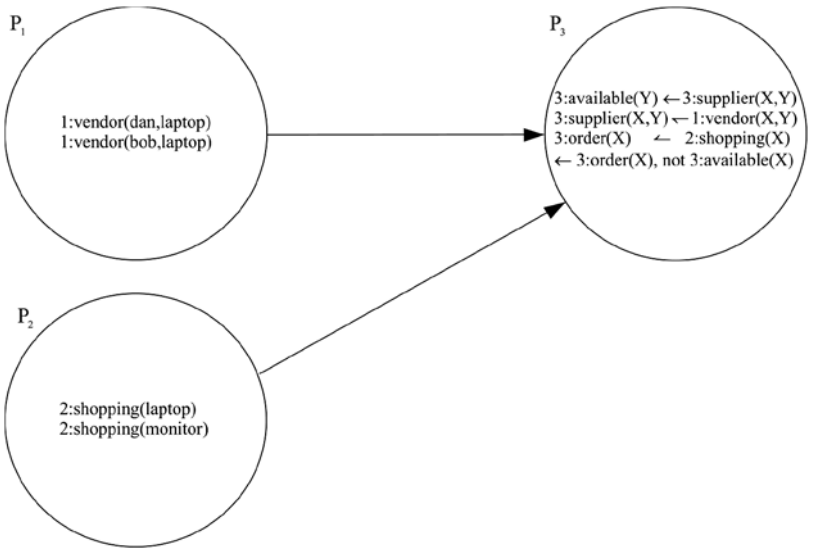


Fig. 3. A P2P system with maximal and minimal mapping rules.

import tuples from the relation *vendor* of P_1 into the relation *supplier* of P_2. Intuitively, the rule states that if 1:$vendor(X, Y)$ is *true* in the source peer the atom 2:$supplier(X, Y)$ can be imported in the target peer (that is 2:$supplier(X, Y)$ is *true* in the target peer) only if it implies the satisfaction of some constraints that otherwise would be violated;

— the standard rule 2:$available(Y) \leftarrow$ 2:$supplier(X, Y)$, stating that a device Y is available if there is a supplier X of Y;

— the integrity constraint $\leftarrow$ 2:$order(X)$, *not* 2:$available(X)$, stating that the order of a device cannot exist if it is not available.

Peer P_2 is inconsistent. The integrity constraint is violated as the ordered device *laptop* is not available (there is no supplier of laptops). The device *laptop* needs to be provided by a supplier. Therefore, P_2 "needs" to import from its neighbors' minimal sets of atoms in order to restore consistency. The intuition is that either 1:$vendor(dan, laptop)$ or 1:$vendor(bob, laptop)$ can be imported into P_2 to satisfy the constraint (but not both).

The two concepts proposed in Caroprese *et al.* (2006a); Caroprese *et al.* (2006b); Caroprese and Zumpano (2007); Caroprese and Zumpano (2008); Caroprese and Zumpano (2012a) and in Caroprese and Zumpano (2012b) can be merged. The basic idea is that a peer of a P2P system can use each neighbor to extract either as much knowledge as possible (i.e. to integrate its knowledge) or just the portion that is strictly needed (i.e. to repair the knowledge of the system). This unified framework defines a sort of *integrate–repair* strategy.

The following example will intuitively clarify our perspective and will be used as a running example in the rest of the paper.

Example 3.
Consider the P2P system depicted in Figure 3.

- Peer P_1 stores information about vendors of devices and contains the facts: 1:$vendor\ (dan, laptop)$, whose meaning is "*Dan is a vendor of laptops*" and 1:$vendor(bob, laptop)$, whose meaning is "*Bob is a vendor of laptops*";

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

- Peer P_2 stores information about devices that should be ordered:
  2:$shopping(laptop)$ and
  2:$shopping(monitor)$;
- Peer P_3 contains
  - — the integrity constraint
    $\leftarrow$ 3:$order(X)$, *not* 3:$available(X)$, disallowing to import the order of a device that cannot be provided by any supplier.
  - — the standard rule
    3:$available(Y) \leftarrow$ 3:$supplier(X, Y)$, stating that a device Y is available if there is a supplier X of Y;
  - — two mapping rules
    - 3:$order(X) \leftarrowtail$ 2:$shopping(X)$ that, intuitively, allows to import as many orders as possible from P_2 and
    - 3:$supplier(X, Y) \rightarrowtail$ 1:$vendor(X, Y)$ that allows to import minimal sets of supplier from P_1 able to provide the ordered devices.

The intuitive meaning of the P2P system is the following: the fact 2:$shopping(laptop)$, belonging to the adding resource P_2, can be used to derive 3:$order(laptop)$. This fact does not violate the integrity constraint in P_3 thanks to the repair resource P_1 whose role is to try to guarantee the consistency of P_3. In more detail, either 1:$vendor(dan, laptop)$ or 1:$vendor(bob, laptop)$ can be used to derive the fact 3:$available(laptop)$ and thus to satisfy the constraint. Therefore, the preferred scenarios of the system, called max-min weak models, contain besides the base predicates, either the facts {3:$supplier(dan$, $laptop)$, 3:$order(laptop)$, 3:$available(laptop)$} or the facts {3:$supplier$ $(bob, laptop)$, 3:$order$ $(laptop)$, 3:$available(laptop)$}.

Observe that, we cannot act in a similar way with respect to the fact 2:$shopping(monitor)$ belonging to P_2: no repair mechanism can be activated in order to support the derived predicate 3:$order(monitor)$.

Summarizing, the presence of the *repair resource* P_1 allows the system to fix the knowledge imported from P_2.

In the previous example, peer P_3 aims at enriching its knowledge by importing from P_2 as much knowledge as possible, and uses P_1 to (eventually) restore inconsistencies. Therefore, with respect to P_3, peer P_2 acts as an *adding resource*, whereas peer P_1 acts as a *repair resource.*

Different alternative semantics for P2P systems, that will be extensively discussed in Section 6, have been proposed in the literature. In any case, in each of them mapping rules are used as a vehicle to import data.

Our approach, as well as in general P2P data management systems, can be viewed as a special case of multi-context systems (MCS) as it models autonomous logic-based entities (peers) that interchange pieces of information using mapping rules. In any case, the essential feature of P2P system is that each peer may leave and join the system arbitrarily. Due to this specific dynamic nature, the focus in the P2P context is not that of finding the explanations of inconsistencies, but just cope with them. In our work, due to the two different forms of mapping rules, each peer is given the possibility to decide how to interact with a neighbor peer: as a source used to maximize its own knowledge or as a source used to fix its own knowledge. This specific notion has not a counterpart in any of the previous works in the literature, neither in the field of MCS, nor in the field of P2P systems.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Contributions. The paper presents a logic-based framework for modeling the interaction among peers. It is assumed that each peer consists of a database, a set of standard logic rules, a set of mapping rules of two possible types, and a set of integrity constraints. In such a context, a query can be posed to any peer in the system and the answer is provided by using locally stored data and all the information that can be consistently imported from the neighborhood.

In synthesis, the main contributions of the paper are the following:

- The introduction of two different forms of mapping rules: *maximal mapping rules* used to import maximal sets of atoms while preventing inconsistency anomalies and *minimal mapping rules* used to fix the knowledge by importing minimal sets of atoms allowing to restore consistency. In other words, maximal mapping rules state that it is preferable to import as long as no local inconsistencies arise, whereas minimal mapping rules state that it is preferable not to import unless a local inconsistency exists. By using these two forms of mapping rules, a generic peer is able to decide how to interact with a neighbor peer: as a source used to maximize its own knowledge or as a source used to fix its own knowledge.
- The proposal of the *Max Weak Model Semantics*, in which mapping rules are used to import *as much knowledge as possible* from its neighborhood without violating local integrity constraints.
- The proposal of the *Min Weak Model Semantics*, in which the P2P system can be locally inconsistent and the information provided by the neighbors is used to restore consistency, that is, to only integrate the missing portion of a correct, but incomplete database.
- The *Max-Min Weak Model Semantics* that unifies the previous two different perspectives captured by the Max and Min Weak Model Semantics. This more general declarative semantics allows to characterize each peer in the neighborhood as a resource used either to enrich (integrate) or to fix (repair) the knowledge, so as to define a kind of *integrate–repair* strategy for each peer in the P2P setting.
- The definition of an alternative characterization of the Max-Min Weak Model Semantics (resp. Max Weak Model Semantics and Min Weak Model Semantics) obtained by rewriting mapping rules into prioritized rules. Therefore, a P2P system PS is rewritten into an equivalent prioritized logic program, $Rew(PS)$, such that the max-min weak models of PS (resp. maximal weak models and minimal weak models) are the preferred stable models of $Rew(PS)$.
- Results on the complexity of answering queries. The paper, by considering analogous results on stable model semantics for *prioritized logic programs* (PLP), proves that for disjunction-free ($\vee - free$) prioritized programs deciding whether an interpretation M is a max-min weak model (resp. maximal weak models and minimal weak model) of PS is $coNP$ complete, deciding whether an atom is *true* in some max-min weak model (resp. maximal weak models and minimal weak model) is Σ_2^p-complete, whereas deciding whether an atom is *true* in every preferred model is Π_2^p-complete (Sakama and Inoue 2000). Moreover, the paper also provides results on the existence of a max-min weak model (resp. maximal weak model and minimal weak model) showing that the problem is in Σ_2^p.
- An extensive section, Discussion, reporting different features of the proposal. In more detail, the practical aspects of the proposal are highlighted and several additional and alternative issues, arising from the basic framework, are presented: a technique allowing to deal with P2P systems locally inconsistent; a deterministic

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

semantics, derived from the Max Weak Model Semantics, allowing to assign a unique three value model to particular types of P2P systems; a polynomial distributed algorithm for its computation; a system prototype.

Structure of the paper. The remainder of the paper is organized as follows. Section 2 introduces relevant background information. Section 3 describes the syntax of P2P systems. Section 4 describes alternative semantics, namely the *Max Weak Model Semantics* in Caroprese *et al.* (2006a), the *Min Weak Model Semantics* in Caroprese and Zumpano (2012b) and introduces a new formal declarative semantics, called *Max-Min Weak Model Semantics*, that unifies the previous two into a more general perspective. Moreover, it introduces, for each of the proposed semantics, an alternative characterization, modeled in terms of logic programs with priorities. Section 5 presents results on computational complexity, Section 6 focuses on some relevant discussions related to practical aspects of the proposed framework, and Section 7 introduces a comprehensive discussion of related works. Finally, Section 8 reports concluding remarks and directions for further research.

2 Background

We assume that there are finite sets of *predicates*, constants and variables (Abiteboul *et al.* 1995). A *term* is either a constant or a variable. An *atom* is of the form $p(t_1, \ldots, t_n)$, where p is a predicate and $t_1, \ldots, t_n$ are terms. A *literal* is either an atom A or its negation *not* A. As in this work we use the *Closed World Assumption*, we adopt *negation as failure*. A *rule* is of the form:

- $H \leftarrow \mathcal{B}$, where H is an atom and $\mathcal{B}$ is a conjunction of literals or
- $\leftarrow \mathcal{B}$, where $\mathcal{B}$ is a conjunction of literals.

H is called *head* of the rule and $\mathcal{B}$ is called *body* of the rule. A rule of the form $\leftarrow \mathcal{B}$ is also called *constraint*. A program P is a finite set of rules. P is said to be positive if it is negation free.

An *exclusive disjunctive rule* is the form $A \oplus A' \leftarrow \mathcal{B}$ and it is a notational shorthand for $A \leftarrow \mathcal{B} \wedge not\ A'$, $A' \leftarrow \mathcal{B} \wedge not\ A$ and $\leftarrow A \wedge A'$.[1] Its intuitive meaning is that if $\mathcal{B}$ is *true* then exactly one of A and A' must be *true*.

It is assumed that programs are *safe* – that is, variables occurring in the head or in negated body literals are range restricted as they occur in some positive body literal.

An atom (resp. literal, rule, program) is *ground* if no variable occurs in it. A ground atom is also called *fact*. The set of ground instances of an atom a (resp. literal l, rule r, program P), denoted by $ground(a)$ (resp. $ground(l)$, $ground(r)$, $ground(P)$) is built by replacing variables with constants in all possible ways. An interpretation is a set of facts. The truth value of ground atoms, literals, and rules with respect to an interpretation M is as follows: $val_M(A) = (A \in M)$, $val_M(not\ A) = not\ val_M(A)$, $val_M(L_1, \ldots, L_n) = min\{val_M(L_1), \ldots, val_M(L_n)\}$, and $val_M(A \leftarrow L_1, \ldots, L_n) = val_M(A) \geq val_M(L_1, \ldots, L_n)$, where A is an atom, $L_1, \ldots, L_n$ are literals, and $true > false$. An interpretation M is a model for a program P, if all rules in $ground(P)$ are *true* w.r.t. M. A model M of a program P is said to be minimal if there is no model N of P such that $N \subset M$. We denote the set of minimal models of a program P with $MM(P)$. Given an interpretation M and a predicate g, $M[g]$ denotes the set of g-tuples in M. The semantics of a pos-

[1] We use for the operator *and* both "," and "∧."

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

itive program P is given by its unique minimal model which can be computed by applying the *immediate consequence operator* $\mathbf{T}_P$ until the fixpoint is reached ($\mathbf{T}_P^{\infty}(\emptyset)$). The semantics of a program with negation P is given by the set of its stable models, denoted as $SM(P)$. An interpretation M is a *stable model* of P if M is the unique minimal model of the positive program P^M, where P^M is obtained from $ground(P)$ by (i) removing all rules r such that there exists a negative literal *not* A in the body of r and A is in M and (ii) removing all negative literals from the remaining rules (Gelfond and Lifschitz 1988). It is well known that stable models are minimal models (i.e. $SM(P) \subseteq MM(P)$).

2.1 Prioritized logic programs

Several works have investigated various forms of priorities into logic languages (Brewka and Eiter 1999; Brewka *et al.* 2003; Delgrande *et al.* 2003; Sakama and Inoue 2000). In this paper, we refer to the extension proposed in Sakama and Inoue (2000).

A *preference relation* $\succeq$ among ground atoms is defined as follows. For any ground atoms e_1 and e_2, if $e_1 \succeq e_2$ then we say that e_1 *has a higher priority than* e_2. $e_1 \succ e_2$ stands for $e_1 \succeq e_2$ and $e_2 \not\succeq e_1$. The statement $e_1 \succeq e_2$ is called a *priority*. The statement $p_1(x) \succeq p_2(y)$, where x and y are tuples containing variables, stands for every priority $p_1(s) \succeq p_2(t)$, where s and t are instances of x and y, respectively.

If $p(x) \succ p(y)$, $p(x)$, and $p(y)$ do not have common ground instances. Indeed, assuming that there is a ground atom $p(s)$ which is an instance of $p(x)$ and $p(y)$, the statements $p(s) \succeq p(s)$ and $p(s) \not\succeq p(s)$ would hold at the same time (a contradiction).

Given a set Φ of priorities, we define the closure Φ^* as the set of priorities which are reflexively or transitively derived using priorities in Φ.

Let $\mathcal{M}$ be a class of sets of ground atoms and Φ a set of priorities. The relation $\sqsupseteq$ is defined over the sets of $\mathcal{M}$ as follows. For any sets M_1, M_2, and M_3 of $\mathcal{M}$,

- $M_1 \sqsupseteq M_1$;
- $M_1 \sqsupseteq M_2$ if $\exists\, e_1 \in M_1 \setminus M_2$, $\exists\, e_2 \in M_2 \setminus M_1$ such that $(e_1 \succeq e_2) \in \Phi^*$ and $\nexists\, e_3 \in M_2 \setminus M_1$ such that $(e_3 \succ e_1) \in \Phi^*$;
- if $M_1 \sqsupseteq M_2$ and $M_2 \sqsupseteq M_3$, then $M_1 \sqsupseteq M_3$.

If $M_1 \sqsupseteq M_2$ holds, then we say that M_1 is *preferable* to M_2 w.r.t. Φ. Moreover, we write $M_1 \sqsupset M_2$ if $M_1 \sqsupseteq M_2$ and $M_2 \not\sqsupseteq M_1$. A set M is a *preferred* set of $(\mathcal{M}, \Phi)$ if M is in $\mathcal{M}$ and there is no set N in $\mathcal{M}$ such that $N \sqsupset M$. The class of preferred sets of $(\mathcal{M}, \Phi)$ will be denoted by $PS(\mathcal{M}, \Phi)$. A PLP is of the form $(P, \Phi_1, \ldots, \Phi_n)$, where P is a logic program and $\Phi_1, \ldots, \Phi_n$, with $n \geq 1$, are sets of priorities. The *preferred stable models* of $(P, \Phi_1, \ldots, \Phi_n)$ denoted as $PSM(P, \Phi_1, \ldots, \Phi_n)$ are the stable models of P selected by applying consecutively the sets of priorities $\Phi_1, \ldots, \Phi_n$. More formally,

- $PSM(P, \Phi_1) = PS(SM(P), \Phi_1)$;
- $PSM(P, \Phi_1, \ldots, \Phi_n) = PS(PSM(P, \Phi_1, \ldots, \Phi_{n-1}), \Phi_n)$.

3 P2P systems

A peer identifier is a number $i \in \mathbb{N}^+$. A *(peer) predicate* is a pair $i{:}p$, where i is a peer identifier and p is a predicate.[2] A *(peer) atom* A is of the form $i{:}p(X)$, where

[2] Whenever the reference to a *peer predicate* (resp. *peer atom, peer literal, peer fact, peer rule, peer standard rule, peer integrity constraint, peer maximal mapping rule, and peer minimal mapping rule*) is clear from the context, the term *peer* can be omitted.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

i is a peer identifier, $p(X)$ is an atom, and X is a list of terms. A *(peer) literal* is a peer atom A or its negation *not* A. A conjunction $\mathcal{B} = i{:}p_1(X_1), \ldots, i{:}p_m(X_m)$, *not* $i{:}p_{m+1}(X_{m+1}), \ldots,$ *not* $i{:}p_n(X_n), \varphi$, where φ is a conjunction of built-in atoms,[3] will be also denoted as $i{:}(p_1(X_1), \ldots, p_m(X_m),$ *not* $p_{m+1}(X_{m+1}), \ldots,$ *not* $p_n(X_n), \varphi)$.

Definition 1 (*Peer Rule*). A *(Peer) rule* can be of one of the following types:

1. *(Peer) standard rule.*
   It is of the form $H \leftarrow \mathcal{B}$, where $H = i{:}h(X)$ and $\mathcal{B} = j{:}(p_1(X_1), \ldots, p_m(X_m),$ *not* $p_{m+1}(X_{m+1}), \ldots,$ *not* $p_n(X_n), \varphi)$.
2. *(Peer) integrity constraint.*
   It is of the form $\leftarrow \mathcal{B}$, where $\mathcal{B} = i{:}(p_1(X_1), \ldots, p_m(X_m),$ *not* $p_{m+1}(X_{m+1}), \ldots,$ *not* $p_n(X_n), \varphi)$,
3. *(Peer) maximal mapping rule.*
   It is of the form $H \leftharpoonup \mathcal{B}$, where $H = i{:}h(X)$, $\mathcal{B} = j{:}(p_1(X_1), \ldots, p_m(X_m), \varphi)$ and $i \neq j$.
4. *(Peer) minimal mapping rule.*
   It is of the form $H \rightharpoondown \mathcal{B}$, where $H = i{:}h(X)$, $\mathcal{B} = j{:}(p_1(X_1), \ldots, p_m(X_m), \varphi)$ and $i \neq j$.

In the previous definition, i (resp. j) is the *peer identifier* (resp. *source peer identifier*) of the rule, H is the *head* of the rule, and $\mathcal{B}$ is the *body* of rule. With the term *mapping rule*, we refer to a maximal mapping rule or to a minimal mapping rule. The concepts of *ground rule*, *fact*, and *interpretation* are similar to those reported in Section 2. Given a fact $i{:}p(x)$, i is its peer identifier. In our setting, a predicate is of exactly one of the following three types: *base predicate*, *derived predicate*, and *mapping predicate*. A derived predicate is a predicate occurring in the head of a standard rule, and a mapping predicate is a predicate occurring in the head of a mapping rule. If a predicate is neither a derived predicate nor a mapping predicate, then it is a base predicate.

An atom $i{:}p(X)$ is a *base atom* (resp. *derived atom*, *mapping atom*) if $i{:}p$ is a base predicate (resp. derived predicate, mapping predicate).

The intuitive meaning of a standard rule is that *whenever its body is true, its head has to be true*. This means that an interpretation M satisfies a standard rule r if for each ground instance r' of it, M does not satisfy the body of r' or M satisfies the head of r'.

The intuitive meaning of an integrity constraint is that *its body has to be false*. Therefore, an interpretation M satisfies an integrity constraint c if for each ground instance c' of it, M does not satisfy the body of c'.

The intuitive meaning of a maximal rule is that *whenever its body is true, its head has to be true if it does not violate (directly or indirectly) any integrity constraint.*

Finally, the intuitive meaning of a minimal rule is that *whenever its body is true, its head has to be true if it prevents the violation (directly or indirectly) of some integrity constraint.*

In the following sections, we will see how the semantics of a maximal mapping rule and a minimal mapping rule can be captured by an *exclusive disjunctive rule* and a *priority*.

Given an interpretation M, $M[D]$ (resp. $M[LP]$, $M[MP]$) denotes the subset of base atoms (resp. derived atoms, mapping atoms) in M.

[3] A *built-in atom* is of the form $\theta(X, Y)$, where X and Y are terms and $\theta \in \{<, >, \leq, \geq, =, \neq\}$. It is also denoted as $X\ \theta\ Y$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Definition 2 (*P2P system*). A *peer* P_i, with a peer identifier i, is a tuple $\langle D_i, LP_i, MP_i, IC_i\rangle$, where

- D_i is a set of facts whose peer identifier is equal to i (*local database*);
- LP_i is a set of standard rules whose peer identifier and source peer identifier are equal to i;
- MP_i is a set of mapping rules whose peer identifier is equal to i;
- IC_i is a set of constraints over predicates defined by D_i, LP_i, and MP_i whose peer identifier is equal to i.

A *P2P system* PS is a set of peers $\{P_1, \ldots, P_n\}$ s.t. for each source peer identifier j occurring in its mapping rules, $j \in [1..n]$.

Given a peer $P_i = \langle D_i, LP_i, MP_i, IC_i\rangle$, we denote as

- $\overline{MP_i}$ the subset of maximal mapping rules in MP_i;
- $\underline{MP_i}$ the subset of minimal mapping rules in MP_i.

Clearly, $MP_i = \overline{MP_i} \cup \underline{MP_i}$. Without loss of generality, we assume that every mapping predicate is defined by only one mapping rule of the form $i{:}p(X) \leftharpoondown j{:}q(X)$ (resp. $i{:}p(X) \leftharpoonup j{:}q(X)$). Indeed, a mapping predicate $i{:}p$ consisting of n rules of the form $i{:}p(X) \Leftarrow_k j_k{:}\mathcal{B}_k$, with $\Leftarrow_k \in \{\leftharpoonup, \leftharpoondown\}$ and $k \in [1..n]$, can be rewritten into $2 \cdot n$ rules of the form $i{:}p_k(X) \Leftarrow_k j_k{:}\mathcal{B}_k$ and $i{:}p(X) \leftarrow i{:}p_k(X)$ with $k \in [1..n]$. Observe that, $i{:}p$ becomes a derived predicates and $i{:}p_k(X)$, with $k \in [1..n]$, are new mapping predicates. Moreover, there is no loss of generality in considering mapping rules having a positive body. Indeed, allowing negation in the body of mapping rules, a mapping rule $H \Leftarrow \mathcal{B}(X)$, where $\Leftarrow \in \{\leftharpoonup, \leftharpoondown\}$ and $\mathcal{B}(X) = j{:}(p_1(X_1), \ldots, p_m(X_m), not\ p_{m+1}(X_{m+1}), \ldots, not\ p_n(X_n), \varphi)$, could be rewritten into the mapping rule $H \Leftarrow j{:}c(X)$ and the standard rule $j{:}c(X) \leftarrow \mathcal{B}(X)$.

Given a P2P system $PS = \{P_1, \ldots, P_n\}$, where $P_i = \langle D_i, LP_i, MP_i, IC_i\rangle$ with $i \in [1..n]$, the sets D, LP, MP, IC, $\overline{MP}$, and $\underline{MP}$ denote, respectively, the global sets of ground facts, standard rules, mapping rules, integrity constraints, maximal mapping rules, and minimal mapping rules, that is,

- $D = \bigcup_{i\in[1..n]} D_i$,
- $LP = \bigcup_{i\in[1..n]} LP_i$,
- $MP = \bigcup_{i\in[1..n]} MP_i$,
- $IC = \bigcup_{i\in[1..n]} IC_i$,
- $\overline{MP} = \bigcup_{i\in[1..n]} \overline{MP_i}$,
- $\underline{MP} = \bigcup_{i\in[1..n]} \underline{MP_i}$.

In the rest of the paper, with a little abuse of notation, PS will be denoted both with the tuple $\langle D, LP, MP, IC\rangle$ and with the set $D \cup LP \cup MP \cup IC$.

Moreover, we call a P2P system only containing maximal mapping rules, a *maximal P2P system* and a P2P system only containing minimal mapping rules, a *minimal P2P system*.

A peer $P_i = \langle D_i, LP_i, MP_i, IC_i\rangle$ is *locally consistent* if $SM(D_i \cup LP_i \cup IC_i) \neq \emptyset$. A P2P system whose peers are locally consistent is *locally consistent*. A peer (resp. P2P system) that is not locally consistent is *locally inconsistent*.

Given a mapping rule $r = H \leftharpoonup \mathcal{B}$ (resp. $r = H \leftharpoondown \mathcal{B}$), the corresponding standard logic rule $H \leftarrow \mathcal{B}$ will be denoted as $St(r)$.

Analogously, given a set of mapping rules MP, $St(MP) = \{St(r) \mid r \in MP\}$, and given a P2P system $PS = D \cup LP \cup MP \cup IC$, $St(PS) = D \cup LP \cup St(MP) \cup IC$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

In this context, an interpretation is a set of peer facts. The truth value of a peer fact (resp. literal, rule, maximal mapping rule, and minimal mapping rule) with respect to an interpretation M is as follows:

- $val_M(A) = (A \in M)$,
- $val_M(not\ A) = not\ val_M(A)$,
- $val_M(L_1, \ldots, L_n) = min\{val_M(L_1), \ldots, val_M(L_n)\}$,
- $val_M(H \leftarrow \mathcal{B}) = val_M(H) \geq val_M(\mathcal{B})$,
- $val_M(H \leftharpoonup \mathcal{B}) = val_M(H) \leq val_M(\mathcal{B})$,
- $val_M(H \rightharpoondown \mathcal{B}) = val_M(H) \leq val_M(\mathcal{B})$.

Therefore, while a standard rule is satisfied if its body is *false* or its body is *true* and its head is *true*, a mapping rule is satisfied if its body is *true* or its body is *false* and its head is *false*.

4 Semantics for P2P systems

Recent literature proposed different semantics for P2P systems that will be discussed in Section 7. The simplest semantics for a P2P system is the first-order logic (FOL) semantics obtained by interpreting mapping rules as standard rules. The FOL semantics of a P2P system $PS = \langle D, LP, MP, IC\rangle$ is given by the set of minimal models of $(D \cup LP \cup St(MP) \cup IC)$. The problem with the FOL semantics is that it leads to global inconsistency ($MM(D \cup LP \cup St(MP) \cup IC) = \emptyset$) as soon as an atom imported in a peer causes a violation of one of its integrity constraints.

It's clear that *more robust semantics*, derived by assuming *more flexible behaviors* of mapping rules, are needed.

It is worth noting that, in the FOL semantics, classical negation is used. In this paper, instead, we adopt negation as failure suited for all the nonmonotonic semantics here presented.

Our insight is that a peer of a P2P system can use its mapping rules to import from its neighborhood either *as much knowledge as possible preserving its consistency* or just *the knowledge that is strictly needed to restore the consistency of the system.*

Starting from this idea, in this section we first present two alternative semantics for P2P systems: the *Max Weak Model Semantics* and the *Min Weak Model Semantics.*

In the Max Weak Model Semantics, the peers of a P2P system only have maximal mapping rules and use them to import maximal sets of facts not violating local integrity constraints.

In the Min Weak Model Semantics, the peers of a P2P system only have minimal mapping rules and use them to import minimal sets of facts that are strictly needed to restore the consistency of the system.

In the Max-Min Weak Model Semantics, the peers of a P2P system have maximal and minimal mapping rules and unifies the two previous perspectives. A peer can use each neighbor as a resource either to enrich (integrate) or to fix (repair) its knowledge, adopting a kind of *integrate–repair* strategy.

All these semantics guarantee that a P2P system that is locally consistent admits at least a model, that is, remains consistent.

In order to present the three different semantics, first of all we introduce the concept of weak model that is common to all of them.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

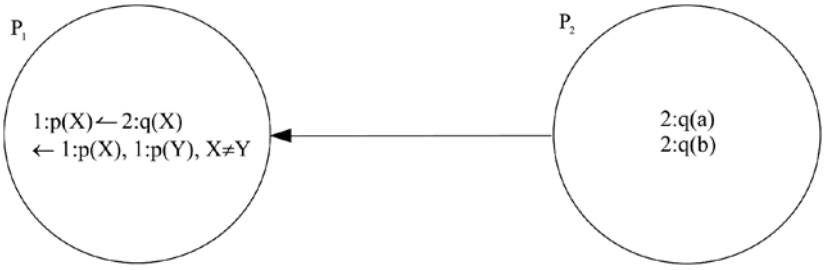


Fig. 4. A system PS.

Definition 3 (Weak Model). *Given a P2P system $PS = D \cup LP \cup MP \cup IC$, an interpretation M is a* weak model *for PS if $\{M\} = MM(St(PS^M))$, where PS^M is the program obtained from $ground(PS)$ by*

- *removing all peer rules r such that a negative literal not A occurs in the body of r and A is not in M;*
- *removing from the remaining peer rules each negative literal;*
- *removing all mapping rules whose head is not in M;*

The set of weak models of PS will be denoted as $WM(PS)$.

Observe that $St(PS^M)$ is an Horn program and it can be partitioned into a set of standard rules Π, a set of integrity constraints Σ and a set of facts D (i.e. $St(PS^M) = \Pi \cup \Sigma \cup D$).

As $\Pi \cup D$ is a positive normal program, it admits exactly one minimal model N. Therefore, N is the minimal model of $St(PS^M)$ if $N \models \Sigma$, otherwise $St(PS^M)$ does not admit any minimal model. If $N = M$ then M is a weak model of PS.

Note that the definition of weak model presents interesting analogies with the definition of stable model (see Section 2).

Indeed, given a logic program P, an interpretation M is a stable model of P if M is the minimal model of the reduct P^M, where the reduct is obtained by removing from $ground(P)$ each rule r such that a negative literal *not* A occurs in the body of r and A is not in M and removing from the remaining rules each negative literal. The fact that M is a minimal model of the reduct P^M ensures that each atom $H \in M$ is *supported*; that is, there is a rule in $ground(P)$ whose head is H and whose body is satisfied by M.

Similarly, in Definition 3 the fact that M is a minimal model of $St(PS^M)$ ensures that each atom is supported. In particular, for each mapping atom $H \in M$ there is a mapping rule in $ground(MP)$ whose head is H and whose body is satisfied by M.

Example 4.
Consider the P2P system PS depicted in Figure 4. P_2 contains the facts $2{:}q(a)$ and $2{:}q(b)$, whereas P_1 contains the maximal mapping rule $1{:}p(X) \leftarrow 2{:}q(X)$ and the constraint $\leftarrow 1{:}p(X), 1{:}p(Y), X \neq Y$. The weak models of the system are $M_1 = \{2{:}q(a), 2{:}q(b)\}$, $M_2 = \{2{:}q(a), 2{:}q(b), 1{:}p(a)\}$, and $M_3 = \{2{:}q(a), 2{:}q(b), 1{:}p(b)\}$.

We shall denote with $M[D]$ (resp. $M[LP]$, $M[MP]$, $M[\overline{MP}]$, $M[\underline{MP}]$) the set of ground atoms of M which are defined in D (resp. LP, MP, $\overline{MP}$, $\underline{MP}$).

Given a pair $P = (A, B)$, where A and B are generic objects, $P[1]$ (resp. $P[2]$) denotes the object A (resp. B).

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

The next proposition shows that for a P2P system $PS = D \cup LP \cup MP \cup IC$ having only positive rules in LP, checking if an interpretation M is a weak model is simpler because a simpler reduct involving only $ground(MP)$ can be used.

Proposition 1.
Given a P2P system $PS = D \cup LP \cup MP \cup IC$ s.t. no negation occurs in LP, an interpretation M is a weak model *for PS if and only if $\{M\} = MM(St(PS_M))$, where PS_M is the program obtained from $ground(PS)$ by removing all mapping rules whose head is not in M.*

Proof
$St(PS^M)$ can be obtained from $St(PS_M)$ by simply removing from $ground(IC)$ each negative literal *not* A s.t. $A \notin M$ that is if $St(PS_M) = \Pi \cup \overline{\Sigma} \cup D$, then $St(PS^M) = \Pi \cup \Sigma \cup D$, where Σ is obtained from $\overline{\Sigma}$ by removing each negative literal *not* A s.t. $A \notin M$.
($\Rightarrow$) As M is a weak model for PS, it is the minimal model of $\Pi \cup D$ and $M \models \Sigma$. As all the negative literals *not* A occurring in $\overline{\Sigma}$ are s.t. $A \notin M$, it follows that $M \models \overline{\Sigma}$. Therefore, M is a minimal model of $St(PS_M)$.
($\Leftarrow$) If M is the minimal model of $St(PS_M)$, it is the minimal model of $\Pi \cup D$ and $M \models \overline{\Sigma}$. Let us consider an integrity constraint $ic \in \Sigma$ (observe that, the body of ic only contains positive literals). There are the following cases:

- $ic \in ground(IC)$. In this case, no negative atom has been removed from the original integrity constraint in order to obtain ic. We have that $ic \in \overline{\Sigma}$ and $M \models ic$.
- $ic \notin ground(IC)$. In this case ic has been obtained from a ground integrity constraint $\overline{ic} \in ground(IC)$. All negative literals *not* A removed from $\overline{ic}$ in order to obtain ic are not in M (otherwise ic could not belong to Σ). Moreover, $\overline{ic} \in \overline{\Sigma}$. As $M \models \overline{ic}$ and each negative literal occurring in $\overline{ic}$ is true w.r.t. M, at least a positive literal occurring in $\overline{ic}$ has to be false w.r.t. M. It follows that $M \models ic$.

Therefore, $M \models \Sigma$ and M is a minimal model of $St(PS^M)$; that is, M is a weak model for PS. □

4.1 Max Weak Model Semantics

In previous works (Caroprese *et al.* 2006a; Caroprese *et al.* 2006b; Caroprese and Zumpano 2007; Caroprese and Zumpano 2008; Caroprese and Zumpano 2012a), the authors introduced the *Max Weak Model Semantics*.

We recall that a *maximal mapping rule* (see Definition 1) is of the form $H \leftharpoonup \mathcal{B}$. Intuitively, $H \leftharpoonup \mathcal{B}$ means that if the body conjunction $\mathcal{B}$ is *true* in the source peer, the atom H will be imported in the target peer (i.e. H is *true* in the target peer) only if it does not imply (directly or indirectly) the violation of some constraints. In this section, we assume that all mapping rules of a P2P system $PS = \langle D, LP, MP, IC \rangle$ are maximal mapping rules; that is, PS is a maximal P2P system.

Example 5.
Consider a P2P system consisting of two peers P_1 and P_2, where
$P_1 = \langle \{1{:}q(b)\}, \emptyset, \emptyset, \emptyset, \emptyset \rangle$,
$P_2 = \langle \{2{:}s(a)\}, \{2{:}r(X) \leftarrow 2{:}p(X);\ \ 2{:}r(X) \leftarrow 2{:}s(X)\},$
$\{2{:}p(X) \leftharpoonup 1{:}q(X)\},\ \{\leftarrow 2{:}r(X), 2{:}r(Y), X \neq Y\}\rangle$.
In this case, the fact $2{:}p(b)$ cannot be imported in P_1 as it indirectly violates the integrity constraint.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Definition 4 (*Maximal Weak Model*). Given a maximal P2P system PS and two weak models M and N of PS, M is said *max-preferable* to N, and is denoted as $M \sqsupseteq_{Max} N$, if $M[MP] \supseteq N[MP]$. Moreover, if $M \sqsupseteq_{Max} N$ and $N \not\sqsupseteq_{Max} M$, then $M \sqsupset_{Max} N$. A weak model M of PS is *maximal* if there is no weak model N of PS such that $N \sqsupset_{Max} M$. The set of maximal weak models of PS will be denoted as $MaxWM(PS)$.

In the Max Weak Model Semantics, peers import maximal sets of facts not violating integrity constraints. Therefore, each peer of the system can be thought as an *integration resource*. We will show that a locally consistent P2P system always admits a maximal weak model while a locally inconsistent P2P system not always has this property. A generalization of our semantics that guarantees the existence of at least a model even for locally inconsistent P2P system will be presented in a following section.

Example 6.
In Example 4, the maximal weak models are M_2 and M_3.

The Max Weak Model Semantics easily allows to express a classical problem, the *three-colorability problem*, as follows.

Example 7.
Three-colorability. We are given two peers: P_1, containing a set of nodes and a set of colors which are defined by the unary relations 1:*node* and 1:*color* respectively, and P_2, containing a set of edges defined by the binary relation 2:*edge*, the mapping rule:

$$2{:}colored(X, C) \quad \leftarrow 1{:}node(X), 1{:}color(C)$$

and the integrity constraints:

$$\leftarrow 2{:}colored(X, C_1),\ 2{:}colored(X, C_2),\ C_1 \neq C_2,$$
$$\leftarrow 2{:}edge(X, Y),\ 2{:}colored(X, C),\ 2{:}colored(Y, C),$$

stating, respectively, that a node cannot be colored with two different colors and two connected nodes cannot be colored with the same color. The mapping rule states that the node x can be colored with the color c, only if in doing this no constraint is violated, that is if the node x is colored with a unique color and there is no adjacent node colored with the same color. Each maximal weak model computes a maximal subgraph which is three-colorable.

The following proposition shows an important property of relation $\sqsupseteq_{Max}$.

Proposition 2.
For any maximal P2P system $PS = D \cup LP \cup MP \cup IC$ s.t. no negation occurs in LP, $\sqsupseteq_{Max}$ defines a partial order on the set of weak models of PS.

Proof
We prove that relation $\sqsupseteq_{Max}$ is antisymmetric and transitive.

- *(Antisymmetry)* Let us consider two weak models M and N in $WM(PS)$. We prove that if $M \sqsupseteq_{Max} N$ and $N \sqsupseteq_{Max} M$, then $M = N$. As $M \sqsupseteq_{Max} N$, then $M[MP] \supseteq N[MP]$. Similarly, as $N[MP] \sqsupseteq_{Max} M[MP]$, then $N[MP] \supseteq M[MP]$. It follows that $N[MP] = M[MP]$. As M and N are weak models, by Proposition 1, $\{M\} = MM(St(PS_M))$ and $\{N\} = MM(St(PS_N))$. Moreover, as $N[MP] = M[MP]$, $St(PS_M) = St(PS_N)$. It follows that $M = N$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

- *(Transitivity)* Let us consider three weak models M, N, and S in $WM(PS)$. We prove that if $M \sqsupseteq_{Max} N$ and $N \sqsupseteq_{Max} S$, then $M \sqsupseteq_{Max} S$.
  As $M \sqsupseteq_{Max} N$, then $M[MP] \supseteq N[MP]$. Similarly, as $N[MP] \sqsupseteq_{Max} S[MP]$, then $N[MP] \supseteq S[MP]$. It follows that $M[MP] \supseteq S[MP]$ and then $M \sqsupseteq_{Max} S$.

□

If the standard rules of a P2P system contain negation, in general $\sqsupseteq_{Max}$ is not antisymmetric. To prove it, let's consider a P2P system only containing a peer P_1, with the standard rules $1{:}p \leftarrow not\ 1{:}q$ and $1{:}q \leftarrow not\ 1{:}p$. The P2P system admits two weak models: $M = \{1{:}p\}$ and $N = \{1{:}q\}$. As $MP[M] = MP[N] = \emptyset$, it follows that $M \sqsupseteq_{Max} N$ and $N \sqsupseteq_{Max} M$, but $M \neq N$.

The next theorem shows that consistent maximal P2P systems always admit maximal weak models.

Theorem 1.
For every locally consistent maximal P2P system, $MaxWM(PS) \neq \emptyset$.

Proof
Let us consider a set M such that $\{M\} \in MM(D \cup LP \cup IC)$, that is, a minimal model of a P2P system obtained from PS by deleting all mapping rules. As PS is locally consistent, such a model exists. Let Π be the logic program obtained by deleting from $ground(D \cup LP \cup IC)$ all peer rules whose body is false w.r.t. M and by removing from the remaining rules the negative literals (observe that, they are *true* w.r.t. M). Π is an Horn program and admits only one minimal model. This minimal model has to be M. Moreover, as M does not contain any mapping atoms, $\Pi = St(PS^M)$. It follows that $\{M\} = MM(St(PS^M))$. This means that M is a weak model for PS. As there is at least a weak model of PS, then $MaxWM(PS) \neq \emptyset$. □

If a P2P system contains at least a locally inconsistent peer, the Max Weak Model semantics does not guarantee the existence of a maximal weak model.

Example 8.
Let us consider a P2P system containing only peer $P_1 = \langle \{1{:}a, 1{:}b\}, \emptyset, \emptyset, \{\leftarrow 1{:}a, 1{:}b\} \rangle$. Clearly, P_1 is locally inconsistent and there is no way to import mapping atoms able to restore its consistency. Observe that, the only way to make the peer consistent is to remove at least one fact from its local database. In the following, we present an extension of our framework allowing deletions of facts from local databases.

4.1.1 An alternative characterization of the Max Weak Model Semantics

In this section, we present an alternative characterization of the Max Weak Model Semantics based on the rewriting of mapping rules into prioritized rules (Brewka *et al.* 2003; Sakama and Inoue 2000).

Definition 5 (*Rewriting of a Maximal P2P System into a Prioritized Logic Program*).
Given a maximal P2P system $PS = D \cup LP \cup MP \cup IC$ and a maximal mapping rule $r = i{:}p(x) \leftarrow \mathcal{B}$, then

- $Rew(r)$ denotes the pair $(i{:}p(x) \oplus i{:}p'(x) \leftarrow \mathcal{B},\ i{:}p(x) \succeq i{:}p'(x))$, consisting of a disjunctive mapping rule and a priority statement,
- $Rew(MP) = (\{Rew(r)[1] |\ r \in MP\}, \{Rew(r)[2] |\ r \in MP\})$, and
- $Rew(PS) = (D \cup LP \cup Rew(MP)[1] \cup IC, Rew(MP)[2])$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

In the above definition, the atom $i{:}p(x)$ (resp. $i{:}p'(x)$) means that the fact $i{:}p(x)$ is imported (resp. not imported) in the peer P_i.

Intuitively, the rewriting of the maximal mapping rule states that if $\mathcal{B}$ is *true* in the *source peer* then two alternative actions can be performed in the *target peer*: $i{:}p(x)$ can be either imported or not imported; but the presence of the priority statement $i{:}p(x) \succeq i{:}p'(x)$) establishes that the action of *importing* $i{:}p(x)$ is preferable over the action of *not importing* $i{:}p(x)$.

Example 9.
The rewriting of the P2P system in Example 4 is

$$\begin{aligned} Rew(PS) = (\{&2{:}q(a),\ 2{:}q(b),\\ &1{:}p(X) \oplus 1{:}p'(X) \leftarrow 2{:}q(X),\\ &\leftarrow 1{:}p(X),\ 1{:}p(Y),\ X \neq Y\},\\ &\{1{:}p(X) \succeq 1{:}p'(X\}). \end{aligned}$$

$Rew(PS)[1]$ has three stable models:
$M_0 = \{2{:}q(a),\ 2{:}q(b),\ 1{:}p'(a),\ 1{:}p'(b)\}$,
$M_1 = \{2{:}q(a),\ 2{:}q(b),\ 1{:}p(a),\ 1{:}p'(b)\}$,
$M_2 = \{2{:}q(a),\ 2{:}q(b),\ 1{:}p'(a),\ 1{:}p(b)\}$.

The set of preferred stable models are $\{M_1, M_2\}$.

Example 10.
The rewriting of the mapping rule of Example 7 consists of the rule:

$$2{:}colored(X,C) \oplus 2{:}colored'(X,C) \leftarrow 1{:}node(X), 1{:}color(C),$$

and the preference:

$$2{:}colored(X,C) \succeq 2{:}colored'(X,C).$$

Given a maximal P2P system PS and a preferred stable model M for $Rew(PS)$, we denote with $St(M)$ the subset of non-primed atoms of M and we say that $St(M)$ is a preferred stable model of PS. We denote the set of preferred stable models of $Rew(PS)$ as $PSM(PS)$. The following theorem shows the equivalence of preferred stable models and maximal weak models.

Theorem 2 (*Equivalence Between Preferred Stable Models and Maximal Weak Models*).
For every maximal P2P system PS, $PSM(PS) = MaxWM(PS)$.

Proof

1. $(PSM(PS) \subseteq MaxWM(PS))$
   Let $M \in PSM(PS)$ and $N = St(M)$. First we prove that N is a weak model. Let us consider a ground mapping rule m and its rewriting $Rew(m)$. The rule $Rew(m)[1] = A \oplus A' \leftarrow \mathcal{B}$ is equivalent to the rules $r = A \leftarrow \mathcal{B} \wedge not\ A'$, $r' = A' \leftarrow \mathcal{B} \wedge not\ A$ and the constraint $\leftarrow A, A'$. There are three cases:
   - $A, A' \notin M$. In this case, $M \not\models \mathcal{B}$. Then the bodies of r and r' are false and so $r, r' \notin (Rew(PS)[1])^M$.
   - $A \in M$ and $A' \notin M$. In this case the body of r' is false and $r' \notin (Rew(PS)[1])^M$. Moreover, $A \leftarrow \mathcal{B} \in (Rew(PS)[1])^M$.
   - $A' \in M$ and $A \notin M$. In this case the body of r is false and $r \notin (Rew(PS)[1])^M$. Moreover, $A' \leftarrow \mathcal{B} \in (Rew(PS)[1])^M$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Then, by construction, we have that $(Rew(PS)[1])^M = St(PS^N) \cup \{A' \leftarrow \mathcal{B} \mid A' \in M \ \wedge \ A \leftharpoonup \mathcal{B} \in ground(PS)\}$.
We have that

- the minimal model of $(Rew(PS)[1])^M$ is M, as M is a stable model of $(Rew(PS)[1])^M$;
- $M = N \cup \{A' \mid A' \in M\}$;
- non-primed atom A can be only inferred by rules in $St(PS^N)$; and
- no primed atom A' occurs in the body of any rule of $St(PS^N)$.

Therefore, the minimal model of $St(PS^N)$ is N and N is a weak model of PS.
Now we prove that N is a maximal weak model of PS. Let us assume by contradiction that there is a weak model L such that $L[MP] \supset N[MP]$. Then the ground mapping rules that will be deleted from $ground(PS)$ to derive PS^L are a subset of those that will be deleted to derive PS^N.

Let us consider the set $K = L \cup \{A' \mid A \notin L \ \wedge \ A \leftharpoonup \mathcal{B} \in ground(PS) \ \wedge \ L \models \mathcal{B}\}$. By construction, K is the minimal model of $Rew(PS)[1]^K$. Then K is a stable model of $Rew(PS)[1]$.

Observe that there must exist two atoms $A \in K$ and $A' \in M$ and, by construction, there cannot exist two atoms $B \in M$ and $B' \in, K$. Moreover, $ground(Rew(PS)[2])$ contains the preference $A \succ A'$. Therefore, $K \sqsupset M$ and M is not a preferred stable model of $Rew(PS)$. This is a contradiction.

2. $(PSM(PS) \supseteq MaxWM(PS))$
Let $N \in MaxWM(PS)$ and $M = N \cup \{A' \mid A \notin N \wedge A \leftharpoonup \mathcal{B} \in ground(PS) \wedge N \models \mathcal{B}\}$. First we prove that M is a stable model of $Rew(PS)[1]$; that is, it is the minimal model of $(Rew(PS)[1])^M$.
By construction,
$(Rew(PS)[1])^M = St(PS^N) \cup \{A' \leftarrow \mathcal{B} \mid A \notin N \wedge A \leftharpoonup \mathcal{B} \in ground(PS) \wedge N \models \mathcal{B}\}$.
We have that

- the minimal model of $St(PS)^N$ is N, as N is a weak model of PS;
- the minimal model of $\{A' \leftarrow \mathcal{B} \mid A \notin N \ \wedge \ A \leftharpoonup \mathcal{B} \in ground(PS) \ \wedge \ N \models \mathcal{B}\}$ is $M \setminus N$;
- non-primed atoms A can be only inferred by rules in $St(PS^N)$; and
- no primed atom A' occurs in the body of any rule of $St(PS^N)$.

Therefore, the minimal model of $(Rew(PS)[1])^M$ is M and M is a stable model of $Rew(PS)[1]$.

Now we prove that M is a preferred stable model for $Rew(PS)$. Let us assume by contradiction that there is a stable model L for $Rew(PS)$ s.t. $L \ \sqsupset \ M$. From point 1 and preferences in $Rew(PS)[2]$, we have that $St(L)$ is a weak model for PS and $St(L)[MP] \supset St(M)[MP]$; that is, $St(M)$ is not a maximal weak model for PS. This is a contradiction. □

This characterization of the Max Weak Model Semantics makes evident that *importing* a mapping atom is preferable over *not importing* it and provides a computational mechanism allowing to derive the maximal weak models of a maximal P2P system.

Example 11.
Consider the P2P system of Example 9, we have
$PSM(PS) = \{\{2{:}q(a),\ 2{:}q(b),\ 1{:}p(a)\}, \{2{:}q(a),\ 2{:}q(b),\ 1{:}p(b)\}\}$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

This example shows that the preferred stable models of PS coincide with its maximal weak models.

4.2 Min Weak Model Semantics

In Caroprese and Zumpano (2012b) the authors introduced the *Min Weak Model Semantics*: a peer can be locally inconsistent and the P2P system it joins provides support to restore its consistency. The basic idea, yet very simple, is the following: an inconsistent peer, in the interaction with other peers, just imports the missing part of its local database which is *correct*, but *incomplete*.

The proposal of the Min Weak Model Semantics stems from the observations that in real-world P2P systems, peers often use the available import mechanisms to extract knowledge from the rest of the system only if this knowledge is strictly needed to repair inconsistencies of the system.

A *minimal mapping rule* (see Definition 1) is of the form $H \leftarrowtail \mathcal{B}$. Intuitively, $H \leftarrowtail \mathcal{B}$ means that if the body conjunction $\mathcal{B}$ is *true* in the source peer, the atom H is imported in the target peer (that is H is *true* in the target peer) only if it implies (directly or indirectly) the satisfaction of some constraints that otherwise would be violated. In this section, we assume that all mapping rules of a P2P system $PS = \langle D, LP, MP, IC \rangle$ are minimal mapping rules; that is, PS is a minimal P2P system.

Definition 6 (*Minimal Weak Model*). Given a minimal P2P system PS and two weak models M and N of PS, M is said *min-preferable* to N, and is denoted as $M \sqsupseteq_{Min} N$, if $M[MP] \subseteq N[MP]$. Moreover, if $M \sqsupseteq_{Min} N$ and $N \not\sqsupseteq_{Min} M$, then $M \sqsupset_{Min} N$. A weak model M of PS is said to be *minimal* if there is no weak model N of PS such that $N \sqsubset_{Min} M$. The set of minimal weak models will be denoted by $MinWM(PS)$.

Next example will clarify the concept of minimal weak model.

Example 12.
Consider the P2P system PS presented in Example 2. The weak models of the system are the following:

$$M_1 = \{1{:}vendor(dan, laptop),\ 1{:}vendor(bob, laptop),\ 2{:}order(laptop),\\ 2{:}supplier(dan, laptop),\ 2{:}available(laptop)\},$$

$$M_2 = \{1{:}vendor(dan, laptop),\ 1{:}vendor(bob, laptop),\ 2{:}order(laptop),\\ 2{:}supplier(bob, laptop),\ 2{:}available(laptop)\}, \text{ and}$$

$$M_3 = \{1{:}vendor(dan, laptop),\ 1{:}vendor(bob, laptop),\ 2{:}order(laptop),\\ 2{:}supplier(dan, laptop),\ 2{:}supplier(bob, laptop),\ 2{:}available(laptop)\},$$

whereas the minimal weak models are M_1 and M_2 because they contain minimal subsets of mapping atoms (resp. $2{:}\{supplier(dan, laptop)\}$ and $\{2{:}supplier(bob, laptop)\}$).

We observe that, adopting the Min Weak Model Semantics, if each peer of a P2P system is locally consistent then no mapping atom is inferred. Clearly, not always a minimal weak model exists. This happens when there is at least a peer which is locally inconsistent and there is no way to import mapping atoms that could repair its local database so that its consistency can be restored.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

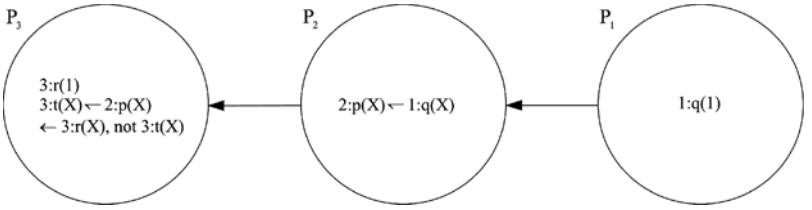


Fig. 5. A P2P system with maximal and minimal mapping rules.

Example 13.
Let us consider the simple P2P system presented in Example 8. Also adopting the Min Weak Model Semantics, PS does not admit any minimal weak model.

It is important to observe that a peer uses its minimal mapping rules to import minimal sets of atoms, allowing the satisfaction of integrity constraints belonging not only to it but also to other peers.

Example 14.
Consider the P2P system depicted in Figure 5.

- Peer P_1 contains the fact $1{:}q(1)$.
- Peer P_2 contains the minimal mapping rule $2{:}p(X) \leftarrowtail 1{:}q(X)$.
- Peer P_3 contains the fact $3{:}r(1)$, the minimal mapping rule $1{:}t(X) \leftarrowtail 2{:}p(X)$ and the integrity constraint $\leftarrow 3{:}r(X),\ not\ 3{:}t(X)$.

Peer P_2 imports the atom $2{:}p(1)$ from P_1 to guarantee the satisfaction of an integrity constraint belonging to P_3. Min Weak Model Semantics assigns to the system its unique minimal weak model: $\{1{:}q(1),\ 2{:}p(1), 3{:}r(1), 3{:}t(1)\}$.

The following proposition shows an important property of relation $\sqsupseteq_{Min}$.

Proposition 3.
For any maximal P2P system, $PS = D \cup LP \cup MP \cup IC$ s.t. no negation occurs in LP, $\sqsupseteq_{Min}$ defines a partial order on the set of weak models of PS.

Proof
We prove that relation $\sqsupseteq_{Min}$ is antisymmetric and transitive.

- *(Antisymmetry)* Let us consider the weak models M and N in PS. We prove that if $M \sqsupseteq_{Min} N$ and $N \sqsupseteq_{Min} M$, then $M = N$. As $M \sqsupseteq_{Min} N$, then $M[MP] \subseteq N[MP]$. Similarly, as $N[MP] \sqsupseteq_{Min} M[MP]$, then $N[MP] \subseteq M[MP]$. It follows that $N[MP] = M[MP]$. As M and N are weak models, by Proposition 1, $\{M\} = MM(St(PS_M))$ and $\{N\} = MM(St(PS_N))$. Moreover, as $N[MP] = M[MP]$, $St(PS_M) = St(PS_N)$. It follows that $M = N$.
- *(Transitivity)* Let us consider the weak models M, N, and S in PS. We prove that if $M \sqsupseteq_{Min} N$ and $N \sqsupseteq_{Min} S$, then $M \sqsupseteq_{Min} S$.
  As $M \sqsupseteq_{Min} N$, then $M[MP] \subseteq N[MP]$. Similarly, as $N[MP] \sqsupseteq_{Min} S[MP]$, then $N[MP] \subseteq S[MP]$. It follows that $M[MP] \subseteq S[MP]$ and then $M \sqsupseteq_{Min} S$.

□

In the Min Weak Model Semantics, a peer may import an atom from a neighbor peer even if such atom is not needed to repair its own inconsistency, but is needed to restore the consistency of a different peer. In this way, the system behaves in a global way.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

4.2.1 An alternative characterization of the Min Weak Model Semantics

Similarly to Section 4.1.1, here we present an alternative characterization of the Min Weak Model Semantics based on the rewriting of mapping rules into prioritized rules (Brewka *et al.* 2003; Sakama and Inoue 2000).

Definition 7 (*Rewriting of a Maximal P2P System into a Prioritized Logic Program*). Given a minimal P2P system PS and a mapping rule $r = i{:}p(x) \leftarrow \mathcal{B}$, then

- $Rew(r)$ denotes the pair $(i{:}p(x) \oplus i{:}p'(x) \leftarrow \mathcal{B},\ i{:}p'(x) \succeq i{:}p(x))$, consisting of a disjunctive mapping rule and a priority statement,
- $Rew(MP) = (\{Rew(r)[1] |\ r \in MP\}, \{Rew(r)[2] |\ r \in MP\})$, and
- $Rew(PS) = (D \cup LP \cup Rew(MP)[1] \cup IC,\ Rew(MP)[2])$.

In the above definition, the atom $i{:}p(x)$ (resp. $i{:}p'(x)$) means that the fact $i{:}p(x)$ is imported (resp. not imported) in the peer P_i.

Intuitively, the rewriting of the mapping rule states that if $\mathcal{B}$ is *true* in the *source peer,* then two alternative actions that can be performed in the *target peer*: $i{:}p(x)$ can be either imported or not imported; but the presence of the priority statement,$i{:}p'(x) \succeq i{:}p(x))$ establishes that the action of *not importing* $i{:}p(x)$ is preferable over the action of *importing* $i{:}p(x)$.

Example 15.
Consider again the system presented in Example 2. The rewriting of the system is

$$\begin{aligned} Rew(PS) = \{&1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}order(laptop),\\ &2{:}supplier(X,Y) \oplus 2{:}supplier'(X,Y) \leftarrow 1{:}vendor(X,Y),\\ &1{:}available(Y) \leftarrow 1{:}supplier(X,Y),\\ &\leftarrow 1{:}order(X),\ not\ 1{:}available(X)\},\\ &\{1{:}supplier'(X,Y) \succeq 1{:}supplier(X,Y)\}). \end{aligned}$$

$Rew(PS)[1]$ has three stable models:

$$\begin{aligned} M_1 = \{&1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}order(laptop),\\ &2{:}supplier(dan, laptop), 2{:}supplier'(bob, laptop), 2{:}available(laptop)\},\\ M_2 = \{&1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}order(laptop),\\ &2{:}supplier'(dan, laptop), 2{:}supplier(bob, laptop), 2{:}available(laptop)\},\\ M_3 = \{&1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}order(laptop),\\ &2{:}supplier(dan, laptop), 2{:}supplier(bob, laptop), 2{:}available(laptop)\}. \end{aligned}$$

The preferred stable models are M_1 and M_2.

The following theorem shows the equivalence of preferred stable models and minimal weak models.

Theorem 3 (*Equivalence Between Preferred Stable Models and Minimal Weak Models*). For every minimal P2P system PS, $PSM(PS) = MinWM(PS)$.

Proof

1. $(PSM(PS) \subseteq MinWM(PS))$
   Let $M \in PSM(PS)$ and $N = St(M)$. First we prove that N is a weak model. Let us consider a ground mapping rule m and its rewriting $Rew(m)$. The rule $Rew(m)[1] =$

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

$A \oplus A' \leftarrow \mathcal{B}$ is equivalent to the rules $r = A \leftarrow \mathcal{B} \wedge not\ A'$, $r' = A' \leftarrow \mathcal{B} \wedge not\ A$, and the constraint $\leftarrow A, A'$. There are three cases:

- $A, A' \notin M$. In this case, $M \not\models \mathcal{B}$. Then the bodies of r and r' are false and so $r, r' \notin (Rew(PS)[1])^M$.
- $A \in M$ and $A' \notin M$. In this case, the body of r' is false and $r' \notin (Rew(PS)[1])^M$. Moreover, $A \leftarrow \mathcal{B} \in (Rew(PS)[1])^M$.
- $A' \in M$ and $A \notin M$. In this case, the body of r is false and $r \notin (Rew(PS)[1])^M$. Moreover, $A' \leftarrow \mathcal{B} \in (Rew(PS)[1])^M$.

Then, by construction, we have that $(Rew(PS)[1])^M = St(PS^N) \cup \{A' \leftarrow \mathcal{B} \mid A' \in M \wedge A \leftarrow \mathcal{B} \in ground(PS)\}$.
We have that

- the minimal model of $(Rew(PS)[1])^M$ is M, as M is a stable model of $(Rew(PS)[1])^M$;
- $M = N \cup \{A' \mid A' \in M\}$;
- non-primed atoms A can be only inferred by rules in $St(PS^N)$; and
- no primed atom A' occurs in the body of any rule of $St(PS^N)$.

Therefore, the minimal model of $St(PS^N)$ is N and N is a weak model of PS.

Now we prove that N is a minimal weak model of PS. Let us assume by contradiction that there is a weak model L such that $L[MP] \subset N[MP]$. Then the ground mapping rules that will be deleted from $ground(PS)$ to derive PS^L are a superset of those that will be deleted to derive PS^N.

Let us consider the set $K = L \cup \{A' \mid A \notin L \wedge A \leftarrow \mathcal{B} \in ground(PS) \wedge L \models \mathcal{B}\}$. By construction, K is the minimal model of $Rew(PS)[1]^K$. Then K is a stable model of $Rew(PS)[1]$. Observe that must exist two atoms $A' \in K$ and $A \in M$ and, by construction, there cannot exist two atoms $B' \in M$ and $B \in K$. Moreover, $ground(Rew(PS)[2])$ contains the preference $A' \succ A$. Therefore, $K \sqsupset M$ and M is not a preferred stable model of $Rew(PS)$. This is a contradiction.

2. $(PSM(PS) \supseteq MinWM(PS))$
Let $N \in MinWM(PS)$ and $M = N \cup \{A' \mid A \notin N \wedge A \leftarrow \mathcal{B} \in ground(PS) \wedge N \models \mathcal{B}\}$. First we prove that M is a stable model of $Rew(PS)[1]$; that is, it is the minimal model of $(Rew(PS)[1])^M$.
By construction, $(Rew(PS)[1])^M = St(PS^N) \cup \{A' \leftarrow \mathcal{B} \mid A \notin N \wedge A \leftarrow \mathcal{B} \in ground(PS) \wedge N \models \mathcal{B}\}$. We have that

- the minimal model of $St(PS)^N$ is N, as N is a weak model of PS;
- the minimal model of $\{A' \leftarrow \mathcal{B} \mid A \notin N \wedge A \leftarrow \mathcal{B} \in ground(PS) \wedge N \models \mathcal{B}\}$ is $M \setminus N$;
- non-primed atoms A can be only inferred by rules in $St(PS^N)$; and
- no primed atom A' occurs in the body of any rule of $St(PS^N)$.

Therefore, the minimal model of $(Rew(PS)[1])^M$ is M and M is a stable model of $Rew(PS)[1]$.

Now we prove that M is a preferred stable model for $Rew(PS)$. Let us assume by contradiction that there is a stable model L for $Rew(PS)$ s.t. $L \sqsupset M$. From point 1 and preferences in $Rew(PS)[2]$, we have that $St(L)$ is a weak model for PS and $St(L)[MP] \subset St(M)[MP]$; that is, $St(M)$ is not a minimal weak model for PS. This is a contradiction.

□

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

This characterization of the Min Weak Model Semantics makes evident that *not importing* a mapping atom is preferable over *importing* it and provides a computational mechanism allowing to derive the minimal weak models of a minimal P2P system.

4.3 Max-Min Weak Model Semantics

This section presents a unified semantics for P2P systems, the Max-Min Weak Model Semantics, that represents a generalization of those introduced in Sections 4.1 and 4.2. A peer, which can be locally inconsistent, can use two import mechanisms for importing knowledge from other peers: maximal mapping rules to import maximal sets of mapping atoms not violating local integrity constraints and minimal mapping rules to restore local consistency. These two mechanisms can be combined and used in the same peer. With this semantics, a peer can consider each of its neighbors as a resource used either to enrich (integrate) or to fix (repair) its knowledge, so as to define a kind of *integrate–repair* strategy.

Example 16.
Consider again the P2P system presented in Example 3. As we observed, peer P_3 is locally consistent. It imports from P_1 all the orders that can be satisfied by suppliers imported from peer P_2. Moreover, a minimum set of suppliers will be imported in P_3.

The fact 3:$order(laptop)$ will be imported in P_3 from P_1 because there is at least a supplier of the object "laptop" that can be imported in P_3 from P_2. Instead, there is no way to import a supplier of the object "monitor." Therefore, the fact 3:$order(monitor)$ will not be imported in P_3. Finally, there are two possible ways to import a supplier of the object "laptop": importing from P_1 either the fact 3:$supplier(dan, laptop)$ or the fact 3:$supplier(bob, laptop)$.

Definition 8 (*Max-Min Weak Model*). Given a P2P system PS and two weak models M and N of PS, we say that M is *max-min-preferable* to N, and we write $M \sqsupseteq N$, if

- $M[\overline{MP}] \supset N[\overline{MP}]$ or
- $M[\overline{MP}] = N[\overline{MP}]$ and $M[\underline{MP}] \subseteq N[\underline{MP}]$.

Moreover, if $M \sqsupseteq N$ and $N \not\sqsupseteq M$, we write $M \sqsupset N$. A weak model M is said to be *max-min* if there is no weak model N such that $N \sqsubset M$. The set of max-min weak models will be denoted by $MaxMinWM(PS)$.

The above definition states that a weak model is a max-min weak model if it maximizes the set of atoms imported by means of maximal mapping rules while minimizing the set of atoms imported my means of minimal mapping rules (used to maintain local consistency). The approach follows the *classical* and *natural* strategy of enriching as much as possible the knowledge of an information source (by means of maximal mapping rules) guaranteeing consistency (by using minimal mapping rules).

Example 17.
Consider the P2P system PS presented in Example 3. The weak models of the system are the following:

$M_1 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor)\},$

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

$M_2 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}shopping(laptop),$
$2{:}shopping(monitor), 3{:}supplier(dan, laptop), 3{:}available(laptop)\},$
$M_3 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}shopping(laptop),$
$2{:}shopping(monitor), 3{:}supplier(bob, laptop), 3{:}available(laptop)\},$
$M_4 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}shopping(laptop),$
$2{:}shopping(monitor), 3{:}supplier(dan, laptop), 3{:}supplier(bob, laptop),$
$3{:}available(laptop)\},$
$M_5 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}shopping(laptop),$
$2{:}shopping(monitor), 3{:}supplier(dan, laptop), 3{:}order(laptop),$
$3{:}available(laptop)\},$
$M_6 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}shopping(laptop),$
$2{:}shopping(monitor), 3{:}supplier(bob, laptop), 3{:}order(laptop),$
$3{:}available(laptop)\},$
$M_7 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop), 2{:}shopping(laptop),$
$2{:}shopping(monitor), 3{:}supplier(dan, laptop), 3{:}supplier(bob, laptop),$
$3{:}order(laptop), 3{:}available(laptop)\},$

whereas the max-min weak models are M_5 and M_6.

The following proposition shows an important property of relation $\sqsupseteq$. It is easy to show that a locally consistent P2P system always admits a max-min weak model.

Proposition 4.
For any maximal P2P system $PS = D \cup LP \cup MP \cup IC$ s.t. no negation occurs in LP, $\sqsupseteq$ defines a partial order on the set of weak models of PS.

Proof
We prove that relation $\sqsupseteq$ is antisymmetric and transitive.

- *(Antisymmetry)* Let us consider the weak models M and N in PS. We prove that if $M \sqsupseteq N$ and $N \sqsupseteq M$, then $M = N$. As $M \sqsupseteq N$, then $M[\overline{MP}] \supset N[\overline{MP}]$ or $M[\overline{MP}] = N[\overline{MP}]$ and $M[\underline{MP}] \subseteq N[\underline{MP}]$. Similarly, as $N \sqsupseteq M$, then $N[\overline{MP}] \supset M[\overline{MP}]$ or $N[\overline{MP}] = M[\overline{MP}]$ and $N[\underline{MP}] \subseteq M[\underline{MP}]$. As the conditions $M[\overline{MP}] \supset N[\overline{MP}]$ and $N[\overline{MP}] \supset M[\overline{MP}]$ cannot hold at the same time, it follows that $N[\overline{MP}] = M[\overline{MP}]$ and $N[\underline{MP}] = M[\underline{MP}]$, that is, $N[MP] = M[MP]$. As M and N are weak models, by Proposition 1, $\{M\} = MM(St(PS_M))$ and $\{N\} = MM(St(PS_N))$. Moreover, as $N[MP] = M[MP]$, $St(PS_M) = St(PS_N)$. It follows that $M = N$.
- *(Transitivity)* Let us consider the weak models M, N, and S in PS. We prove that if $M \sqsupseteq N$ and $N \sqsupseteq S$, then $M \sqsupseteq S$.
  If $M[\overline{MP}] \supset N[\overline{MP}]$ and $N[\overline{MP}] \supset S[\overline{MP}]$, then $M[\overline{MP}] \supset S[\overline{MP}]$ and so $M \sqsupseteq S$.
  If $M[\overline{MP}] \supset N[\overline{MP}]$, $N[\overline{MP}] = S[\overline{MP}]$, and $N[\underline{MP}] \subseteq S[\underline{MP}]$, then $M[\overline{MP}] \supset S[\overline{MP}]$ and so $M \sqsupseteq S$.
  If $M[\overline{MP}] = N[\overline{MP}]$, $M[\underline{MP}] \subseteq N[\underline{MP}]$, and $N[\overline{MP}] \supset S[\overline{MP}]$, then $M[\overline{MP}] \supset S[\overline{MP}]$ and so $M \sqsupseteq S$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

If $M[\overline{MP}] = N[\overline{MP}]$, $M[\underline{MP}] \subseteq N[\underline{MP}]$, $N[\overline{MP}] = S[\overline{MP}]$, and $N[\underline{MP}] \subseteq S[\underline{MP}]$, then $M[\overline{MP}] = S[\overline{MP}]$, $M[\underline{MP}] \subseteq S[\underline{MP}]$, and so $M \sqsupseteq S$.

□

4.3.1 An alternative characterization of the Max-Min Weak Models

Similarly to Sections 4.1.1 and 4.2, this section presents an alternative characterization of the Max-Min Weak Model Semantics based on the rewriting of mapping rules into prioritized rules (Brewka *et al.* 2003; Sakama and Inoue 2000). Given an atom $A = i{:}p(x_1, \ldots, x_n)$, we denote as A' the atom $i{:}p'(x_1, \ldots, x_n)$.

Definition 9 (*Rewriting of a P2P System into a Prioritized Logic Program*). Given a P2P system $PS = D \cup LP \cup MP \cup IC$ and the mapping rules $r_a = i_a{:}\, p_a(x_a) \leftarrow \mathcal{B}_a$ and $r_b = i_b{:}\, p_b(x_b) \leftarrow \mathcal{B}_b$, then

- $Rew(r_a)$ denotes the pair $(i_a{:}\, p_a(x_a) \oplus i_a{:}\, p'_a(x_a) \leftarrow \mathcal{B}_a,\ i_a{:}\, p_a(x_a) \succeq i_a{:}\, p'_a(x_a))$,
- $Rew(r_b)$ denotes the pair $(i_b{:}\, p_b(x_b) \oplus i_b{:}\, p'_b(x_b) \leftarrow \mathcal{B}_b,\ i_b{:}\, p'_b(x_b) \succeq i_b{:}\, p_b(x_b))$,
- $Rew(\overline{MP})$ denotes the pair $(\{Rew(r)[1] |\ r \in \overline{MP}\}, \{Rew(r)[2] |\ r \in \overline{MP}\})$,
- $Rew(\underline{MP})$ denotes the pair $(\{Rew(r)[1] |\ r \in \underline{MP}\}, \{Rew(r)[2] |\ r \in \underline{MP}\})$, and
- $Rew(PS)$ denotes the PLP $(D \cup LP \cup Rew(\overline{MP})[1] \cup Rew(\underline{MP})[1] \cup IC,$ $Rew(\overline{MP})[2], Rew(\underline{MP})[2])$.

In the above definition, the atom $i_a{:}\, p_a(x_a)$ (resp. $i_b{:}\, p_b(x_b)$) means that the fact $1_a{:}p_a(x_a)$ is imported in peer P_{i_a} (resp. $1_b{:}p_b(x_b)$ is imported in peer P_{i_b}).

Intuitively, the rewriting of the maximal (resp. minimal) mapping rule states that if $\mathcal{B}_a$ (resp. $\mathcal{B}_b$) is *true* in the *source peer* then two alternative actions can be performed in the *target peer*: $i_a{:}p_a(x_a)$ (resp. $i_b{:}p_b(x_b)$) can be either imported or not imported; but the presence of the priority statement $i_a{:}p_a(x_a) \succeq i_a{:}p'_a(x_a)$ (resp. $i_b{:}p'_b(x_b) \succeq i_b{:}p_b(x_b)$) establishes that the action of *importing* $i_a{:}p_a(x_a)$ is preferable over the action of *not importing* $i_a{:}p_a(x_a)$ (resp. the action of *not importing* $i_b{:}p_b(x_b)$ is preferable over the action of *importing* $i_b{:}p_b(x_b)$).

Observe that $Rew(PS)$ is a PLP with two levels of priorities. The one applied as first models the preference to import as much maximal mapping atoms as possible. The other one, applied over the models selected in the first step, expresses the preference to import as less minimal mapping atoms as possible.

Example 18.
Consider again the system reported in Example 3. The rewriting of the system is
$Rew(PS) =$

$(\{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor)$
$3{:}available(Y) \leftarrow 3{:}supplier(X, Y),$
$3{:}supplier(X, Y) \oplus 3{:}supplier'(X, Y) \leftarrow 1{:}vendor(X, Y),$
$3{:}order(X) \oplus 3{:}order'(X) \leftarrow 2{:}shopping(X),$
$\leftarrow 3{:}order(X), not\ 3{:}available(X)\},$

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

$\{3{:}order(X) \succeq 3{:}order'(X)\},$
$\{3{:}supplier'(X,Y) \succeq 3{:}supplier(X,Y)\}).$

The logic program has the following stable models:

$M_1 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor),$
$3{:}supplier'(dan, laptop), 3{:}supplier'(bob, laptop),$
$3{:}order'(laptop), 3{:}order'(monitor)\},$

$M_2 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor),$
$3{:}supplier(dan, laptop), 3{:}supplier'(bob, laptop),$
$3{:}order'(laptop), 3{:}order'(monitor), 3{:}available(laptop)\},$

$M_3 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor),$
$3{:}supplier'(dan, laptop), 3{:}supplier(bob, laptop),$
$3{:}order'(laptop), 3{:}order'(monitor), 3{:}available(laptop)\},$

$M_4 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor),$
$3{:}supplier(dan, laptop), 3{:}supplier(bob, laptop),$
$3{:}order'(laptop), 3{:}order'(monitor), 3{:}available(laptop)\},$

$M_5 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor),$
$3{:}supplier(dan, laptop), 3{:}supplier'(bob, laptop),$
$3{:}order(laptop), 3{:}order'(monitor), 3{:}available(laptop)\},$

$M_6 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor),$
$3{:}supplier'(dan, laptop), 3{:}supplier(bob, laptop),$
$3{:}order(laptop), 3{:}order'(monitor), 3{:}available(laptop)\},$

$M_7 = \{1{:}vendor(dan, laptop), 1{:}vendor(bob, laptop),$
$2{:}shopping(laptop), 2{:}shopping(monitor),$
$3{:}supplier(dan, laptop), 3{:}supplier(bob, laptop),$
$3{:}order(laptop), 3{:}order'(monitor), 3{:}available(laptop)\}.$

The preferred stable models are M_5 and M_6.

Given a P2P system PS and a preferred stable model M for $Rew(PS)$, we denote with $St(M)$ the subset of non-primed atoms of M and we say that $St(M)$ is a preferred stable model of PS. We denote the set of preferred stable models of PS as $PSM(PS)$.

The following theorem shows the equivalence of preferred stable models and max-min weak models.

Theorem 4 (*Equivalence Between Preferred Stable Models and Max-Min Weak Models*). For every P2P system PS, $PSM(PS) = MaxMinWM(PS)$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Proof

1. $(PSM(PS) \subseteq MaxMinWM(PS))$
   Let $M \in PSM(PS)$ and $N = St(M)$. First we prove that N is a weak model. Let us consider a ground mapping rule m and its rewriting $Rew(m)$. The rule $Rew(m)[1] = A \oplus A' \leftarrow \mathcal{B}$ is equivalent to the rules $r = A \leftarrow \mathcal{B} \wedge not\ A'$, $r' = A' \leftarrow \mathcal{B} \wedge not\ A$ and the constraint $\leftarrow A, A'$. There are three cases:

   - $A, A' \notin M$. In this case, $M \not\models \mathcal{B}$. Then the bodies of r and r' are false and so $r, r' \notin (Rew(PS)[1])^M$.
   - $A \in M$ and $A' \notin M$. In this case the body of r' is false and $r' \notin (Rew(PS)[1])^M$. Moreover, $A \leftarrow \mathcal{B} \in (Rew(PS)[1])^M$.
   - $A' \in M$ and $A \notin M$. In this case the body of r is false and $r \notin (Rew(PS)[1])^M$. Moreover, $A' \leftarrow \mathcal{B} \in (Rew(PS)[1])^M$.

   Then, by construction, we have that $(Rew(PS)[1])^M = St(PS^N) \cup \{A' \leftarrow \mathcal{B} \mid A' \in M \wedge A \leftharpoondown \mathcal{B} \in ground(PS)\}$.
   We have that

   - the minimal model of $(Rew(PS)[1])^M$ is M, as M is a stable model of $(Rew(PS)[1])^M$;
   - $M = N \cup \{A' \mid A' \in M\}$;
   - non-primed atoms A can be only inferred by rules in $St(PS^N)$; and
   - no primed atom A' occurs in the body of any rule of $St(PS^N)$.

   Therefore, the minimal model of $St(PS^N)$ is N and N is a weak model of PS.

   Now we prove that N is a preferred weak model of PS. Let us assume by contradiction that there is a weak model L such that $L[\overline{MP}] \supset N[\overline{MP}]$ or $L[\overline{MP}] = N[\overline{MP}] \wedge L[\underline{MP}] \subset N[\underline{MP}]$. Let us consider these cases:

   - $(L[\overline{MP}] \supset N[\overline{MP}])$. In this case, the ground maximal mapping rules that will be deleted from $ground(PS)$ to derive PS^L are a subset of those that will be deleted to derive PS^N.
     Let us consider the set $K = L \cup \{A' \mid A \notin L \wedge (A \leftharpoondown \mathcal{B} \in ground(PS) \vee A \rightharpoonup \mathcal{B} \in ground(PS)) \wedge L \models \mathcal{B}\}$. By construction, it is the minimal model of $Rew(PS)[1]^K$. Then K is a stable model of $Rew(PS)[1]$. Observe that there must exist two atoms $A \in K$ and $A' \in M$ and, by construction, there cannot exist two atoms $B \in M$ and $B' \in K$. Moreover, $ground(Rew(PS)[2])$ contains the preference $A \succ A'$. Therefore, $K \sqsupset M$ and M is not a preferred stable model of $Rew(PS)$. This is a contradiction.
   - $(L[\overline{MP}] = N[\overline{MP}] \wedge L[\underline{MP}] \subset N[\underline{MP}])$. In this case, the maximal mapping rules that will be deleted from $ground(PS)$ to derive PS^L coincide with those that will be deleted to derive PS^N. Moreover, the ground minimal mapping rules that will be deleted from $ground(PS)$ to derive PS^L are a superset of those that will be deleted to to derive PS^N.
     Let us consider the set $K = L \cup \{A' \mid A \notin L \wedge (A \leftharpoondown \mathcal{B} \in ground(PS) \vee A \rightharpoonup \mathcal{B} \in ground(PS)) \wedge L \models \mathcal{B}\}$. By construction, it is the minimal model of $Rew(PS)[1]^K$. Then K is a stable model of $Rew(PS)[1]$. Observe that there must exist two atoms $A' \in K$ and $A \in M$ and, by construction, there cannot exist two atoms $B' \in M$ and $B \in K$. Moreover, $ground(Rew(PS)[2])$ contains the preference $A' \succ A$. Therefore, $K \sqsupset M$ and M is not a preferred stable model of $Rew(PS)$. This is a contradiction.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

2. $(PSM(PS) \supseteq MaxMinWM(PS))$
Let $N \in MaxMinWM(PS)$ and $M = N \cup \{A' \mid A \notin N \wedge (A \leftarrow\!\!\sim \mathcal{B} \in ground(PS) \vee A \leftarrow\!\!\sim \mathcal{B} \in ground(PS)) \wedge N \models \mathcal{B}\}$. First we prove that M is a stable model of $Rew(PS)[1]$; that is, it is the minimal model of $(Rew(PS)[1])^M$.
By construction, 1 $(Rew(PS)[1])^M = St(PS^N) \cup \{A' \leftarrow \mathcal{B} \mid A \notin N \wedge (A \leftarrow\!\!\sim \mathcal{B} \in ground\,(PS) \vee A \leftarrow\!\!\sim \mathcal{B} \in ground(PS)) \wedge N \models \mathcal{B}\}$. We have that

   - the minimal model of $St(PS)^N$ is N, as N is a weak model of PS;
   - the minimal model of $\{A' \leftarrow \mathcal{B} \mid A \notin N \wedge (A \leftarrow\!\!\sim \mathcal{B} \in ground(PS) \vee A \leftarrow\!\!\sim \mathcal{B} \in ground(PS)) \wedge N \models \mathcal{B}\}$ is $M \setminus N$;
   - non-primed atoms A can be only inferred by rules in $St(PS^N)$; and
   - no primed atom A' occurs in the body of any rule of $St(PS^N)$.

   Therefore, the minimal model of $(Rew(PS)[1])^M$ is M and M is a stable model of $Rew(PS)[1]$.

   Now we prove that M is a preferred stable model for $Rew(PS)$. Let us assume by contradiction that there is a stable model L for $Rew(PS)$ s.t. $L \supset M$. From point 1 and preferences in $Rew(PS)[2]$, we have that $St(L)$ is a weak model for PS and $St(L)[\overline{MP}] \supset St(M)[\overline{MP}]$ or $St(L)[\overline{MP}] = St(M)[\overline{MP}] \wedge St(L)[\underline{MP}] \subset St(M)[\underline{MP}]$; that is, $St(M)$ is not a max-min weak model for PS. This is a contradiction.

□

This characterization of the Max-Min Weak Model Semantics provides a computational mechanism allowing to derive the max-min weak models of a P2P system.

5 Query answers and complexity

We consider now the computational complexity of calculating max-min weak models and answers to queries (Papadimitriou 1994). As a P2P system may admit more than one max-min weak model, the answer to a query is given by considering *brave* or *cautious* reasoning (also known as *possible* and *certain* semantics). Issues related to the distributed computation will be discussed in Section 6.3.

Definition 10.
Given a P2P system PS and a ground peer atom A, then A is *true* under

- brave reasoning if $A \in \bigcup_{M \in MaxMinWM(PS)} M$,
- cautious reasoning if $A \in \bigcap_{M \in MaxMinWM(PS)} M$.

We assume here a simplified framework not considering the distributed complexity as we suppose that the complexity of communications depends on the number of computed atoms which are the only elements exported by peers.

Theorem 5.
Let PS be a P2P system, then

1. deciding whether an interpretation M is a max-min weak model of PS is *coNP* complete;
2. deciding whether a max-min weak model for PS exists is in Σ_2^p;
3. deciding whether an atom A is *true* in some max-min weak model of PS is Σ_2^p complete;
4. deciding whether an atom A is *true* in every max-min weak model of PS is Π_2^p complete.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Proof

1. **(Membership)** We prove that the complementary problem, that is, the problem of checking whether M is not a max-min weak model, is in NP. We can guess an interpretation N and verify in polynomial time that (i) N is a weak model, that is, $\{N\} = MM(St(PS^N))$, and (ii) either M is not a weak model, that is, $\{M\} \neq MM(St(PS^M))$, or $N \sqsupset M$, that is, $N[\overline{MP}] \supset M[\overline{MP}]$ or $N[\overline{MP}] = M[\overline{MP}] \wedge N[\underline{MP}] \subset M[\underline{MP}]$. Therefore, the original problem is in $coNP$.
(Hardness) We will reduce the *SAT problem* to the problem of *checking whether a weak model is not max-min*. Let X be a set of variables and F a CNF (Conjunctive Normal Form) formula over X. Then the problem that will be reduced is checking whether the QBF (Quantified Boolean formula) $(\exists X)\ F$ is $true$. We define a P2P system PS with two peers: P_1 and P_2. Peer P_1 contains the atoms:

$$1{:}variable(x), \text{ for each } x \in X$$
$$1{:}truthValue(true)$$
$$1{:}truthValue(false)$$

The relation $1{:}variable$ stores the variables in X and the relation $1{:}truthValue$ stores the truth values $true$ and $false$.
Peer P_2 contains the atoms:

$2{:}variable(x)$, for each $x \in X$,
$2{:}positive(x, c)$, for each $x \in X$ and clause c in F s.t. x occurs non-negated in c,
$2{:}negated(x, c)$, for each $x \in X$ and clause c in F s.t. x occurs negated in c,

the mapping rule:

$$2{:}assign(X, V) \leftarrow 1{:}variable(X), 1{:}truthValue(V),$$

stating that the truth value V could be assigned to the variable X,
the standard rules:

$$2{:}clause(C) \leftarrow 2{:}positive(X, C),$$
$$2{:}clause(C) \leftarrow 2{:}negated(X, C),$$
$$2{:}holds(C) \leftarrow 2{:}positive(X, C), 2{:}assign(X, true),$$
$$2{:}holds(C) \leftarrow 2{:}negated(X, C), 2{:}assign(X, false),$$
$$2{:}assignment \leftarrow 2{:}assign(X, V),$$

defining a *clause* from the occurrences of its positive and negated variables (first and second rule), whether a clause holds with a given assignment of values (third and fourth rule) and whether an assignment of values actually exists (fifth rule), and the integrity constraints:

$$\leftarrow 2{:}assign(X, true),\ 2{:}assign(X, false),$$
$$\leftarrow 2{:}clause(C),\ not\ 2{:}holds(C),\ 2{:}assignment,$$
$$\leftarrow 2{:}variable(X),\ not\ 2{:}assign(X, true),\ not\ 2{:}assign(X, false),\ 2{:}assignment,$$

stating that two different truth values cannot be assigned to the same variable (first constraint), that if there is an assignment then there cannot be an unsatisfied clause (second constraint) and cannot be an unevaluated variable (third constraint). Let D the set of atoms in PS, MP the set of mapping rules in PS, LP the set of standard rules in PS, and IC the set of integrity constraints in PS. Let M be the minimal model of $D \cup LP \cup IC$, that is, the model containing no mapping atom. As PS is locally consistent, M is a weak model of PS. Observe that the integrity constraints in PS are

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

satisfied when no mapping atom is imported in P_2 that is if no assignment of values is performed for the variables in X. If F is not satisfiable, then there is no way to import mapping atoms in P_2 preserving consistency because the second constraint will be violated. In this case, M is a max-min weak model. If F is satisfiable, there is a weak model N whose set of mapping atoms corresponds to an assignment of values to the variables in X that satisfies F. Clearly, as $MP[N] \supset MP[M]$, M is not a max-min weak model. Moreover, if M is not a max-min weak model, there must be another weak model N whose set of mapping atoms corresponds to an assignment of values to the variables in X that satisfies F. In other words, F is satisfiable if and only if M is not a max-min weak model.

2. Let us guess an interpretation M. By (1), deciding whether M is a max-min weak model can be decided by a call to a $coNP$ oracle.
3. From Theorem 4, an atom A is $true$ in some max-min weak model of PS if and only if it is true in some preferred stable model of PS. The complexity of this problem has been presented in Sakama and Inoue (2000). For disjunction-free ($\vee - free$) PLPs, deciding whether an atom is $true$ in some preferred stable model is Σ_2^p complete.
4. From Theorem 4, an atom A is $true$ in every max-min weak model of PS if and only if it is true in every preferred stable model of PS. The complexity of this problem has been presented in (Sakama and Inoue 2000). For disjunction-free ($\vee - free$) PLPs, deciding whether an atom is $true$ in every preferred stable model is Π_2^p complete.

□

6 Discussion

This section introduces some useful discussions on some features of the proposed semantics.

6.1 Dealing with locally inconsistent P2P systems

The framework presented so far does not guarantee that a locally inconsistent P2P system (i.e. containing at least a locally inconsistent peer) has a weak model. Indeed, there could be locally inconsistent peers that cannot reach a consistent state by importing sets of atoms from other peers. This happens because the only mechanism modeled by our original framework is an "enriching" mechanism that does not allow deletions of atoms from local databases.

Therefore, if a peer is locally inconsistent and there is no way to import mapping atoms able to restore its consistency, the peer remains inconsistent because no atom can be deleted.

In this section, we present an extension of our framework that *simulates* deletions of atoms by using maximal mapping rules.

Informally, the idea is to create for each peer an *auxiliary peer*, move the database from the original peer to the auxiliary one and equip the original peer with a set of maximal mapping rules allowing to import, from the auxiliary peer, maximal sets of atoms not violating its integrity constraints.

Definition 11.
Let $PS = \{P_1, \ldots, P_n\}$ and $P_i = \langle D_i, LP_i, MP_i, IC_i \rangle$, with $i \in [1..n]$, a peer in PS.
Then, $Split(P_i)$ is the set containing the following peers:

- $P'_{(i+n)} = \langle \{(i+n) : p(X) \mid i : p(X) \in D_i\}, \emptyset, \emptyset, \emptyset \rangle$,
- $P'_i = \langle \emptyset, LP_i, MP_i \cup \widehat{MP_i}, IC_i \rangle$, where $\widehat{MP_i} = \{i : p(X) \leftarrow (i+n) : p(X) \mid i : p(X) \in D_i\}$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

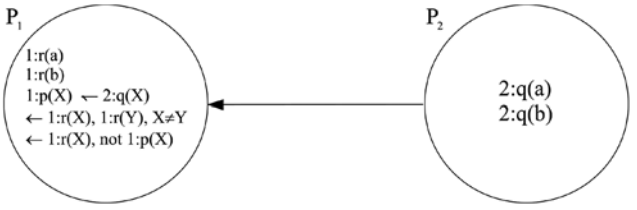


Fig. 6. The system PS.

Moreover,

$$Split(PS) = \bigcup_{P_i \in PS} Split(P_i).$$

In the previous definition, peer P'_i is derived from peer P_i by deleting its local database D_i and inserting a set of maximal mapping rules allowing to import facts into the old base relations (which now are mapping relations) from the auxiliary peer $P'_{(i+n)}$. Given a P2P system PS, we define $\widehat{MP} = \bigcup_{P_i \in PS} \widehat{MP_i}$.

We now present a generalization of our semantics allowing to deal with locally inconsistent P2P systems.

Definition 12.
Let $PS = \{P_1, \ldots, P_n\}$. The *generalized weak models* of PS, denoted as $GWM(PS)$, are obtained from the weak models of $Split(PS)$ by removing all the atoms $i{:}A$ with $i > n$.

Definition 13 (*Generalized Max-Min Weak Model*). Given a P2P system PS and two generalized weak models M and N of PS, we say that M is *G-preferable* to N, and we write $M \sqsupseteq_G N$, if

- $M[\widehat{MP}] \supset N[\widehat{MP}]$ or
- $M[\widehat{MP}] = N[\widehat{MP}]$ and $M[\overline{MP}] \supset N[\overline{MP}]$ or
- $M[\widehat{MP}] = N[\widehat{MP}]$ and $M[\overline{MP}] = N[\overline{MP}]$ and $M[\underline{MP}] \subseteq N[\underline{MP}]$.

Moreover, if $M \sqsupseteq_G N$ and $N \not\sqsupseteq_G M$, we write $M \sqsupset_G N$. A weak model M is said to be a *generalized Max-Min weak model* if there is no weak model N such that $N \sqsubset_G M$. The set of generalized Max-Min weak models will be denoted by $GMinMaxWM(PS)$.

Example 19.
Consider the P2P system PS depicted in Figure 6. P_2 contains the facts $2{:}q(a)$ and $2{:}q(b)$, whereas P_1 contains the minimal mapping rule $1{:}p(X) \leftarrow 2{:}q(X)$; the constraint $\leftarrow 1{:}r(X), 1{:}r(Y), X \neq Y$ stating that the base relation $1{:}r$ can contain at most one tuple, the constraint $\leftarrow 1{:}r(X),\ not\ \ 1{:}p(X)$ stating that if P_1 contains the fact $1{:}r(X)$ then the fact $1{:}p(X)$ has to be derived; and the facts $1{:}r(a)$ and $1{:}r(b)$.

This P2P system is inconsistent because the local database of peer P_1 violates the constraint $\leftarrow 1{:}r(X), 1{:}r(Y), X \neq Y$.

$Split(PS)$ is depicted in Figure 7.
The generalized Max-Min weak models of PS are the following:

- $M_1 = \{2{:}q(a), 2{:}q(b), 1{:}r(a), 1{:}p(a)\}$,
- $M_2 = \{2{:}q(a), 2{:}q(b), 1{:}r(b), 1{:}p(b)\}$.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

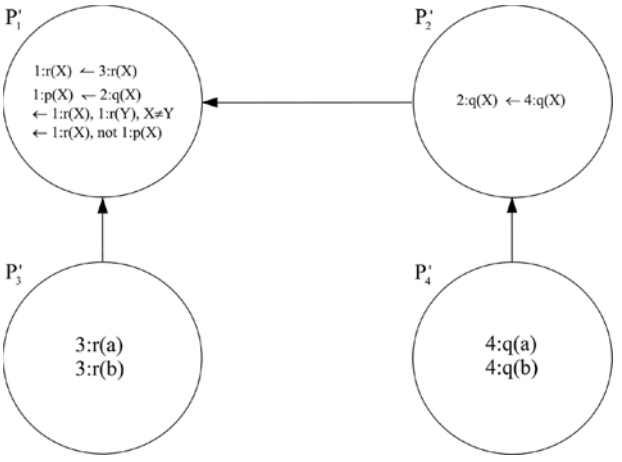


Fig. 7. The system $Split(PS)$.

Observe that if LP contains negation, this technique does not guarantee that consistency can be achieved.

6.2 Computing the Max Weak Model Semantics

This section recalls an alternative characterization of the Max Weak Model Semantics that allows to model a maximal P2P system $PS = \langle D, LP, MP, IC \rangle$, where LP only contains *positive* peer standard rules, with a single disjunctive logic program $Rew_t(PS)$ embedding the set of priorities presented in Section 4.2 (Caroprese and Zumpano 2007).

In Caroprese and Zumpano (2007), it has been shown that the stable models of $Rew_t(PS)$ corres-pond to the maximal weak models of PS.

Using this rewriting, the computation of the preferred weak models of a P2P system PS can be performed in a *centralized way* by using an inference engine like DLV (Leone *et al.* 2005) able to process $Rew_t(PS)$ and compute its stable models.

Although this approach is still not pragmatic, because the computation is centralized and its complexity is prohibitive for real cases, the program $Rew_t(PS)$ can be used as a starting point for a *distributed technique*, as will be pointed out in Section 6.3.

The formal details of this approach are out of the scope of this section and can be found in (Caroprese and Zumpano 2007). Here we want to show how to make our approach pragmatic by implementing a derived version of Max Weak Model Semantics into a real P2P system.

Let's firstly introduce some concepts and definitions. Given a peer atom $A = i{:}p(x)$, A^t denotes the atom $i{:}p^t(x)$ and A^v denotes the atom $i{:}p^v(x)$. A^t is called *testing atom*, whereas A^v is called *violating atom*.

A testing atom A^t corresponds to the mapping atom A that could be derived in the target peer. While A is derived only if its existence does not cause any inconsistency, A^t is always derived in order to test whether A can be inferred safely. If the presence of A violates at least one integrity constraint, the corresponding violating atom A^v is derived. In this case, the atom A^v blocks the derivation of A and the inconsistency that A would cause is prevented.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Definition 14.
Given a conjunction

$$\mathcal{B} = A_1, \ldots, A_h, not\ A_{h+1}, \ldots, not\ A_n, B_1, \ldots, B_k, not\ B_{k+1}, \ldots, not\ B_m, \varphi, \quad (1)$$

where A_i ($i \in [1..\ n]$) is a mapping atom or a derived atom, B_i ($i \in [1..\ m]$) is a base atom and φ is a conjunction of built in atoms, we define

$$\mathcal{B}^t = A_1^t, \ldots, A_h^t, not\ A_{h+1}^t, \ldots, not\ A_n^t, B_1, \ldots, B_k, not\ B_{k+1}, \ldots, not\ B_m, \varphi. \quad (2)$$

From the previous definition, it follows that given a negation free conjunction of the form

$$\mathcal{B} = A_1, \ldots, A_h, B_1, \ldots, B_k, \ldots, \varphi \quad (3)$$

then

$$\mathcal{B}^t = A_1^t, \ldots, A_h^t, B_1, \ldots, B_k, \varphi. \quad (4)$$

Definition 15 (*Rewriting of an integrity constraint*). Given an integrity constraint[4] $c = \leftarrow \mathcal{B}$, its rewriting is defined as $Rew_t(c) = \{A_1^v \vee \cdots \vee A_h^v \leftarrow \mathcal{B}^t\}$.

If the body $\mathcal{B}^t$ (that is of the form (2)), in the previous definition is *true*, then it can be deduced that at least one of the violating atoms $A_1^v, \ldots, A_h^v$ is *true*. This states that in order to avoid inconsistencies, at least one of the atoms $A_1, \ldots, A_h$ cannot be inferred.

Definition 16 (*Rewriting of a standard rule*). Given a standard rule[4] $s = H \leftarrow \mathcal{B}$, its rewriting is defined as $Rew_t(s) = \{H \leftarrow \mathcal{B};\ H^t \leftarrow \mathcal{B}^t;\ A_1^v \vee \cdots \vee A_h^v \leftarrow \mathcal{B}^t, H^v\ \}$.

In order to find the mapping atoms that, if imported, generate some inconsistencies (i.e. in order to find their corresponding violating atoms), all possible mapping testing atoms are imported and the derived testing atoms are inferred. In the previous definition, if $\mathcal{B}^t$ (that is of the form (4)) is *true* and the violating atom H^v is *true*, then the body of the disjunctive rule is *true* and therefore it can be deduced that at least one of the violating atoms $A_1^v, \ldots, A_h^v$ is *true* (i.e. to avoid such inconsistencies at least one of atoms $A_1, \ldots, A_h$ cannot be inferred).

Definition 17 (*Rewriting of a maximal mapping rule*). Given a mapping rule[5] $m = H \leftharpoonup \mathcal{B}$, its rewriting is defined as $Rew_t(m) = \{H^t \leftarrow \mathcal{B};\ H \leftarrow H^t, not\ H^v\ \}$.

Intuitively, to check whether a mapping atom H generates some inconsistencies, if imported in its target peer, a testing atom H^t is imported in the same peer. Rather than violating some integrity constraint, it (eventually) generates, by rules obtained from the rewriting of standard rules and integrity constraints, the atom H^v. In this case, H cannot be inferred and inconsistencies are prevented.

Definition 18 (*Rewriting of a Maximal P2P system*). Given a Maximal P2P system $PS = D \cup LP \cup MP \cup IC$, then

- $Rew_t(MP) = \bigcup_{m \in \mathcal{MP}} Rew_t(m)$,
- $Rew_t(LP) = \bigcup_{s \in \mathcal{LP}} Rew_t(s)$,

[4] Recall that $\mathcal{B}$ is of the form (1).
[5] Recall that $\mathcal{B}$ is of the form (3).

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

- $Rew_t(IC) = \bigcup_{c \in \mathcal{IC}} Rew_t(c)$,
- $Rew_t(PS) = D \cup Rew_t(LP) \cup Rew_t(MP) \cup Rew_t(IC)$.

Example 20.
Let us consider the maximal P2P system presented in Example 4. From Definition (18), we obtain

$$\begin{aligned} Rew_t(PS) = \{&2{:}q(a);\ 2{:}q(b);\\ &1{:}p^t(X) \leftarrow 2{:}q(X);\\ &1{:}p(X) \leftarrow 1{:}p^t(X), \mathit{not}\ 1{:}p^v(X);\\ &1{:}p^v(X) \vee 1{:}p^v(Y) \leftarrow 1{:}p^t(X), 1{:}p^t(Y), X \neq Y\}. \end{aligned}$$

The stable models of $Rew_t(PS)$ are

$$M_1 = \{2{:}q(a), 2{:}q(b), 1{:}p^t(a), 1{:}p^t(b), 1{:}p^v(a), 1{:}p(b)\},$$
$$M_2 = \{2{:}q(a), 2{:}q(b), 1{:}p^t(a), 1{:}p^t(b), 1{:}p(a), 1{:}p^v(b)\}.$$

Definition 19.
[Total Stable Model] Given a P2P system PS and a stable model M for $Rew_t(PS)$, the interpretation obtained by deleting from M its violating and testing atoms, denoted as $T(M)$, is a *total stable model* of PS. The set of total stable models of PS is denoted as $TSM(PS)$.

Example 21.
For the P2P system PS reported in Example 20, $TSM(PS) = \{\{2{:}q(a), 2{:}q(b), 1{:}p(b)\}, \{2{:}q(a), 2{:}q(b), 1{:}p(a)\}\}$.

In Caroprese and Zumpano (2007), it has been shown that the set of total stable models is equivalent to the set of maximal weak models, that is, $TSM(PS) = MaxWM(PS)$.

Observe that this rewriting technique allows computing the maximal weak models of a P2P system with an arbitrary topology. The topology of the system will be encoded in its rewriting. As an example, if a system PS is *cyclic*, its rewriting $Rew_t(PS)$ could be *recursive*.

6.3 A system prototype

The rewriting presented in the previous section has been used in Caroprese and Zumpano (2007); Caroprese and Zumpano (2017b); Caroprese and Zumpano (2017a) as a starting point to implement a system prototype of a P2P system based on a deterministic version of our maximal Weak Model Semantics.

The first important observation is that a P2P system may admit many maximal weak models whose computational complexity has been shown to be prohibitive.

Therefore, it is needed to look for a more pragmatic solution for assigning a semantics to a P2P system. Starting from this observation, a deterministic model whose computation is guaranteed to be polynomial time has been proposed in Caroprese and Zumpano (2007); Caroprese and Zumpano (2017b); Caroprese and Zumpano (2017a). The new proposed semantics, called *well-founded semantics*, assigns to a P2P system its *Well-Founded Model*, a three-valued partial deterministic model that captures the intuition that if an atom is *true* in a maximal weak model and it is *false* in another one, then it is *undefined* in the well-founded model (Gelder 1989; Lonc and Truszczynski 2000).

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

It has been shown that, given a maximal P2P system PS whose standard rules are positive, the rewriting $Rew_t(PS)$ presented in Section 6.2 is Head Cycle Free (Ben-Eliyahu and Dechter 1992). Therefore, it can be normalized obtaining a normal program that we denote as $Rew_w(PS)$.

The next step is to adopt for $Rew_w(PS)$ a Well-Founded Model Semantics. The program $Rew_w(PS)$ admits a well-founded model W that can be computed in polynomial time.

Example 22.
Consider the P2P system presented in Example 4. The normal version $Rew_w(PS)$ of the rewriting $Rew_t(PS)$ presented in Example 20 is

$$\begin{aligned} Rew_w(PS) = \{&2{:}q(a);\ 2{:}q(b);\\ &1{:}p^t(X) \leftarrow 2{:}q(X);\\ &1{:}p(X) \leftarrow 1{:}p^t(X), \mathit{not}\ 1{:}p^v(X);\\ &1{:}p^v(X) \leftarrow 1{:}p^t(X), 1{:}p^t(Y), X \neq Y\};\\ &1{:}p^v(Y) \leftarrow 1{:}p^t(X), 1{:}p^t(Y), X \neq Y\}. \end{aligned}$$

The well-founded semantics of PS is given by the well-founded model of $Rew_w(PS)$, $W = \langle\{2{:}q(a), 2{:}q(b)\}, \emptyset\rangle$.[6] The facts $2{:}q(a)$ and $2{:}q(b)$ are *true*, while the facts $1{:}p(a)$ and $1{:}p(b)$ are *undefined*.

Although the adoption of a well-founded model for a maximal P2P system represents a step forward in the implementation of a real system prototype – as it can be computed in polynomial time – it is evident that the evaluation of a unique logic program requires a *centralized computation* and this is not realistic: a *distributed computation* is needed.

In Caroprese and Zumpano (2017b); Caroprese and Zumpano (2017a), a technique allowing to compute the well-founded model in a distributed way has been presented.

The basic idea is that each peer computes its own portion of $Rew_w(PS)$, sending the result to the other peers. In more detail, if a peer receives a query, then it recursively queries the peers to which it is connected through mapping rules, before being able to calculate its answer.

Formally, a local query submitted by a user to a peer does not differ from a remote query submitted by another peer. The only substantial difference is in the construction of the answer that, in the case of remote query, must be returned to the requesting peer. Once retrieved the necessary data from neighbor peers, the peer computes its well-founded model and evaluates the query (either local or remote) on that model. If the query is a remote query, the answer is sent to the requesting peer.

Details on the architecture and implementation of this system prototype can be found in Caroprese and Zumpano (2017a). The paper in Caroprese and Zumpano (2017a) also reports an application scenario related to the integration of biomedical data from PubMed (http://www.ncbi.nlm. nih.gov/pubmed/). The experiment has been conducted by considering three peers: *Peer1*, *Peer2*, and *Peer3*. *Peer1* contains information about papers related to the HIV virus, *Peer2* contains information about papers related to the Ebola virus, and *Peer3* integrates data provided by *Peer1* and *Peer2*. The final aim of this experiment is the integration of all the papers related to both HIV and Ebola virus into a unique data source in *Peer3*.

[6] The first component of the pair is the set of *true* facts while the second one is the set of *false* facts.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

7 Related works

7.1 Semantic peer data management systems

The present paper is placed among the works on *semantic peer data management systems*. This research topic formally started with the work in Halevy *et al.* (2003) in which the problem of schema mediation in a P2P system is investigated. A formalism, PPL, for mediating peer schemas, which uses the GAV (Global As View) and LAV (Local As View) (Lenzerini 2002) formalism to specify mappings, is proposed. Mappings relate two conjunctive queries expressed in terms of the schema of disjoint peers. The semantics is assigned using classical FOL and query answering is defined by extending the notion of certain answer. More specifically, certain answers for a peer are those that are true in every global instance that is consistent with local data. This choice implies, as a consequence, the need for the consistency of each peer with respect to the whole P2P system. As for a comparison, in this paper we do not adopt an FOL interpretation for a P2P system and tolerate inconsistencies. A mapping rule (Halevy *et al.* 2003) is a logical implication between two peers. It is often the case that preference is given to external data over internal data, or equivalently by using the concept of *trust* given in Bertossi and Bravo (2017) $(P, less, Q)$ denotes that the peer P trusts itself less than Q. This paper formalizes a different proposal: mapping rules are a means used either to import maximal sets of atoms while preventing inconsistency anomalies or to fix the knowledge by importing the minimal sets of atoms allowing to restore consistency. In both cases, in our basic framework, we implicitly satisfy the preference that a peer trusts more its own data over data provided by other peers.

In Calvanese *et al.* (2004), a sound, complete, and terminating procedure that returns the certain answers to a query submitted to a peer is proposed. The paper presents a semantics for a P2P system, based on epistemic logic. Mapping rules between two peers P_1 and P_2 are of the form $CQ_1 \rightarrow CQ_2$, where CQ_1 and CQ_2 are conjunctive queries over the schema of P_1 and P_2. An advantage of this framework is that certain answers of fixed conjunctive queries posed on a peer can be computed in polynomial time. The proposal does not manage local inconsistency. Each peer has to be consistent with respect to its integrity constraints, otherwise the entire P2P system is considered inconsistent. Moreover, if inconsistencies arise due to mapping rules, the whole P2P system is considered inconsistent. An extension of the epistemic theory that ensures *local inconsistency tolerance* has been presented in Calvanese *et al.* (2008). The paper extends the epistemic theory with an additional operator so as to tolerate local inconsistency. More specifically, it ignores a peer inconsistent with respect to its own local constraints. No consistency restoration process is proposed; on the contrary, in our proposal an inconsistent peer is not cut off of the system, but can be *repaired* by means of mapping rules so as to restore consistency.

In Franconi *et al.* (2003); Franconi *et al.* (2004b); Franconi *et al.* (2004a), a characterization of P2P database systems and a model-theoretic semantics dealing with inconsistent peers is proposed. The basic idea is that if a peer does not have models, all the (ground) queries submitted to the peer are *true* (i.e. are *true* with respect to all models). Thus, if some databases are inconsistent, it does not mean that the entire system is inconsistent. The semantics in Franconi *et al.* (2003) coincides with the epistemic semantics in Calvanese *et al.* (2004); Calvanese *et al.* (2005). The semantics in Franconi *et al.* (2003); Franconi *et al.* (2004b); Franconi *et al.* (2004a) also provides a distributed algorithm to compute queries; the setting assumes the existence of a super peer instructor that updates peers' data. The proposal is not inconsistency tolerant as the arise of an inconsistency causes the entire P2P system to became inconsistent. As for a comparison with the present proposal, the works in

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

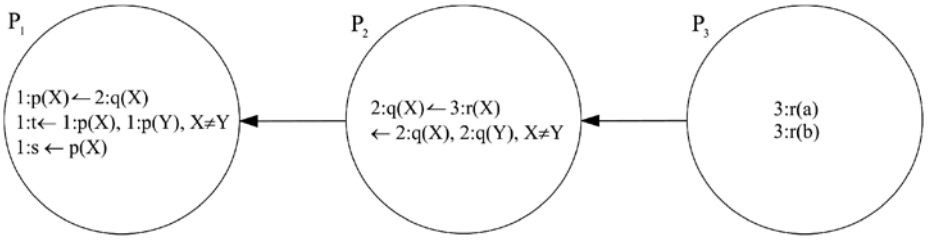


Fig. 8. A P2P system.

Calvanese *et al.* (2004); Calvanese *et al.* (2005) Calvanese *et al.* (2008) and in Franconi *et al.* (2003); Franconi *et al.* (2004b) Franconi *et al.* (2004a) are significantly different.

Consider the P2P system in Figure 8. The epistemic semantics proposed in Calvanese *et al.* (2004); Calvanese *et al.* (2008) states that both atoms 2:$q(a)$ and 2:$q(b)$ are imported in peer P_2 which becomes inconsistent. In this case, the semantics assumes that the whole P2P system is inconsistent and every atom is *true* as it belongs to all minimal models. Consequently, 1:t and 1:s are *true*. The semantics proposed in Franconi *et al.* (2003) assumes that only P_2 is inconsistent as it has no model. Thus, as the atoms 2:$q(a)$ and 2:$q(b)$ are *true* in P_2 (they belong to all models of P_2) and the atoms 1:$p(a)$ and 1:$p(b)$ can be derived in P_1. Finally, 1:t and 1:s are *true*. The Maximal Weak Model Semantics, here proposed, states that with 3:$r(a)$ and 3:$r(b)$ being *true* in P_3, either 2:$q(a)$ or 2:$q(b)$ could be imported in P_2 (but not both, otherwise the integrity constraint is violated) and, consequently, only one tuple is imported in the relation 1:p of the peer P_1. Note that, whatever is the derivation in P_2, 1:s is derived in P_1 while 1:t is not derived. Therefore, the atoms 1:s and 1:t are, respectively, *true* and *false* in P_1. Considering Example 2, as the peer P_2 is inconsistent the semantics presented in Calvanese *et al.* (2004); Calvanese *et al.* (2008) cut it off from the system, whereas our semantics restore consistency in P_2 by importing either 2:$supplier(dan, laptop)$ or 2:$supplier(bob, laptop)$.

In all previous proposals, mapping rules are of "import kind." None of them uses mapping rules to fix the knowledge of a correct, but incomplete database as we do by means of minimal mapping rules. They all adopt the traditional classical idea of importing knowledge and use mapping rules as logical implications. In this paper, we follow a different perspective. Maximal mapping rules are used by a peer to import as much knowledge as possible, while preventing inconsistencies, whereas minimal mapping rules are used by a peer as a means to restore consistency by importing minimal sets of data. Moreover, the combined use of both maximal and minimal mapping rules allows to characterize each peer in the neighborhood as a resource used either to enrich (integrate) or to fix (repair) the knowledge, so as to define a kind of *integrate–repair* strategy. This feature has no counterpart in the abovementioned proposals.

7.2 Preferences in P2P data management systems

In a more general perspective, interesting semantics for data exchange systems, that offer the possibility of explicitly modeling some preference criteria while performing the data integration process, has been proposed in Bertossi and Bravo (2004); Bertossi and Bravo (2007); Bertossi and Bravo (2017); Caroprese and Zumpano (2008); Caroprese and Zumpano (2011).

In Bertossi and Bravo (2004); Bertossi and Bravo (2007); Bertossi and Bravo (2017), a semantics is proposed that allows for cooperation among pairwise peers that are related to

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

each other by means of data exchange constraints (i.e. mapping rules) and trust relationships. The decision by a peer on what other data to consider (besides its local data) does not depend only on its data exchange constraints, but also on the trust relationship that it has with other peers. Given a peer P in a P2P system, a solution for P is a database instance that respects the exchange constraints and trust relationship P has with its "immediate neighbors." Trust relationships are of the form $(P, less, Q)$ stating that P trusts itself less than Q, $(P, more, Q)$ stating that P trusts itself more than Q, and $(P, same, Q)$ stating that P trusts itself the same as Q. These trust relationships are static and are used in the process of collecting data in order to establish preferences in the case of conflicting information.

The introduction of preference criteria among peers is out of the scope of this paper, and in the present proposal, no explicit preference is formally defined. In any case, note that, an implicit preference is embedded into maximal and minimal mapping rules. Specifically, maximal mapping rules state that it is preferable to import as long as no local inconsistencies arise; whereas minimal mapping rules state that it is preferable not to import unless a local inconsistency exists. In addition, a second implicit level of preference exists in our proposal. Each peer trusts more local data over imported data; therefore, our framework always gives more preference to local data over data imported by external peers. This setting can be easily modified in order to cope with the different perspectives in which a generic peer trusts less or the same its own data w.r.t. data provided by external peers.

We have proposed in recent papers extensions of the Max Weak Model Semantics that allow to explicitly express preferences between peers: in Caroprese and Zumpano (2008), a mechanism is defined that allows to set different degrees of reliability for neighbor peers. More specifically, the paper extends the Max Weak Model Semantics with a mechanism that allows to set priorities among mapping rules. While collecting data, it is quite natural for a source peer to associate different degrees of reliability to the portion of data provided by its neighbor peers. Starting from this simple observation, the paper in Caroprese and Zumpano (2008) enhances the Max Weak Model Semantics by using priority levels among mapping rules in order to select the maximal weak models containing a maximum number of mapping atoms according to their importance. Trusted Weak Models can be computed as stable models of a logic program with weak constraints (Buccafurri *et al.* 2000; Calimeri *et al.* 2006). Both in Caroprese and Zumpano (2008) and in Bertossi and Bravo (2004); Bertossi and Bravo (2017), the mechanism is *rigid* in the sense that the preference among conflicting sets of atoms that a peer can import from only depends on the priorities (trust relationship) fixed at design time. To overcome static preferences, in Caroprese and Zumpano (2011) "dynamic" preferences, that allow to select among different scenarios looking at the properties of data provided by the peers, is introduced. The work in Caroprese and Zumpano (2011) allows to model concepts like "*in the case of conflicting information, it is preferable to import data from the neighbor peer that can provide the maximum number of tuples"* or "*in the case of conflicting information, it is preferable to import data from the neighbor peer such that the sum of the values of an attribute is minimum*" without selecting *a priori* preferred peers.

7.3 Relationship to MCS

General P2P data management systems are related to MCS. An MCS consists of a set of contexts and a set of inference rules (known as mapping or bridge rules) that allows the information flow between different contexts. The general nonmonotonic MCS model has been defined in Brewka and Eiter (2007). The paper proposes a general framework for multi-context reasoning that enables to combine arbitrary monotonic and nonmonotonic

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

logics. Information flows among contexts by means of nonmonotonic bridge rules and several different notions of equilibrium for acceptable belief have been investigated.

In Eiter *et al.* (2010); Eiter *et al.* (2014), inconsistencies are analyzed in MCSs, in order to understand where and why they occur and how they can be managed. Each context is assumed to be consistent; therefore, the reason of inconsistencies just relies on the application of mapping rules. The paper introduces two approaches of explaining inconsistencies in MCSs in terms of bridge rules: the first notion characterizes inconsistencies in terms of mapping rules that need to be altered to restore consistency, and the second notion looks for combinations of rules which cause inconsistency. The two notions, following the classical terminology in Reiter (1987), are called respectively *diagnosis* and *explanation*.

Using the concept of diagnosis, it is possible to capture the semantics of maximal P2P systems in terms of MCSs. An MCS M is a collection of contexts $(C_1, \ldots, C_n)$, where $C_i = (L_i, kb_i, br_i)$, L_i is a *logic*, kb_i is a knowledge base and br_i is a set of *bridge rules*, for $i \in [1..n]$. An equilibrium S is a tuple $(S_1, \ldots, S_n)$, where, for $i \in [1..n]$, S_i is the knowledge derived from kb_i and the heads of the bridge rules in br_i whose bodies are satisfied by S. We consider equilibria that are minimal under component-wise set inclusion.[7]

A diagnosis is a pair (D_1, D_2), where D_1 and D_2 are subsets of $\bigcup_{i\in[1..n]} br_i$, such that removing from M the bridge rules in D_1 and adding to M the bridge rules of D_2 in conditional form (obtained from the rules in D_2 by removing the bodies), M is consistent. A maximal P2P system $PS = \{P_1, \ldots, P_n\}$, where $P_i = \langle D_i, LP_i, MP_i, IC_i\rangle$, with $i \in [1..n]$, can be modeled with an MCS system $M = (C_1, \ldots, C_n)$, where, for $i \in [1..n]$, $C_i(L,\ kb_i,\ ground(St(MP_i)))$, L is the ASP logic, and kb_i is obtained by removing the peer identifier from $(D_i \cup LP_i \cup IC_i)$. One can show that the set of maximal weak models of PS correspond to the minimal equilibria of the MCS obtained by removing from M the bridge rules of diagnosis of the form $(D_1, \emptyset)$, where D_1 is minimal. A P2P system cannot be modeled by an MCS (in its basic form) if it contains minimal mapping rules.

Example 23.
Let's consider the P2P system presented in Example 4. It can be modeled by an MCS M having two ASP contexts, $C_1 = (L, kb_1, br_1)$ and $C_2 = (L, kb_2, br_2)$, where L is the ASP logic, $kb_1 = \{\bot \leftarrow p(X),\ p(Y),\ X \neq Y\}$, $br_1 = \{1 : p(a) \leftarrow 2 : q(a), 1 : p(b) \leftarrow 2 : q(b)\}$, $kb_2 = \{q(a),\ q(b)\}$, and $br_2 = \emptyset$. Clearly, M is inconsistent because it does not admit any equilibrium. Indeed, the atoms $p(a)$ and $p(b)$ are derived in C_1, its integrity constraint is violated, and it does not have an acceptable state. M admits two minimal diagnosis of the form $(D_1, \emptyset)$. The first one is $(\{1 : p(a) \leftarrow 2 : q(a)\}, \emptyset)$. If we remove its bridge rule from M, we obtain an MCS having only one minimal equilibrium: $(\{p(b)\}, \{q(a), q(b)\})$. It corresponds to the maximal weak model M_3.

The second one is $(\{1 : p(b) \leftarrow 2 : q(b)\}, \emptyset)$ and removing its bridge rule from M, we obtain an MCS having the only minimal equilibrium $(\{p(a)\}, \{q(a), q(b)\})$. It corresponds to the maximal weak model M_2.

As for additional element of discussion, our proposal falls within the area of P2P system, in which a generic peer is a kind of dynamic context whose presence is not guaranteed in the system; that is, a peer may enter and leave the system, arbitrarily. Therefore, the focus in the P2P context (and also in our paper) is not that of finding the explanations of inconsistencies, but just to cope with them. Moreover, in our work a generic peer is given the possibility to decide how to interact with a neighbor peer: the use of maximal mapping

[7] For precise definitions of these concepts, see (Brewka and Eiter 2007).

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

rules states that it is preferable to import as long as no local inconsistencies arise; whereas the use of minimal mapping rules states that it is preferable not to import unless a local inconsistency exists.

This specific notion has not a counterpart in any of the above works in the field of MCS.

In Bikakis *et al.* (2011), a fully distributed approach for reasoning in Ambient Intelligent Environ-ments, based on the MCS paradigm, has been proposed. The paper refers to the propositional case and inconsistencies are managed by prioritizing mapping rules that cause inconsistency, and specifically the decision which mapping rule to ignore is based, for every context, on the imposed strict total order of all contexts. Specifically, the user is forced to establish, at design time, the preference ordering on all the contexts and, as a consequence, this allows to obtain a unique solution in polynomial time. As for a comparison, our approach models autonomous logic-based entities (peers) that interchange pieces of information using mapping rules. The essential feature of a P2P system is that each peer may leave and join the system arbitrarily. Due to this specific dynamic nature, our proposal avoids forcing any *a priori* preference ordering and as a consequence may admit many preferred weak models, whose computational complexity is in the second level of the polynomial hierarchy. In addition, the work in Bikakis *et al.* (2011) does not deal with the case in which the peer is locally inconsistent, whereas we can cope with this issue. More generally, our proposal supports information flow between different agents through mapping rules, enables reasoning with inconsistent local information (minimal model semantics), and handles agents that provide mutually inconsistent information. On the other hand, it assumes that all peers share a common alphabet of constants and does not include any notion of privacy. In addition, the present proposal in its basic framework does not include any notion of preference between peers, which could be used to resolve potential conflicts caused by mutually inconsistent information sources and does not provide an algorithm for distributed computation. These last two features have been already investigated in other works of the same authors and have been briefly discussed in Section 6. More specifically, different extensions including preference criteria and aggregate functions have been proposed in Caroprese and Zumpano (2008); Caroprese and Zumpano (2011); Caroprese and Zumpano (2012a) and a distributed computation assigning semantics to a P2P system in polynomial time has been presented in Caroprese and Zumpano (2017b); Caroprese and Zumpano (2017a).

Context theories can also be modeled as theories of defeasible logic (Antoniou and Williams 1997; Marek and Truszczynski 1993; Nute 1994) and mappings as *defeasible rules,* and a preference ordering on the system contexts is used to solve conflicts.

8 Concluding remarks and directions for further research

In this paper, we have proposed three different semantics for P2P deductive databases.

In the *Max Weak model Semantics*, a peer imports maximal sets of atoms from its neighborhood to enrich its knowledge while maintaining inconsistency anomalies.

In the *Min Weak model semantics*, the P2P system can be locally inconsistent, and the information provided by the neighbors is used in order to restore consistency, that is, to only integrate with a missing portion of knowledge a correct but incomplete database.

In addition, the present paper unifies the previous two different perspectives captured by the maximal and minimal Weak Model Semantics into the *Max-Min Weak Model Semantics*. This declarative semantics, being more general, allows to characterize each peer in the neighborhood as a resource used either to enrich (integrate) or to fix (repair) the knowledge, so as to define a kind of *integrate–repair* strategy for each peer in the P2P setting.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

The paper also introduces an alternative characterization of the Max-Min Weak Model Semantics (resp. Max Weak Model Semantics and Min Weak Model Semantics) by rewriting a P2P system into an equivalent PLP. Results on the complexity of answering queries are also presented. The paper, by considering analogous results on stable model semantics for PLPs, proves that for disjunction-free ($\vee - free$) prioritized programs deciding whether an interpretation M is a max-min weak model of PS is $coNP$ complete; deciding whether an atom is *true* in some preferred model is Σ_2^p-complete, whereas deciding whether an atom is *true* in every preferred model is Π_2^p-complete (Sakama and Inoue 2000). Moreover, the paper also provides results on the existence of a max-min weak model showing that the problem is in Σ_2^p.

Our work opens several avenues for future research. As a direction for further research, the work could be enriched by the introduction of preference criteria and explicit level of trusts so as to allow, in the presence of multiple alternatives, the selection of data satisfying specific criteria and/or provided by the most reliable sources.

References

Abiteboul, S. and Duschka, O. M. 1998. Complexity of answering queries using materialized views. In *Proc. of the 17th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, Seattle, Washington, 1–3 Jun. 1998, 254–263.

Abiteboul, S., Hull, R. and Vianu, V. 1995. *Foundations of Databases*. Addison-Wesley Publishing Company, Inc., Boston (USA).

Adjiman, P., Chatalic, P., Goasdoué, F., Rousset, M. and Simon, L. 2006. Distributed reasoning in a peer-to-peer setting: Application to the semantic web. *Journal of Artificial Intelligence Research 25*, 269–314.

Antoniou, G. and Williams, M. 1997. *Nonmonotonic Reasoning*. MIT Press Cambridge, Massachusetts (USA).

Arenas, M., Bertossi, L. and Chomicki, J. 1999. Consistent query answers in inconsistent databases. ACM, New York, USA. In *Proc. of the 18th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, May 31 - June 2, Philadelphia, Pennsylvania (USA), 68–79.

Ariew, R. 1976. Ockham's Razor: A Historical and Philosophical Analysis of Ockham's Principle of Parsimony, University of Illinois, Champaign-Urbana.

Ben-Eliyahu, R. and Dechter, R. 1992. Propositional semantics for disjunctive logic programs. In *Proc. of the Joint International Conference and Symposium on Logic Programming (JICSLP 1992)*, Nov. 1992, Washington, DC, 813–827.

Bernstein, P. A., Giunchiglia, F., Kementsietsidis, A., Mylopoulos, J., Serafini, L. and Zaihrayeu, I. 2002. Data management for peer-to-peer computing : A vision. In *Proc. of the 5th International Workshop on the Web and Databases, WebDB 2002*, in conjunction with ACM PODS/SIGMOD 2002. Informal proceedings, Madison, Wisconsin, 6–7 Jun. 2002, 89–94.

Bertossi, L. E. and Bravo, L. 2004. Query answering in peer-to-peer data exchange systems. In *Current Trends in Database Technology – EDBT 2004 Workshops, EDBT 2004 Workshops PhD, DataX, PIM, P2P&DB, and ClustWeb*, Heraklion, Crete, Greece, 14–18 Mar. 2004. Revised Selected Papers, 476–485.

Bertossi, L. E. and Bravo, L. 2007. The semantics of consistency and trust in peer data exchange systems. In *Proc. of the 14th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2007)*, Yerevan, Armenia, 15–19 Oct. 2007, 107–122.

Bertossi, L. E. and Bravo, L. 2017. Consistency and trust in peer data exchange systems. *Theory and Practice of Logic Programming 17,* 2, 148–204.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Bikakis, A., Antoniou, G. and Hassapis, P. 2011. Strategies for contextual reasoning with conflicts in ambient intelligence. *Knowledge and Information Systems 27,* 1, 45–84.

Binas, A. and McIlraith, S. A. 2008. Peer-to-peer query answering with inconsistent knowledge. In *Proc. of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR 2008)*, Sydney, Australia, 16–19 Sep. 2008, 329–339.

Brewka, G. and Eiter, T. 1999. Preferred answer sets for extended logic programs. *Artificial Intelligence 109,* 1–2, 297–356.

Brewka, G. and Eiter, T. 2007. Equilibria in heterogeneous nonmonotonic multi-context systems. In *22nd AAAI Conference on Artificial Intelligence (AAAI07)*, R. C. Holte and A. Howe, Eds. AAAI Press, 1–2, 385–390.

Brewka, G., Niemelä, I. and Truszczynski, M. 2003. Answer set optimization. In *IJCAI-03, Proc. of the 18th International Joint Conference on Artificial Intelligence*, Acapulco, Mexico, 9–15 Aug. 2003, 867–872.

Buccafurri, F., Leone, N. and Rullo, P. 2000. Enhancing Disjunctive Datalog by constraints. *IEEE Transactions on Knowledge and Data Engineering 12,* 5, 845–860.

Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. 2005. Inconsistency tolerance in P2P data integration: An epistemic logic approach. In *Database Programming Languages, 10th International Symposium, DBPL 2005*, Trondheim, Norway, 28–29 August 2005. Revised selected papers, 90–105.

Calì, A., Calvanese, D., De Giacomo, G. and Lenzerini, M. 2004. Data integration under integrity constraints. *Information Systems 29*, 2, 147–163.

Calì, A., Lembo, D. and Rosati, R. 2003. On the decidability and complexity of query answering over inconsistent and incomplete databases. In *Proc. of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, San Diego, CA, USA, 9–12 Jun. 2003, 260–271.

Calimeri, F., Faber, W., Pfeifer, G. and Leone, N. 2006. Pruning operators for disjunctive logic programming systems. *Fundamenta Informaticae 71*, 2–3, 183–214.

Calvanese, D., Damaggio, E., De Giacomo, G., Lenzerini, M. and Rosati, R. 2003. Semantic data integration in P2P systems. In *Databases, Information Systems, and Peer-to-Peer Computing, First International Workshop, DBISP2P*, Berlin, Germany, 7–8 Sep. 2003. Revised Papers, 77–90.

Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. 2008. Inconsistency tolerance in P2P data integration: An epistemic logic approach. *Information Systems 33,* 4–5, 360–384.

Calvanese, D., De Giacomo, G., Lenzerini, M. and Rosati, R. 2004. Logical foundations of peer-to-peer data integration. In *Proc. of the 23rd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, Paris, France, 14–16 Jun. 2004, 241–251.

Caroprese, L., Greco, S. and Zumpano, E. 2006a. A logic programming approach to querying and integrating P2P deductive databases. In *Proc. of the 19th International Florida Artificial Intelligence Research Society Conference*, Melbourne Beach, Florida, USA, 11–13 May 2006, 31–36.

Caroprese, L., Molinaro, C. and Zumpano, E. 2006b. Integrating and querying P2P deductive databases. In *Tenth International Database Engineering and Applications Symposium (IDEAS 2006)*, Delhi, India, 11–14 Dec. 2006, 285–290.

Caroprese, L. and Zumpano, E. 2007. Consistent data integration in P2P deductive databases. In *1st International Conference on Scalable Uncertainty Management (SUM 2007*), Washington, DC, USA, 10–12 Oct. 2007, 230–243.

Caroprese, L. and Zumpano, E. 2008. Modeling cooperation in P2P data management systems. In *Foundations of Intelligent Systems, 17th International Symposium, ISMIS 2008*, Toronto, Canada, 20–23 May 2008, 225–235.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Caroprese, L. and Zumpano, E. 2011. Aggregates and priorities in P2P data management systems. In *15th International Database Engineering and Applications Symposium (IDEAS 2011)*, Lisbon, Portugal, 21–27 Sep. 2011, 1–7.

Caroprese, L. and Zumpano, E. 2012a. Handling preferences in P2P systems. In *Foundations of Information and Knowledge Systems – 7th International Symposium, FoIKS 2012*, Kiel, Germany, 5–9 Mar. 2012, 91–106.

Caroprese, L. and Zumpano, E. 2012b. Restoring consistency in P2P deductive databases. In *Scalable Uncertainty Management – 6th International Conference, SUM 2012*, Marburg, Germany, 17–19 Sep. 2012, 168–179.

Caroprese, L. and Zumpano, E. 2017a. P2P deductive databases: A system prototype. In *Proc. of the 19th International Conference on Information Integration and Web-based Applications & Services, iiWAS 2017*, Salzburg, Austria, 4–6 Dec. 2017, 258–265.

Caroprese, L. and Zumpano, E. 2017b. P2P deductive databases: Well founded semantics and distributed computation. In *New Trends in Databases and Information Systems – ADBIS 2017 Short Papers and Workshops, AMSD, BigNovelTI, DAS, SW4CH, DC*, Nicosia, Cyprus, 24–27 Sep. 2017, 91–99.

Chatalic, P., Nguyen, G. H. and Rousset, M. 2006. Reasoning with inconsistencies in propositional peer-to-peer inference systems. In *ECAI 2006, 17th European Conference on Artificial Intelligence, Including Prestigious Applications of Intelligent Systems (PAIS 2006)*, Riva del Garda, Italy, 29 Aug.–1 Sep. 2006, 352–356.

Delgrande, J. P., Schaub, T. and Tompits, H. 2003. A framework for compiling preferences in logic programs. *Theory and Practice of Logic Programming 3,* 2, 129–187.

Eiter, T., Fink, M., Schüller, P. and Weinzierl, A. 2010. Finding explanations of inconsistency in multi-context systems. In *Proc. of the 12th International Conference on Principles of Knowledge Representation and Reasoning, KR 2010*, Toronto, Ontario, Canada, 9–13 May 2010.

Eiter, T., Fink, M., Schüller, P. and Weinzierl, A. 2014. Finding explanations of inconsistency in multi-context systems. *Artificial Intelligence 216*, 233–274.

Fagin, R., Kolaitis, P. G., Miller, R. J. and Popa, L. 2005. Data exchange: Semantics and query answering. *Theoretical Computer Science 336,* 1, 89–124.

Fagin, R., Kolaitis, P. G. and Popa, L. 2005. Data exchange: Getting to the core. *ACM Transactions on Database Systems 30,* 1, 174–210.

Franconi, E., Kuper, G. M., Lopatenko, A. and Serafini, L. 2003. A robust logical and computational characterisation of peer-to-peer database systems. In *Databases, Information Systems, and Peer-to-Peer Computing, First International Workshop, DBISP2P*, Berlin, Germany, 7–8 Sep. 2003. Revised Papers, 64–76.

Franconi, E., Kuper, G. M., Lopatenko, A. and Zaihrayeu, I. 2004a. A distributed algorithm for robust data sharing and updates in P2P database networks. In *Current Trends in Database Technology – EDBT 2004 Workshops, EDBT 2004 Workshops PhD, DataX, PIM, P2P&DB, and ClustWeb,* Heraklion, Crete, Greece, 14–18 Mar. 2004. Revised Selected Papers, 446–455.

Franconi, E., Kuper, G. M., Lopatenko, A. and Zaihrayeu, I. 2004b. Queries and updates in the coDB peer to peer database system. In *(e)Proc. of the 30th International Conference on Very Large Data Bases*, Toronto, Canada, 31 Aug.–3 Sep. 2004, 1277–1280.

Fuxman, A., Kolaitis, P. G., Miller, R. J. and Tan, W. C. 2006. Peer data exchange. *ACM Transactions on Database Systems 31,* 4, 1454–1498.

Gelder, A. V. 1989. The alternating fixpoint of logic programs with negation. In *Proc. of the 8th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, Philadelphia, PA, USA29–31 Mar 1989, 1–10.

Gelfond, M. and Lifschitz, V. 1988. The stable model semantics for logic programming. In *Proc. of the 5th International Conference and Symposium on Logic Programming,* Seattle, Washington, 15–19 Aug. 1988 (2 Volumes), 1070–1080.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press

Greco, G., Greco, S. and Zumpano, E. 2003. A logical framework for querying and repairing inconsistent databases. *IEEE Transactions on Knowledge and Data Engineering 15,* 6, 1389–1408.

Gribble, S. D., Halevy, A. Y., Ives, Z. G., Rodrig, M. and Suciu, D. 2001. What can database do for peer-to-peer? In *Proc. of the 4th International Workshop on the Web and Databases, WebDB 2001*, Santa Barbara, California, USA, 24–25 May 2001, in conjunction with ACM PODS/SIGMOD 2001. Informal proceedings, 31–36.

Halevy, A. Y., Ives, Z. G., Suciu, D. and Tatarinov, I. 2003. Schema mediation in peer data management systems. In *Proc. of the 19th International Conference on Data Engineering*, Bangalore, India, 5–8 Mar. 2003, 505–516.

Halevy, A. Y., Ives, Z. G., Suciu, D. and Tatarinov, I. 2005. Schema mediation for large-scale semantic data sharing. *VLDB Journal, 14,* 1, 68–83.

Lenzerini, M. 2002. Data integration: A theoretical perspective. In *Proc. of the 21st ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, Madison, Wisconsin, USA, 3–5 Jun. 2002, 233–246.

Leone, N., Greco, G., Ianni, G., Lio, V., Terracina, G., Eiter, T., Faber, W., Fink, M., Gottlob, G., Rosati, R., Lembo, D., Lenzerini, M., Ruzzi, M., Kalka, E., Nowicki, B. and Staniszkis, W. 2005. The INFOMIX system for advanced integration of incomplete and inconsistent data. In *Proc. of the ACM SIGMOD International Conference on Management of Data*, Baltimore, Maryland, USA, 14–16 Jun. 2005, 915–917.

Lonc, Z. and Truszczynski, M. 2000. On the problem of computing the well-founded semantics. In *1st International Conference on Computational Logic – CL 2000,* London, UK, 24–28 Jul. 2000, 673–687.

Madhavan, J. and Halevy, A. Y. 2003. Composing mappings among data sources. In *VLDB 2003, Proc. 29th International Conference on Very Large Data Bases,* Berlin, Germany, 9–12 Sep. 2003, 572–583.

Marek, V. W. and Truszczynski, M. 1993. *Nonmonotonic Logic - Context-dependent Reasoning.* Artificial Intelligence, Springer 1993, ISBN 978-3-540-56448-5, pp. I–XII, 1–417.

Nute, D. 1994. Defeasible logic. In *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning.* Oxford University Press, Inc. New York, NY, USA, 353–395.

Papadimitriou, C. H. 1994. *Computational Complexity.* Addison-Wesley. ISBN 978-0-201-53082-7,
pp. I–XV, 1–523.

Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence 32,* 1, 57–95.

Rousset, M., Adjiman, P., Chatalic, P., Goasdoué, F. and Simon, L. 2006. Somewhere: A scalable peer-to-peer infrastructure for querying distributed ontologies. In *On the Move to Meaningful Internet Systems 2006: CoopIS, DOA, GADA, and ODBASE, OTM Confederated International Conferences, CoopIS, DOA, GADA, and ODBASE 2006*, Montpellier, France, 29 Oct.–3 Nov. 2006, Part I.

Sakama, C. and Inoue, K. 2000. Prioritized logic programming and its application to commonsense reasoning. *Artificial Intelligence 123,* 1–2, 185–222.

Tatarinov, I. and Halevy, A. Y. 2004. Efficient query reformulation in peer-data management systems. In *Proc. of the ACM SIGMOD International Conference on Management of Data,* Paris, France, 13–18 Jun. 2004, 539–550.

https://doi.org/10.1017/S1471068419000073 Published online by Cambridge University Press