On Conservative Enforced Updates *

Michael Dekhtyar[1] , Alexander Dikovsky[2] and Nicolas Spyratos[3]

[1] Dept. of CS, Tver State Univ. 3 Zheljabova str. Tver, Russia, 170000
dekhtyar@tversu.ac.ru
[2] Keldysh Institute for Applied Math. 4 Miusskaya sq. Moscow, Russia, 125047
dikovsky@spp.keldysh.ru
[3] Université de Paris-Sud, LRI, U.R.A. 410 du CNRS, Bât. 490 ·
F-91405 Orsay Cedex, France
Nicolas.Spyratos@lri.fr

Abstract. A new method is proposed of restoring integrity constraints after committing an external update of a data base. This method is conservative in the sense that it performs the minimum of the necessary changes in the initial data base. The method is based on a novel idea of a so called preference strategy, i.e. an "oracle" which resolves globally all the conflicts, taking into account the update to be committed.

1 Introduction

There is an emerging variety of applications for which automatic enforcement of constraints is more appropriate than the traditional policy (according to which updates are rejected in case of inconsistency). Among these applications we can mention the active databases, the temporal databases and the data warehouses. In this work we propose a new approach to the automatic restoration of the integrity constraints of a database after external updates. Our approach applies to the databases with data represented by sets of literals, and integrity constraints represented by general logic programs with clauses of the form $l :- l_1, ..., l_n$, where $l, l_1, ..., l_n$ are literals. The intuitive meaning of a constraint clause is that in the case where all the literals $l_1, ..., l_n$ are in the database, the literal l must also be in the database. An external update is a pair $\Delta = (D^+, D^-)$, where D^+ is the set of literals to be added to the database and D^- is the set of literals to be removed from the database. This fits the frame proposed in [4], and like in [4] we are looking for a solution of the following problem: given an initial data base state I, a set of integrity constraints $\mathcal{P}$, and an external update Δ, a new data base state I_1 should be found, in which the update Δ is committed, and which conforms to the theory of $\mathcal{P}$. We call this problem an "enforced update problem". Various approaches to this problem have been proposed. In the early papers dealing with propositional data bases the so called revision algorithms transforming I into I_1 were described, as well as some natural criteria of minimal change produced by the revision (cf. [8, 3]). More recent papers deal

* This work was sponsored partly by INTAS (Grant 94-2412) and by the Russian Fundamental Studies Foundation (Grants 96-01-00395, 97-01-00973).

with the first order constraints, and propose solutions differing from each other in several aspects. One of the main aspects is the type of a model created by the resolving algorithm. E.g., in [9, 11] and in [12] the resulting model I_1 is stable, whereas in [1] it is well founded. The other aspect is the character of a change produced in I to receive I_1. E.g., in [9] the symmetric difference of I and I_1 is minimal with respect to the other models of constraints. The next aspect is the termination of the resolving algorithm. In [9, 11] the algorithm is partial in general, and converges under some natural restrictions. In [1] the algorithm is total. One more aspect is the type of the considered data bases. E.g., in [9, 11] and in [12] the total databases are considered, whereas in [1] the data bases are incomplete. Another aspect is the determinacy of the solution, and respectively, the confluence or the choice between several possible solutions in the case of nondeterminism. E.g., in [10] the algorithm is sequential and deterministic, in [9, 11, 12] the algorithms are nondeterministic and the resulting solutions are multiple, in [1] the algorithm is deterministic and not sequential.

We resolve the enforced update problem for incomplete data bases and with respect to the general Horn logic programs representing the integrity constraints. Unlike the other approaches our approach is concentrated first of all on the aspect of the minimality of the change produced to the initial data base. The "good model" properties of the resulting data base are regarded as second rate. We provide total nondeterministic algorithms which for any external update consistent with the constraints, deliver multiple solutions such that:

- in each solution the update is committed,
- each solution is a model of the constraints (i.e. the constraints are restored after the update),
- only minimal changes are performed in the initial data base in order to make the resulting data base satisfy the first two claims.

We call such solutions *conservative* and do not give preference of one conservative solution to the others (a real choice may depend on an implementation). The idea of the conservative solution of the enforced update problem sources from [6], where it is resolved in the case of constraints of restricted form: $l :- l_1$, (called there "literal rules"), and for updates restricted to the addition of literals to the database (the addition of an atom a is understood there as the "insertion" of a, and the addition of $\neg a$ is understood as the "deletion" of a).

The solution provided in this paper is novel from the technical point of view. It is based on the so called preference strategy, i.e. an "oracle" which resolves globally all the conflicts, taking into account the update Δ to be committed. This "combinatorial" solution hasn't any certain direction (from goal / from facts), and doesn't follow any semantical scheme (inflationary closure, stable closure, well founded closure, etc.).

The rest of the paper is organized as follows: Section 2 contains basic definitions and notations. Section 3 presents the update problem, basic results and algorithms. Section 4 discusses incrementality of the update algorithms, while Section 5 deals with their computational complexity.