

Some Open Problems Related to the Implied Constraint Problem

Stephen J. Hegner
Department of Computer Science and Electrical Engineering
Votey Building
University of Vermont
Burlington, VT 05405
U. S. A.

Phone: (802)656-3330
CSnet: hegner%uvm.edu@csnet-relay
USEnet: ..!uvm-gen!hegner

Address through August 1990:

Universitetet i Oslo
Matematisk Institutt
Postboks 1053 Blindern
0316 Oslo 3
NORWAY
Phone: +47-2-45-58-88
Net: m.hegner_s@use.uio.uninett

In this note I will identify several open problems of interest to me on the general topic of identifying the constraints on a view of a database schema. Informally, given a base schema and a view mapping, we wish to characterize exactly those states of the view schema which can arise as the image of a state of the base schema. More formally, we may proceed as follows.

A *database schema* is a pair $\mathbf{D} = (\mathbf{R}, \mathbf{C})$, with \mathbf{R} a finite first-order language and \mathbf{C} a set of sentences in the language of \mathbf{R} . In the traditional relational setting, \mathbf{R} is just the set of relation names of the schema and \mathbf{C} is the set of constraints. However, we do not explicitly disallow function symbols in \mathbf{R} ; all of the problems given here make equal sense with function symbols allowed. We let $\text{DB}(\mathbf{D}) = \text{DB}(\mathbf{R})$ denote the set of all structures of the language of \mathbf{R} (the *databases* of \mathbf{D}) and we let $\text{LDB}(\mathbf{D})$ denote the set of all models of \mathbf{C} (the *legal databases* of \mathbf{D}).

If $\mathbf{D}_1 = (\mathbf{R}_1, \mathbf{C}_1)$ and $\mathbf{D}_2 = (\mathbf{R}_2, \mathbf{C}_2)$ are database schemata, a *database mapping* $\gamma : \mathbf{D}_1 \rightarrow \mathbf{D}_2$ is just an interpretation of the language of \mathbf{R}_2 into the language of \mathbf{R}_1 , in the sense given in a mathematical context in [?] and in a database context in [?]. (This just amounts to expressing γ as a query in the relational calculus.) Such an interpretation induces in a natural way a mapping $\gamma^* : \text{DB}(\mathbf{D}_1) \rightarrow \text{DB}(\mathbf{D}_2)$. We call γ *correct* if $\gamma^*(\text{LDB}(\mathbf{D}_1)) \subseteq \text{LDB}(\mathbf{D}_2)$, and we call γ a *view* if $\gamma^*(\text{LDB}(\mathbf{D}_1)) = \text{LDB}(\mathbf{D}_2)$.

We consider problems in which we are given a schema $\mathbf{D}_1 = (\mathbf{R}_1, \mathbf{C}_1)$ (the *base schema*), a second finite first-order language \mathbf{R}_2 , (the *view language*), and an interpretation γ of the language of \mathbf{R}_2 in the language of \mathbf{R}_1 . We are interested in identifying various ways of finding and/or testing a set of sentences Φ such that $\gamma : \mathbf{D}_1 \rightarrow (\mathbf{R}_2, \Phi)$ is correct and/or a view. In other words, we seek to determine the constraints $\mathbf{C}_2 = \Phi$ on the “view” schema $\mathbf{D}_2 = (\mathbf{R}_2, \mathbf{C}_2)$, implied, through γ , by the constraints on the base schema \mathbf{D}_1 . We identify four specific problems. Except for the work of Klug [?] and Jacobs, Aronson, and Klug [?] on the testing problem, I am not aware of any work at all on these problems.

The Testing Problem This is the weakest form of the implied constraint problem which we consider. We are given a single sentence φ in the language of \mathbf{R}_2 , and we ask whether or not $\gamma : \mathbf{D}_1 \rightarrow (\mathbf{R}_2, \{\varphi\})$ is correct. In general, this problem is undecidable. (Just take $\mathbf{R}_1 = \mathbf{R}_2$, $\gamma = \text{identity}$, and we are asking whether $\mathbf{C}_1 \models \varphi$ in the language of \mathbf{R}_1 ; the implication problem for first-order logic). However, the question is semi-decidable; if $\gamma : \mathbf{D}_1 \rightarrow (\mathbf{R}_2, \{\varphi\})$ is correct, we can detect this. The proof, as well as much further discussion on this issue, may be found in [?]. To achieve decidability, we must enforce some restrictions; the open problem is to determine useful ones. In other words, for “interesting” classes of database schemata and mappings, establish necessary and/or sufficient conditions for the testing problem to be decidable.

The Finiteness Problem In this problem, we assume that \mathbf{C}_1 is a finite set, and we ask if there is a *finite* Φ such that $\gamma : \mathbf{D}_1 \rightarrow (\mathbf{R}_2, \Phi)$ is a view. We note that it is quite easy to produce very simple examples such that no such finite Φ exists. For example, take $\mathbf{R}_1 = \{R[ABCD]\}$, $\mathbf{C}_1 = \{B \rightarrow D, C \rightarrow D, DA \rightarrow B\}$, $\mathbf{R}_2 = \{R[ABC]\}$, and $\gamma = \pi_{ABC}$.

The problem of identifying conditions under which there exists such a finite Φ is therefore nontrivial, even in very simple and practical settings.

The Axiomatizability Problem Here we continue to assume that \mathbf{C}_1 is finite, and ask if Φ exists at all as a set of first-order sentences. It is not difficult to construct simple examples in which \mathbf{C}_1 is finite, and yet no first-order Φ exists. For example, take $\mathbf{R}_1 = \{R[AB]\}$, $\mathbf{C}_1 = \{A \rightarrow B, B \rightarrow A\}$, $\mathbf{R}_2 = \{R[A], R[B]\}$, and $\gamma = (\pi_A, \pi_B)$. Then Φ must express the constraint that the cardinality of the instance of $R[A]$ is the same as that of $R[B]$. This is not first-order in general, although it is when restricted to at-most-countable databases. As a general open problem, we seek necessary and or sufficient conditions such that the view is first-order axiomatizable. As a more specific open problem, we ask if there is a nice natural example of a view which is not first-order, even when attention is restricted to databases which are at most countable.

The Algorithmic Problem Here we ask specifically how to compute Φ effectively. More to the point, identify important special cases of base schemata and database mappings such that the axiomatization of the view is recursive. As a related open problem, we may ask also if it is possible to find cases in which the computation is also tractable (polynomial time).

Variations for Logic Databases In addition to the traditional framework of relational databases, this general problem also has application to logic databases. Within the latter framework, we regard \mathbf{C}_1 as the set of sentences comprising the “base” logic database, and we ask how to compute a set of sentences Φ which define the “view” logic database. The additional complication that arises from this point of view is that \mathbf{C}_1 will now vary as the base schema is updated. We must therefore additionally characterize admissible *families* of constraints \mathbf{C}_1 so that Φ is effectively computable for *each* member of this family.

References

- [Ende72] Enderton, H. B. *A Mathematical Introduction to Logic*, Academic Press, 1972.
- [JaAK82] Jacobs, B. E. , A. R. Aronson, and A. C. Klug, “On interpretations of relational languages and solutions to the implied constraint problem,” *ACM TODS*, **7**,2(1982), pp. 291-315.
- [Klug82] A. C. Klug, “Calculating constraints on relational expressions,” *ACM TODS*, **5**, (1980), pp. 260-290.