

This document is an abbreviated and slightly updated version of a proposal which the author submitted to Vetenskapsrådet in April, 2003. The text of the full proposal is available in PDF format at: <http://www.cs.umu.se/~hegner/Publications/Proposals/vr03.pdf>.

1. Specific Goals of the Project

Two essential features of any large, modern database management system are the capability of performing updates and the capability of delivering limited access through views. The integration of these two aspects — the capability of supporting updates which are specified in views — has long been known to be a very difficult problem. The major problem is not one of existence; it is always possible to reflect an update to a view schema back to the main schema. Rather, the key issue is one of suitability; the reflection of the view update back into the base schema must be acceptable according to certain criteria.

For the most part, previous research on this problem has sought to develop strategies whereby a maximal set of view updates is allowable, with little attention paid to so-called *interface criteria* which impact the amount of knowledge which the user must possess in order to effect and understand a view update. Specifically, most of these existing strategies require that the user have knowledge of parts of the schema, and even of the data, which lie outside of the view. The research proposed herein, on the other hand, keeps central the idea that a view is to be an entity unto itself, and should appear to the user to be no different than a full database system. Such a view presentation is particularly important in the case that the user is not a database expert.

The goal of this project is to build upon the previous work of the author on methods of update support for views which emphasize interface criteria, the so-called *closed views*, by studying the complexity issues surrounding the support of updates via such views. More specifically, circumstances under which the management of view updates is computationally feasible, relative to the complexity of updating the base schema, will be investigated.

2. Overview of the Research Area

The topic of database updates via views has been the subject of investigation for at least a quarter century [DB78]. While some of the more recent work has moved from the classical relational model

to other models, such as entity-relationship and object-oriented models, virtually all investigations have been conducted within the context of open views. The distinction between open and closed views was introduced in [Heg90], and further elaborated in [Heg02] and [Heg03]. Roughly speaking, an *open view* admits the widest possible family of update operations, but requires that the user have knowledge of aspects of the overall schema and/or database which lie outside of the view. A *closed view*, on the other hand, is functionally indistinguishable from the full schema of a database system, and may be understood fully without any knowledge of the larger schema.

The seminal work on closed update strategies, *i.e.*, systematic update strategies on closed views, is the *constant-complement* approach, developed more than twenty years ago by Bancilhon and Spyratos [BS81]. In later works [Heg02] [Heg03], the author extended these initial ideas, obtaining key results on a characterization of precisely when a view complement defines a closed update strategy, in terms of *fully commuting pairs*, and also, making use of the natural order structure of database instances, obtained results which identify conditions under which closed update strategies are unique.

Other than the recent paper [Heg04] of the author, the only work which explicitly addresses complexity issues in the context of the constant-complement strategy is [CP84], which identifies a very high worst-case complexity. However, at the time that [CP84] was written, key details regarding the characterization of the closed views in terms of *fully commuting pairs*, as developed in [Heg02] [Heg03] were not known. In light of these more recent developments, the results are not applicable to closed update strategies as employed in this work. In effect, there is no previous literature on update complexity on views, known to the author, which is relevant to the research proposed herein.

3. Description of the Project

Overview of the research program The overall goal of the research program is to investigate the issue of relative update complexity for closed views. The program will be divided into four steps, as follows.

Step 1: Investigation of the problem in the context of the traditional relational model, with simple SPJ-views (Select-Project-Join).

Step 2: Extension of the results of Step 1 to the case in which foreign-key dependencies are also imposed.

Step 3: Investigation of the problem in the general context of order-based views, as developed in the previous work of the author [Heg02], [Heg03].

Step 4: Application of the results of Steps 1 and 2 to a more general data model: the Higher-Order Entity-Relational Model (HERM) [Tha00].

Steps 3 will be undertaken in parallel with Steps 1 and 2, while Step 4 will be taken after the others are well underway. The reason for initiating this investigation with parallel steps is that there does not exist as sufficiently rich body of knowledge about the relative complexity of updates on closed relational views to admit a pure investigation of the general case. On the other hand, it is anticipated that the results of Steps 1 and 2 will feed the development of Step 3.

Preliminary results for Steps 1 and 3 have already been obtained in the paper [Heg04].

References

The papers listed below which are by the author of this proposal are available in PDF format at the web site <http://www.cs.umu.se/~hegner/Publications/>.

- [BS81] Bancilhon, F. and Spyratos, N., “Update semantics of relational views,” *ACM Trans. Database Systems*, **6**(1981), pp. 557–575.
- [CP84] Cosmadakis, S. and Papadimitriou, C., “Updates of relational views,” *J. Assoc. Comp. Mach.*, **31**(1984), pp. 742–760.
- [DB78] Dayal, U. and Bernstein, P. A., “On the updatability of relational views,” in: *Proceedings of the Fourth International Conference on Very Large Data Bases*, pp. 368–474, 1978.
- [Heg90] Hegner, S. J., “Foundations of canonical update support for closed database views,” in: Abiteboul, S. and Kanellakis, P. C., eds., *ICDT’90, Third International Conference on Database Theory, Paris, France, December 1990*, pp. 422–436, Springer-Verlag, 1990.
- [Heg02] Hegner, S. J., “Uniqueness of update strategies for database views,” in: *Foundations of Information and Knowledge Systems: Second International Symposium, FoIKS 2002, Salzau Castle, Germany, February 2002, Proceedings*, pp. 230–249, Springer-Verlag, 2002.
- [Heg03] Hegner, S. J., “An order-based theory of updates for database views,” *Ann. Math. Art. Intell.*, **39**(2003), , in press.
- [Heg04] Hegner, S. J., “The relative complexity of updates for a class of database views,” in: Seipel, D. and Turull Torres, J. M., eds., *Foundations of Information and Knowledge Systems: Third International Symposium, FoIKS 2004, Vienna, Austria, February 2004, Proceedings*, Springer-Verlag, 2004, to appear.
- [Tha00] Thalheim, B., *Entity-Relationship Modeling*, Springer-Verlag, 2000.