1 Introduction to the Problem

Access to databases is typically provided via windows of limited information, called views. The support of updates via such views has long been known to be a difficult problem which demands compromise. In most approaches, a system policy determines which reflected update to the main schema best supports the view update. Typically, such policies are based upon some notion of minimal change; minimal repairs in logic databases [ADB07], information-based approaches [Heg09a], and even the classical constant-complement strategy [BS81] [Heg04] are examples. Unfortunately, in many situations, such approaches and strategies are not applicable. The set of minimal reflections may be very large, with no meaningful way to rank them. More importantly, the user (more precisely, the *rôle*) requesting the view update may not have the access rights necessary to effect the necessary reflection to the main schema. In such situations, a cooperative approach may be applicable; if a rôle does not have sufficient rights to perform a given view update, the cooperation of other rôles, with update rights on other views, may be enlisted to obtain a suitable reflected update to the main schema.

In this research program, a systematic study of view update via cooperation is to be undertaken. A key feature of this approach is that rather than regarding the views as comprising a decomposition of the main schema, the main schema itself is regarded as a composition of *schema components* [Heg08b]. In this way, the views become first-class citizens within the entire framework, and all necessary concepts may be expressed directly with respect to them.

2 Completed Work

A substantial amount of work has already been undertaken in support of this project. The major results include the following.

Database schema components: The basic model of schema components, upon which this work is based, is developed in [Heg08b]. Although earlier models of database components have been developed [ST04] [Tha05] [ST06], they are rooted in the *co-design* approach [Tha03] in which the applications are embodied in the database schema itself. On the other hand, the work proposed here retains the classical notion of *data independence* in which schema design in independent of particular applications. Therefore, the framework of [Heg08b] is appropriate for the proposed research.

Negotiation for cooperative update: Cooperative update requires a model of negotiation amongst the parties. An initial proof-of-concept model was presented in [HS07]. A much more efficient and realistic model was recently presented in [Heg11], and will be used as the basis for the proposed work.

Component-based concurrency: Support for concurrent transactions is central to modern database systems. Therefore, a mechanism to support cooperative updates concurrently is essential to the success of the approach, and the development of such a mechanism forms a central part of the proposed work. In [Heg10b], a model of concurrency which was motivated by and is particularly suited to the needs of cooperative update is presented.

Optimality of local updates: Although the proposed approach is cooperative, the need to optimize the local updates to each component remains. The necessary notions for such local optimization have been developed in the context of views in a sequence of papers [Heg08a], [Heg09a], [Heg08c], and [Heg09b]; these may be adapted readily to the component-based framework.

View internalization: Traditional database views are defined via a quotient construction on the main schema. Consequently, it is difficult to combine and compare them. In [Heg10a] and [Heg10c], it is shown that a class of very useful views defined by *conjunctive queries* may be represented using a subset structure relative to the main view. This view internalization will be used in the proposed work to simplify many of the constructions, particularly those which relate components to one another.

3 The Proposed Research

The proposed research consists of six related tasks.

Concurrent transactions on schema components: A major focus of the proposed work will be the support of concurrent transactions. Traditional models of database concurrency are oriented towards non-cooperative, batch-oriented transactions which may be rolled back or aborted in the case of conflict. Cooperative update requires a different approach in which much more planning and a more sophisticated notion of data object are warranted to avoid such problems, and in which the need to support the much slower responses of human actors is central. The framework in [Heg10b] will be used as the basis for the development of the appropriate model of transaction and concurrency.

Rôle-based access for schema components: Cooperative update requires an access model which recognizes the interaction of the actors. The approach will be based upon and extend *rôle-based access* [SCFY96], [WO07]. The proposed model will include not only traditional access rights to parts of the database, but also obligations to cooperate in certain ways, based upon an authority hierarchy of rôles.

Efficient representation of nondeterministic update specifications: The model of negotiation forwarded in [HS07] and [Heg11] is founded upon the notion of *nondeterministic update requests*; actors may support not just a single update but rather a family of possibilities. The negotiation process then is tantamount to finding those updates upon which all agree. For this approach to be realizable in practice, an efficient representation for a suitable family of such requests is essential, and the development of such forms a major component of the proposed work.

Relationship to workflow and business processes: The common problem setting of the proposed work and workflow [AH02], and in particular business process modelling using languages such as BPMN [BPM09], is unmistakable. Indeed, many of the major example applications are similar. However, while workflows and business-process languages specify the flow of control explicitly, cooperative update identifies the constraints which govern the underlying workflows without imposing any specific flow of control. As part of this research project, a more formal connection between these approaches will be established.

Support for non-monotonic negotiation: The model of negotiation forwarded in [HS07] and [Heg11] is monotonic, in the sense that further negotiation can only reduce the initial set of possibilities; actors cannot persuade others to expand their set of acceptable supporting updates. This choice was made for reasons of obtaining efficient algorithms. However, non-monotonic negotiation is clearly important, and classes of such which may be realized efficiently must be identified.

Application to consistency maintenance of large data repositories: In early work on this project, Peggy Schmidt, a doctoral student at Christian-Albrechts-Universität zu Kiel in Germany, was a collaborator. Subsequently, she moved to an industrial position and was unable to continue. However, now she has resumed work on her dissertation and will collaborate on an important practical application for the framework, a component-based approach to the maintenance of consistency for large data repositories which may contain multiple representations of the same information.

References

[AH02] Wil van der Aalst and Kees van Hee, *Workflow Management: Models, Methods, and Systems*, MIT Press, 2002.

[ADB07] Ofer Arieli, Marc Denecker, and Maurice Bruynooghe, "Distance semantics for database repair," *Ann. Math. Artif. Intell.*, **50**(2007), pp. 389–415.

[BS81] François Bancilhon and Nicolas Spyratos, "Update semantics of relational views," *ACM Trans. Database Systems*, **6**(1981), pp. 557–575.

[BPM09] "Business process modeling notation v2.0," `http://www.omg.org/spec/BPMN/2.0/Beta1/PDF`, 2009.

[Heg04] Stephen J. Hegner, "An order-based theory of updates for closed database views," *Ann. Math. Art. Intell.*, **40**(2004), pp. 63–125.

[Heg08a] Stephen J. Hegner, "Information-optimal reflections of view updates on relational database schemata," in: Sven Hartmann and Gabriele Kern-Isberner, eds., *Foundations of Information and Knowledge Systems: Fifth International Symposium, FoIKS 2008, Pisa, Italy, February 11-15, 2008, Proceedings*, pp. 112–131, Volume 4932 of Lecture Notes in Computer Science, Springer-Verlag, 2008.

[Heg08b] Stephen J. Hegner, "A model of database components and their interconnection based upon communicating views," in: Hannu Jakkola, Yashui Kiyoki, and Takehiro Tokuda, eds., *Information Modelling and Knowledge Systems XIX*, pp. 79–100, Frontiers in Artificial Intelligence and Applications, IOS Press, 2008.

[Heg08c] Stephen J. Hegner, "Semantic bijectivity and the uniqueness of constant-complement updates in the relatiional context," in: Klaus-Dieter Schewe and Bernhard Thalheim, eds., *International Workshop on Semantics in Data and Knowledge Bases, SDKB 2008, Nantes, France, March 29, 2008, Proceedings*, pp. 172–191, Volume 4925 of Lecture Notes in Computer Science, Springer-Verlag, 2008.

[Heg09a] Stephen J. Hegner, "Information-based distance measures and the canonical reflection of view updates," 2009, invited submission to a special issue of *Annals of Mathematics and Artificial Intelligence*, selected papers from FoIKS 2008.

[Heg09b] Stephen J. Hegner, "Optimal reflection of bidirectional view updates using information-based distance measures," in: Leopoldo Bertossi and Henning Christiansen, eds., *Proceedings of the International Workshop on Logic in Databases (LID 2009), Roskilde University, Denmark, October 29, 2009*, pp. 55–70, 2009.

[Heg10a] Stephen J. Hegner, "Internal representation of database views," *J. Universal Comp. Sci.*, **16**(2010), pp. –, in press.

[Heg10b] Stephen J. Hegner, "A model of independence and overlap for transactions on database schemata," in: Barbara Catania and Bernhard Thalheim, eds., *Advances in Databases and Information Systems, 14th East European Conference, ADBIS 2010, Novi Sad, Serbia, September 20 - 24, 2010, Proceedings*, Volume TBA of Lecture Notes in Computer Science, Springer-Verlag, 2010.

[Heg10c] Stephen J. Hegner, "Optimal complements for a class of relational views," in: Klaus-Dieter Schewe and Bernhard Thalheim, eds., *Fourth International Workshop on Semantics in Data and Knowledge Bases, SDKB 2010, Bordeaux, France, July 5, 2010, Proceedings*, Volume TBA of Lecture Notes in Computer Science, Springer-Verlag, 2010.

[Heg11] Stephen J. Hegner, "A simple model of negotiation for cooperative updates on database schema components," in: Yasushi Kiyoki, Takehiro Tokuda, Anneli Heimbürger, Hannu Jaakkola, and Naofumi Yoshida., eds., *Frontiers in Artificial Intelligence and Applications XX11*, 2011, in press.

[HS07] Stephen J. Hegner and Peggy Schmidt, "Update support for database views via cooperation," in: Yannis Ioannis, Boris Novikov, and Boris Rachev, eds., *Advances in Databases and Information Systems, 11th East European Conference, ADBIS 2007, Varna, Bulgaria, September 29 - October 3, 2007, Proceedings*, pp. 98–113, Volume 4690 of Lecture Notes in Computer Science, Springer-Verlag, 2007.

[SCFY96] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman, "Role-based access control models," *IEEE Computer*, **29**(1996), pp. 38–47.

[ST06] Klaus-Dieter Schewe and Bernhard Thalheim, "Component-driven engineering of database applications," in: *APCCM '06: Proceedings of the 3rd Asia-Pacific conference on Conceptual modelling*, pp. 105–114, Australian Computer Society, Inc., Darlinghurst, Australia, 2006.

[ST04] Peggy Schmidt and Bernhard Thalheim, "Component-based modeling of huge databases," in: András Benczúr, János Demetrovics, and Georg Gottlob, eds., *Advances in Databases and Information Systems: 8th East European Conference, ADBIS 2004, Budapest, Hungary, September 22-25, 2004, Proceedings*, pp. 113–128, Volume 3255 of Lecture Notes in Computer Science, Springer-Verlag, 2004.

[Tha03] Bernhard Thalheim, "Informationssystem-Entwicklung: die integrierte Entwicklung der Strukturering, Funktionalität, Verteilung und Interaktivität von großen Informationssystemen," Technical Report 15-2003, Brandenburgische Technische Universität Cottbus, 2003.

[Tha05] Bernhard Thalheim, "Component development and construction for database design," *Data Knowl. Eng.*, **54**(2005), pp. 77–95.

[WO07] He Wang and Sylvia L. Osborn, "Discretionary access control with the administrative role graph model," in: Volkmar Lotz and Bhavani M. Thuraisingham, eds., *SACMAT 2007, 12th ACM Symposium on Access Control Models and Technologies, Sophia Antipolis, France, June 20-22, 2007, Proceedings*, pp. 151–156, ACM, 2007.