

EFIS'99: Engineering Federated Information Systems

The second International Workshop on Engineering Federated Information Systems EFIS'99 took place in the Baltic town of Kühlungsborn from May 5 to May 7, 1999. After the successful first International Workshop on Engineering Federated Database Systems (EFDBS'97) in Barcelona in June 1997 [1], the goal of this second workshop was to extend the scope in order to cover database as well as non-database information sources. This change in the scope of the workshop was reflected in the change from EFDBS to EFIS.

We regard FIS (Federated Information Systems) as information systems that are integrated to some extent. For realizing such an integration we can utilize middleware tools, database management systems, etc. Obviously, we do not regard those tools themselves as being FIS. In contrast to the classical notions of 'federated database systems' and 'multi-database systems', the term 'federated information system' intends to include not only structured information sources but also semi-structured and even unstructured information. These inclusions often have implications on the three dimensions of autonomy, heterogeneity, and distribution which are frequently used for characterizing database systems architectures. A FIS, in general, may allow for a rather high degree of freedom in each of these dimensions and, therefore, require new techniques for coping with the resulting problems. However, it is obvious that only severe restrictions in particular dimensions (e.g., mediator-based FIS with read-only query access) can help in the actual design of practicable federation solutions, which was one of the major issues for discussion throughout the workshop.

Starting from the research issues identified in the EFDBS'97 Workshop [1] EFIS'99 particularly focussed on topics like the following ones. Of course, a major topic was analysis, design, and implementation of FIS. Furthermore, software architectures for FIS and experience with middleware platforms for building FIS were to the fore as well. Technical support, in particular for the integration of legacy systems, was another important aspect. Experiences with and evaluation of techniques, methods, architectures, and tools for building FIS were also expected and were covered by several talks.

The international programme committee selected 14 papers from the submissions for presentation on the workshop: seven as long papers and seven as short papers. These papers are included in the Workshop Proceedings [3]. Beside a discussion session three invited talks rounded off the workshop programme: H. Gröger (IBM Germany) presented the basic concepts of "IBM DB2 DataJoiner: A commercially available federated database system", F. Saltor (UPC Barcelona, Catalonia) investigated "Semantic Issues in FIS", and J. Grimson (Trinity College Dublin, Ireland) reported on experiences in "Engineering Federated Electronic Healthcare Records".

Out of the accepted papers for this workshop the Programme Committee decided to propose two papers for publication in the "Australian Journal on Information Systems". These two papers are "Constructing Multidatabase Collec-

tions Using an Extended ODMG Object Model” by A. Skehill and M. Roantree and “An Integrated Toolkit for Building Database Federations” by U. Hohenstein and A. Ebert. Object models, mainly extensions to ODMG, are often used as canonical data models in research prototypes for federated database systems, and both papers are follow this approach. The main observations can be summarized as follows:

- The need for explicit federation definition and query/manipulation languages is observed. Such specification languages may enable, for instance, the automatic generation of homogenizing adapters (wrappers).
- The introduction of a metamodel in ODMG 2.0 should assist FIS designers who use the ODMG model as a canonical model. However, this is offset by the fact that the ODMG schema access interface has yet to be implemented by any of the ODMG database vendors.
- It appears that ODMG is not well-suited as a canonical model. Sometimes it is hard to reflect the ODL and OML semantics of specific concepts of commercial ODBMSs in ODMG. Even if it is possible to reflect concepts such as clustering or versions in ODMG for particular ODMBS, there is a need for a common simulation of *all* the advanced concepts of *all* the ODBMS, i.e., some kind of abstraction of those concepts. Otherwise it seems to be impossible to use those features for integration.

The two papers have been carefully revised and extended in order to respect the comments and suggestions made by the EFIS'99 programme committee and reviewers as well as by the EFIS'99 participants during the workshop.

A workshop report focussing on the results and observations made during the extensive discussion on the workshop is available as [2].

August 1999

S. Conrad (Magdeburg University, Germany)

W. Hasselbring (Tilburg University, The Netherlands)

G. Saake (Magdeburg University, Germany)

References

- [1] S. Conrad, B. Eaglestone, W. Hasselbring, M. Roantree, F. Saltor, M. Schönhoff, M. Strässler, and M. Vermeer. Research Issues in Federated Database Systems. *ACM SIGMOD Record*, 26(4):54–56, 1997.
- [2] S. Conrad, W. Hasselbring, U. Hohenstein, R.D. Kutsche, M. Roantree, G. Saake, and F. Saltor. Engineering Federated Information Systems — Report of EFIS'99 Workshop. *ACM SIGMOD Record*, 28(3), 1999. *to appear*.
- [3] S. Conrad, W. Hasselbring, and G. Saake, editors. *Engineering Federated Information Systems (Proceedings of the 2nd Workshop EFIS'99, May 5–7, 1999, Kühlungsborn, Germany)*. infix-Verlag, 1999.