

Executive Summary

Dagstuhl Seminar 08281: “Software Engineering for Tailor-made Data Management”

S. Apel, D. Batory, G. Graefe, G. Saake, and O. Spinczyk

Schloss Dagstuhl, July 6th to July 11th, 2008

Tailoring data management components is not only important for the domain of embedded systems. A database management system (DBMS) that provides exactly the functionality that is needed by the platform, the application scenario, and the stakeholder’s requirements, yields several benefits, e.g., a lean and well-structured code base, robustness, and maintainability.

The desire for tailor-made data management solutions is not new: concepts like kernel-systems or component toolkits have been proposed 20 years ago. However, a view on the current practice reveals that nowadays data management solutions are either monolithic DBMS such as ORACLE and DB2 or special-purpose systems developed from scratch for specific platforms and scenarios, e.g., for embedded systems and sensor networks.

While monolithic DBMS architectures hinder a reasonable and effective long-term evolution and maintenance, special-purpose solutions suffer from the conceptual problem to reinvent the wheel for every platform and scenario or to be too general to be efficient. A mere adaptation of present solutions is impossible from the practical point of view, e.g., it becomes too expensive or simply impractical, which is confirmed by the current practice. Especially in the domain of embedded and realtime systems there are extra requirements on resource consumption, footprint, and execution time. That is, contemporary data management solutions have to be tailorable to the specific properties of the target platform and the requirements and demands made by the stakeholders and the application scenario.

The question that arises is what is different now from the attempts made some years ago that makes us believe that something can change in future. This question and possible answers from the field of modern software engineering have been discussed in the Dagstuhl Seminar “Software Engineering for Tailor-made Data Management”, July 6h to July 11th, 2008.

In 12 sessions, 21 talks have presented recent work and new ideas regarding trends in software engineering and their application to the development of tailor-made data management solutions. On the one hand, there were researchers from the database community who stressed the necessity of tailoring data management solutions, e.g., in order to be able to adapt them at runtime in a service-oriented world or to provide different variants for query processing and optimization. On the other hand, several researchers from the software engineering and programming languages communities presented new ideas to build, manage, and evolve complex systems such as database systems. Continuous discussions have connected the different views of both communities. A recurring idea was to think in terms of families of DMS instead of individual systems. Members of a DMS family share common “features” with the goal of reuse. Several researchers presented mechanisms for identifying, implementing, and managing features or reported from success stories in the field of DMS or related domain.

In two special sessions the participants have been split into three discussion groups addressing the specific topics of tailor-made data management solutions:

“Software Engineering or C for Embedded Systems?” The first discussion group addressed the special needs of the embedded systems industry and the potential of various novel software engineering approaches to advance the state of the art in this field. The focus lay on languages and tools that simplify the development and maintenance of software product-lines such as advanced preprocessors, aspect-oriented programming, feature-oriented programming, model-based code generators, and service-oriented architectures. The group identified the benefits but also open issues of each of these approaches and concluded with a step-by-step roadmap for the adoption of the techniques in industrial embedded systems.

“Dynamic Adaptation in DBMS – Is There a Use Case?” This group discussed a specific branch of tailor-made DBMS research, i.e. the *dynamic* adaptation. After finding a common terminology, the group showed that there are indeed use cases for dynamic adaptation, for

instance maintenance issues of long-running DBMS or self-tuning DBMS. Several technical problems of dynamic adaptation were discussed. Even though the dynamic adaptation of running DBMS code was considered useful, the transition of program state is still an open research problem.

“DBMS in 10 Years – How Can Software Engineering Help?” The third group had the most difficult task: A prediction about the state of the art in DBMS in about 10 years. The group discussed current trends in programming models and languages, hardware, and future application scenarios, each from the perspectives of DB users and DB developers. Among many other topics XML, multi-core CPUs, and autonomous DBMS were addressed. It was concluded that most contemporary software engineering methods address application development, but not specifically low-level systems development as needed by DB developers. Therefore, our joint research effort is very important for the future of DBMS.

In a wrap-up session, the participants have reflected the program and discussions of the seminar. The general opinion was that both communities have learned much from each other. Individual joint projects have been planned, a mailing list has been installed¹, and a continuation of a related workshop series has been encouraged.

¹ dagstuhl-setmdm@lists.cs.tu-dortmund.de