

Cellular DBMS: An Attempt Towards Biologically-Inspired Data Management

Syed Saif ur Rahman, Veit Köppen, Gunter Saake
Department of Technical and Business Information Systems
Faculty of Computer Science
Otto-von-Guericke University of Magdeburg, Germany
{srahman, vkooppen, saake}@ovgu.de



Journal of Digital
Information Management

ABSTRACT: Existing database management systems (DBMS) are complex and less predictable (i.e., the consistency of performance with the increase of functionality and the data growth is not certain). Database researchers acknowledge the need for revisiting DBMS architectures to fulfill the needs of new hardware and application trends. We propose a biologically inspired DBMS architecture called "Cellular DBMS". The Cellular DBMS architecture promises development of highly customizable and autonomous DBMS. This paper explains in detail the design principles for Cellular DBMS architecture. It also explains an aspect-oriented programming based model to equip Cellular DBMS architecture with autonomy. Finally, it presents an extension to decomposed storage model (DSM) for use in Cellular DBMS.

Categories and Subject Descriptors

H.2.2 [Database Management]; Physical Design: H.2.4 [Database Management]; Systems

General Terms

Design, Experimentation, Management, Performance, Reliability.

Keywords: Customization, Adaptability, Biological-Inspiration, Data Management, Embedded Databases, DBMS Architectures

Received: Received 5 August 2009; Revised 29 September 2009; Accepted 8 October 2009

1. Introduction and Motivation

In the past, database research got motivations from arrival of new hardware, software, and applications for further progress. These motivations are still there and will persist in the future. Desire for improvements keeps researchers busy for finding the breakthrough in prevailing requirements. Sometimes we get many breakthroughs in a short period and sometimes we wait for decades to get few. In the prevailing era, we have explosion in the data growth and usage scenarios, because of wide spread usage of internet and advent of new applications (e.g., social networking, virtual worlds). Hardware trends are changing and the processing and storage unit cost have reduced. Many assumptions about secondary storage and mainmemory, etc., made in past are no longer valid and many bottlenecks, such as network communication cost have changed. Leading database researchers found a consensus on the need of revisiting database engines, accommodating architectural shifts in computing hardware and platforms, and finding solutions for new usage scenarios [1]. Cellular DBMS¹ is an effort to contribute to database research in the above-mentioned directions.

¹ "Cellular DBMS", <http://www.witi.cs.uni-magdeburg.de/~srahman/CellularDBMS/index.html>

Existing data management solutions are complex. These solutions have evolved over time and now they provide a multitude of functionalities. These functionalities are tightly coupled within their monolithic architecture [2]. Due to complexity, their performance is less predictable, i.e., the consistency of performance with the increase of functionality and the data growth is not certain and it is difficult to assess, how performance will vary for different hardware, workload, and operating systems, etc. Continuous administration and maintenance is needed to keep them performing at an optimal level, which results in high administrative and maintenance cost. Existing database management systems (DBMS) have dozens of tuning knobs. Internal sub-systems are tightly coupled. Effect of tuning a knob on other knobs and their performance is less predictable [2, 3]. Furthermore, existing DBMS architectures and solutions were designed decades ago considering legacy hardware and their bottlenecks. Now many opportunities exist to redesign existing data management architectures for exploiting features of new hardware.

Database researchers have suggested transition of DBMS from monolithic to a diversified architecture with small, simple, and reusable components of limited functionality with clean inter-component interaction [1, 2]. The Cellular DBMS architecture is designed by considering these suggestions. The Cellular DBMS architecture takes inspiration from biological systems. We want to utilize the mechanisms that exist in biological systems for data management. Using these mechanisms, we want to develop highly customizable and autonomous DBMS with more predictable performance. The vision for Cellular DBMS predictability is shown in Figure 1, i.e., a DBMS should be consistently predictable with the data growth and addition of functionalities. To achieve these goals in Cellular DBMS, we envision integration of techniques from different relevant fields, such as software engineering, distributed data management, computer networks, and parallel processing.

This paper is organized as follows: Section 2 introduces the related concepts required for background information and technical discussion. A detailed related work is provided in Section 3. Cellular DBMS architecture and its design principles are explained in Section 4. Section 5 presents the implementation details. Sample implementation scenarios are discussed in Section 6. Section 7 concludes the paper with some directions to future work.

2. Related Concepts

In this section, we will introduce the related concepts of DBMS, Software Engineering, and Autonomy that are important for improving the understandability of the reader for the topic.

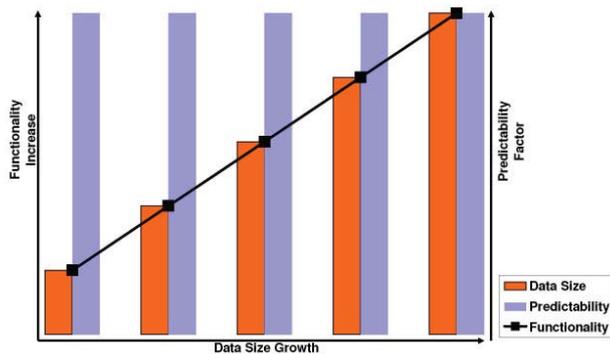


Figure 1. Cellular DBMS goal for predictability.

2.1 DBMS Aspect

In this sub-section, we will briefly introduce the related DBMS aspects for Cellular DBMS, which includes the concepts of storage models and embedded databases.

2.1.1 Storage Models

Storage model selection is an important design decision for DBMS architecture. In this sub-section, we will explain the two most commonly used storage models, i.e., N-Ary Storage Model [4] and Decomposed Storage Model [4] followed by discussion on design decision of selecting decomposed storage model for Cellular DBMS architecture.

N-Ary Storage Model (NSM)

N-Ary Storage Model (NSM) stores data as seen in the relational conceptual schema, i.e., all attributes of a conceptual schema record are stored together [4]. Most of the existing DBMS are NSM based.

Decomposed Storage Model (DSM)

Decomposed Storage Model (DSM) is a transposed storage model [5] that stores all values of the same attribute of the relational conceptual schema relation together [4]. Svensson et al. mentioned the Cantor project [6, 7] as the pioneer for this approach [8]. Copeland and Khoshafian in [4] concluded many advantages of DSM. We listed few of them as follows:

- Simplicity (Copeland and Khoshafian related it to RISC [9])
- Less user involvement
- Less performance tuning requirement
- Reliability
- Increased physical data independence and availability

In literature column-oriented [10], vertical fragmentation [11], vertical partitioning [12], etc., are terms that are also used to refer to DSM.

Discussion

Copeland and Khoshafian in [4] analyzed both approaches and concluded that neither of the two approaches could be an ideal solution for all domains. DSM requires relatively more storage space; however, the required storage can be reduced by using compression techniques [13]. Update and retrieval performance of both models depends on the nature of data and implementation of models. DSM is known for fast retrieval whereas NSM is efficient in fast updates [13]. Copeland and Khoshafian suggest that many disadvantages of DSM can be avoided by using hardware and software techniques, such as differential files, multiple disks, large main-memory, and so on [4]. DSM allows the usage of the CPU cache efficiently [14].

Zukowski et al. in [15] compare the two approaches on most recent hardware for CPU performance trade-offs in block-oriented query processing. Zukowski et al. conclude that it depends on the query to identify, which data layout is better, furthermore, they recommend on-the-fly conversion between these formats for better performance and stress on research on hybrid data layout using best of both approaches. Example of hybrid data layout can be found in PAX [16] and MonetDB/X100 [14].

2.1.2 Embedded Database

An embedded database is a data management solution that is embedded into its user-application. However, the same term is also used for a database that resides in an embedded system [17]. An embedded database is transparent to application end-users. An embedded database possesses many special characteristics as mentioned in literature [17, 18, 19, 20, 21]. We list some important embedded database characteristics as follows:

- Small footprint
- Small set of tasks
- Little maintenance
- Multiple-platforms support
- API based access

Cellular DBMS exploits multiple atomic customizable embedded databases at once to generate a large DBMS.

2.2 Software Engineering Aspect

For designing data management architectures, knowledge of software engineering aspects plays an important role. In the end, software engineering techniques are used to implement a data management system. Many researchers already have considered the software engineering aspect while designing data management architectures and found its impact on design decision as too high [20, 22, 23, 24, 25, 26, 27]. In Cellular DBMS, we also consider the software engineering aspects that can benefit us to achieve the targeted goals.

2.2.1 Software Product Lines

Software Product Line (SPL) engineering is an efficient and cost-effective approach to produce a family of related programs for a domain [28]. A product line shares a common set of features developed from a common set of software artifacts [29]. It has been shown that a high degree of customizability makes an SPL a suitable candidate for the development of data management systems [22]. Rosenmüller et al. in [30] and Saake et al. in [25] demonstrated how SPL overcomes the limitation of customizability and performance for data management in embedded systems that exist in other approaches.

2.2.2 Feature-oriented Programming

Feature-oriented programming (FOP) is a mechanism for developing software product lines where programs are synthesized by composing features [31]. A feature can be defined as “a distinguishable characteristic of a concept that is relevant to some stakeholder” [32]. When an SPL is designed in terms of features, creating a program is simply the selection of all required features and automatic composition of the according feature modules [31].

2.2.3 Aspect-oriented Programming

Aspect-oriented programming (AOP) [33] is a methodology that emerged with the aim to separate crosscutting concerns. AOP ensures code scalability and maintenance by preventing code

tangling and scattering [33]. Using AOP, crosscutting code is separated from the program logic using aspects. These aspects, such as data persistence, transaction management, and data security, either can be provided by a software component or could be required by it [33]. Using join-points, pointcuts, and advice, an aspect weaver brings the program code and aspect code together [34]. Join-points are points in the execution of a program and are events of interest for aspect weaving [34]. Pointcuts is the collection of join-points and is used for selection of related method-execution points [34]. An advice is the intended behavior to be weaved [34].

2.3 Autonomy

Autonomy in data management means the capability of DBMS to monitor, diagnose, and tune itself for consistent performance. Autonomy is an essential feature to reduce the human effort in DBMS administration. Automatic administration can reduce the administration cost for data management of large enterprises as well as for embedded systems. "The embedded vendors all acknowledge the need for automatic administration, but fail to identify precisely how their products actually accomplish this" [21]. Similarly, Chang et al. based on their experiences of the Bigtable [35] implementation stressed the importance of proper system-level monitoring of the system itself and its users to detect and fix problems. Autonomous DBMSs monitor themselves and perform tuning operations automatically based on pre-defined policies. A key motivation of Cellular DBMS architecture is to achieve autonomy for self-tuning data management [2, 3].

3. Related Work

Cellular DBMS is an innovation in its own, but it did not appear from nowhere. All concepts and technologies that are joined together in Cellular DBMS have their counterpart in literature and industry. We believe that few concepts are new, but to make such a claim is unrealistic. For decades, many researchers have worked on similar topics and always found the possibility to have similar findings with different names in different domains. Cellular DBMS inherits the SPL-based approach from FAME-DBMS² [20, 36]. Different aspects of Cellular DBMS have to be covered to convince the reader for the originality of Cellular DBMS, such as inspiration from biological systems, use of AOP, and column-oriented storage.

3.1 Term "Cellular DBMS" in Literature

The Infobionics Knowledge Server³ also known as Infobionics Cellular DBMS claims to be first fluid dynamic solution for managing, navigating, and querying data. The Infobionics Cellular Database Management System places information in individual Data Cells, which can be flexibly compiled via Link Cells into an infinite number of DataSets⁴. However, in patent [37] it is stated as "a system for acquiring knowledge from cellular information. The system has a database comprising a database management module ("DBMS")." The concepts presented for Cellular DBMS in current and related publications [38, 39, 40] are different from the ones used by Infobionics Cellular DBMS. We are inspired from human cellular organization whereas in contrast Infobionics Cellular DBMS

is inspired from human brain. For each cell in Cellular DBMS high customizability, limited functionality, and highly predictable behavior is backbone of the concept. Internal architectural details of Infobionics Cellular DBMS are not publicly available; however, based on the available information in form of patent [37] and press releases, we found our work quite different in terms of both concept and implementation.

Kersten et al. in [41] proposed architecture for Cellular database system. According to the proposal, each cell is a bounded container, i.e., a workstation or a mobile unit linked into a communication infrastructure. It assumes the internet as the underlying communication network. This work also envisions a cell as an autonomous DBMS as we do, however, realization of autonomy is different in our approach. We utilized an AOP based model for implementing autonomy. Furthermore, we suggested freedom of using any customizable embedded database as cell.

In 2003, Kersten et al. along with other researchers in [42] again tried to draw the focus of database community towards Organic databases at VLDB. In 2006, Kersten and Siebes took step forward with the concept of an Organic Database System in [43]. They provided the vision of new database architectures as "an Organic Database System where a large collection of connected, autonomous data cells implement a semantic meaningful store/recall information system" [43]. Verroca et al. in [44] used the term Cellular Database for a solution for cellular network data management. Kodama et al. in [45, 46] proposed a Cellular DBMS that is based on the layer model. It is based on incremental modular abstraction hierarchy. Mechanisms are gradually added in it as a global model. They have applied the cellular model to model web-based information spaces for designing the Cellular DBMS [45].

3.2 Embedded Database

3.2.1 COMET

Tesanovic et al. in [24] proposed the concept of aspectual component-based real-time system development (ACCORD) and applied it successfully in the design and development of a component-based embedded real-time database system (COMET). COMET DBMS was developed for resource-constrained embedded vehicle control systems. COMET DBMS is highly tailorable for different requirements and was developed using component-based and aspect-oriented programming approaches. Cellular DBMS also target real-time embedded domain for its variants. It has similarity with COMET DBMS in terms of use of AOP in this domain.

3.2.2 Berkeley DB

Berkeley DB [19] is a customizable embedded database system. Cellular DBMS takes many inspirations from Berkeley DB. Key/Value pairs, API-based access, main memory database, and small footprint all these concepts have their counterpart in Berkeley DB.

3.2.3 FAME-DBMS

FAME-DBMS [20, 36] is developed based on an SPL approach. The SPL approach promises benefits for the embedded domain as proposed by Leich et al. [22]. Our existing Cellular DBMS implementation is an extension of FAME-DBMS. The conceptual approach of Cellular DBMS can be implemented using any customizable embedded database. Additional features of Cellular DBMS that are not part of FAME-DBMS include column-based storage, different cell type implementations, autonomy, and evolution. We have many new features, such as high-level

² "FAME-DBMS", <http://www.fame-dbms.org/>

³ "The Infobionics Knowledge Server", <http://www.infobionics.com/>

⁴ "Cellular DBMS Seeks Business Intelligence Beta Sites", PRESS RELEASE, infobionics, http://www.infobionics.com/news/news_2/file_item.pdf, Accessed: 21-07-2009

composite cells in development phase and many features, such as implicit learning for autonomy are planned as future work. Data management of embedded system is the focus of FAME-DBMS [20, 36]; in contrast, Cellular DBMS is not confined to it. FAME-DBMS focuses on the derivation of concrete instances of a DBMS by composing features of DBMS product line whereas Cellular DBMS derives one or more instances of any DBMS and exploits them in concert for data management.

3.3 AOP for Autonomy

Use of AOP to implement autonomic behavior is not a new concept. Many researchers in past have used it successfully to develop autonomic systems. Greenwood et al. in [47] outlined the case of the use of dynamic AOP for autonomic systems. Truyen et al. in [48] demonstrated the applicability of AOP for implementing self-adaptive frameworks. Tesanovic et al. in [24] proposed the concept of aspectual component-based real-time system development (ACCORD) and applied it successfully in the design and development of a component-based embedded real-time database system (COMET). In Cellular DBMS, we use an AOP based model to implement autonomic behavior at cell as well as at DBMS level.

3.4 Biological Inspiration

To take inspiration from biological systems in computer science is not a new approach. There have been many attempts by many researchers to take benefits from the concepts in biological systems. An important step in this direction was taken by John von Neumann in his work on self-reproduction and cellular automata [49, 50]. We found a major contribution from Gheorghe Păun and Cristian Calude in the area of membrane computing [51, 52, 53, 54]. Kersten and Siebes proposed an organic database system in [43]. It is similar to our approach, but with many differences as we have already discussed in section 3.1. We want to use the best of it regarding how to take the inspiration from biological systems.

4. Cellular DBMS

A Cellular DBMS is composed of multiple atomic and autonomic customizable embedded database instances, called Cells [38, 39]. The motivation behind this approach is to ensure that a DBMS can be reduced to a fine-grained atomic unit (i.e., a cell) with a predictable behavior and reduced complexity [2]. This approach enables us to assess the behavior of a complete DBMS by accumulating the behavior of all atomic cells.

4.1 DBMS Cell

A Cell is an atomic and autonomic instance of a customized embedded database [38, 39]. Each cell is based on RISCstyle architecture with simple and limited functionality. A cell can be customized based on different criteria, such as hardware, software, application scenario, nature of data, etc. Decisions about cell composition require a detailed analysis of all these criteria. Cellular DBMS architecture restricts cell functionality to a manageable complexity. It ensures that each cell is optimized for its task and is predictable for its performance. Each cell is customized to handle a single kind of data (i.e., data with unique characteristics, e.g., aggregated data as shown in Table 2). If a cell supports handling multiple tables than the same kind of data should be stored in these tables. It ensures customization of each cell according to the kind of data. Multiple cells should be used to handle different kinds of data.

The most fine-grained variant of the cell can handle key/value pairs of data. Variants that are more complex can handle tables

| Kind of Data | Standing | Setup | Transactional | Aggregated |
|---|----------|--------|---------------|------------|
| Data Size | S | S to M | M to R | M to R |
| Read Frequency | H | H | M to H | L to H |
| Write Frequency | O | L | M to H | L to M |
| O=No Write, L=Low, S=Small, M=Medium, H=High, R=Large | | | | |

Table 1. Data categorization

and maintain data dictionary, however, complex variants should be composed by using multiple fine-grained variants of cell. As mentioned above, the simplest cell handles a key/value pair and has definite (optimal) data-handling capacity, however, with the data growth; more cells could be induced into DBMS to extend its data-handling capacity. Virtually each cell uses Binary Fission [55] mechanism to grow. In binary fission, biological cell grows to twice of its starting size and then split-up into two cells, each cell having a complete copy of its essential genetic material. Not exactly, but similarly each DBMS cell splits into two equal halves. One-half is left in the parent cell where as the other half is moved to a newly induced cell.

Deployment of cells depends on many criteria, such as the kind of data, the distribution of computing resources⁵, as well as the hardware of computing resources, etc. For example, in the simplest sensor network scenario, a single cell can be deployed on an individual node. However, for more complex scenario, multiple cells can be deployed on a single node or can be distributed over multiple nodes in a network.

4.2 Types of DBMS Cells

Cellular DBMS defines many different types of cells. Each type differs from the other based on its composition and characteristics. These types enhance the diversity of Cellular DBMS for many data management scenarios. Currently, implementation of Cellular DBMS is a work in progress. More cell types are expected to appear in the future. In this paper, we explain the types that we have defined based on our existing architecture and implementation.

4.2.1 Composite Cells

A cell can be composed of multiple similar or dissimilar cells related to each other as shown in Figure 2. Such composition of cells is termed as Composite Cell. Each composite cell should have limited (optimal) data-handling capacity to ensure it has manageable complexity and predictable performance. With the data growth, more composite cells could be induced into the DBMS to extend its data management capacity. Each composite cell maintains a meta-data of cell composition. Composite cell can be used to implement a table in Cellular DBMS where each column is implemented by a cell that could be of different type, e.g., one column cell contains in-memory data management functionality whereas another column cell can also store persistent data. It can also be used to handle large amount of data that simple cells cannot handle. From software engineering perspective, when using multiple cells, composite cells avoid code replication and allow us to reuse the program code between different cells on a single computing resource.

4.2.2 High-level Composite Cells

In Cellular DBMS, composite cells can be built from simple cells, as well as from composite cells, which results in highlevel composite cells as shown in Figure 2. The reason for such

⁵ From computing resource, we mean any device with processing capability. It may range from small-embedded device to high-end enterprise server machines.

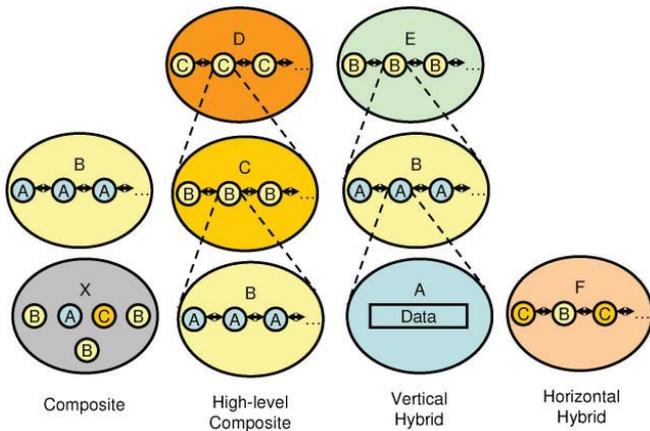


Figure 2. Different types of cells

architecture is to provide hierarchical data management functionalities to manage complexity. According to Cellular DBMS architecture, data-handling capacity of a cell is optimally limited for highly predictable performance and reduced complexity. Cellular DBMS uses high-level composite cell for handling large amount of data. In high-level composite cell, cell composition becomes deeper with the increase in the size of data. Each high-level composite cell maintains a meta-data cell (i.e., stores metadata) that helps in fast retrieval and updation of records. To retrieve data from high-level composite cell, only meta-data cells are used to trace the data cell. Once the targeted data cell is traced, only this cell or its related cells are used for data management.

4.2.3 Hybrid Cell

For diversified data management, Cellular DBMS introduce the concept of Hybrid Cell. We could have horizontal as well as vertical hybrid cells as shown in Figure 2. From horizontal hybrid cell, we mean a composite cell that is composed of different type of cells such that each type is handling a definite data range. For example, we want to store city codes to be used in the contact book of a mobile phone product. If mobile is to be used in European Union (EU), frequency to access city codes of EU countries is much higher as compared to city codes of Australia. Using horizontal hybrid cell, we can store data in a composite cell in such a way that EU city codes should be stored in cell with a type that is suitable for faster access time whereas we store remaining city codes in a cell, which requires less storage space. We can exploit this feature in conjunction with autonomy to move data among different cells based on their usage scenario and available resources.

From vertical hybrid cell, we mean a high-level composite cell that is composed of different type of cells at different levels. For example, we have In-Memory Data Management type cell at the fine-grained level, i.e., level 0. At one level above, i.e., level 1, we have B+Tree composite cell using multiple In-Memory Data Management cells, and finally one more level above, i.e., level 2, we have SortedList using multiple B+Tree composite cell. Vertical hybrid cell can be generated using the evolution approach discussed in this paper; however, implementation of hybrid cells is future work.

4.2.4 Evolving Cell

Evolution in Cellular DBMS means run-time transformation of cells. We term a cell that supports evolution as an "Evolving Cell". Evolution can be constructive as well as destructive. From constructive evolution, we mean the transformation of a cell from one form to another in such a way that the previous

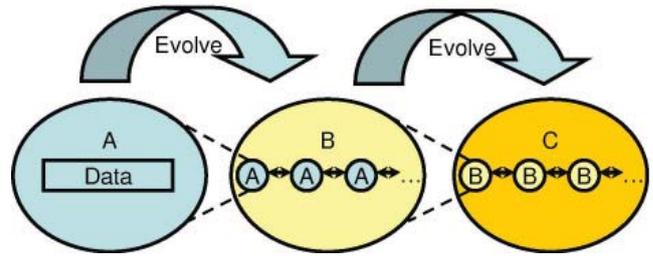


Figure 3. Evolving cell

form becomes an atomic integral unit of new form as shown in Figure 3. New form of such an evolved cell should have larger data-handling capacity. Evolution is a mandatory concept to bring autonomy in Cellular DBMS. For example, consider a cell X that is initially an in-memory data management cell. We also support a SortedList that stores data using multiple in-memory data management cells. SortedList is the simplest composite cell. From evolution, we mean the transformation of cell X to SortedList so that cell X becomes an atomic integral unit of SortedList.

4.3 Clean API and Interaction

From software engineering aspect, providing a consistent API for simple as well as composite cells is an important design criterion, which is required for communication between cells. We argue that, two communicating cells should not care about the concrete type of one-another. On the other hand, simple cells provide limited data management functionality and should exhibit a simple API that reflects the limited functionality, which is in contrast to a consistent API and has to be considered when generating cells.

For solution, we use two different mechanisms. First, we allow only interface extensions, but not modifications of an interface [39]. For example, a DBMS feature might add a method to the interface of the DBMS, but is not allowed to modify the signature of an existing method. This ensures upward compatibility, i.e., we can use cells with a more complex API when cells with a simple API are expected.

The second approach is to generate wrappers for simple cells when complex cells are expected [39]. For example, if a method for creating an index is expected from an in-memory cell without index support, an empty wrapper method can be generated to provide this method. Wrappers are used to achieve only downward compatibility and wrappers that are more complex might be required. Furthermore, it has to be analyzed for which scenarios it is not possible to generate such wrappers.

4.3.1 Distributed Cells

In Cellular DBMS, cells are not confined to a single computing resource. Cells can be distributed across network, or more ambitiously speaking across internet. Important distribution criteria could be size and locality of data. For example, in a complex distributed sensor network scenario, cells are deployed on multiple nodes and collaborate for data management. On each node of such a distributed scenario, a single cell might be used or a composite cell might provide complex data management. Distributed cells interact with each other through API calls over the network. For distributed deployment, we envision a Cellular DBMS using a global data dictionary and statistics as well as distributed monitoring functionality to implement distributed autonomy. However, it has to be further analyzed how distributed deployment of interacting cells can be achieved in Cellular DBMS.

4.4 Cell Classification

Based on our current architecture and implementation, we can also classify cells in two types based on the data they store, i.e., data cell and meta-data cell. Data cell manages data. Meta-data cell is also a data cell, but it stores metadata.

4.5 Design Principles for Autonomy in Cellular DBMS

Autonomy of each cell is an important design principle for Cellular DBMS. Cellular DBMS envision the development of complete autonomous DBMS by accumulating autonomic behavior of all participating cells. For autonomy, the most fundamental functionalities are Monitoring, Diagnostics, and Tuning [56, 57]. According to the proposed architecture, monitoring, diagnostic, and tuning components should also be customizable according to the cell functionalities to ensure reduced monitoring overhead. We present an AOP based model for autonomy at cell-level. We argue based on provided related work in section 3.5 that AOP join-point model can be used to implement efficient monitoring functionality for data management.

According to Cellular DBMS architecture, each cell contains an optional lightweight monitoring functionality. The purpose of monitoring functionality is to monitor the cell for specific parameters. These parameters are defined as a policy for DBMS cell goals. Each cell should be able to adapt to changes based on events identified by the monitoring component. Additional to a cell-level monitoring, there should be a monitoring component at composite cell as well. It should get feedback from an individual cell monitoring component and should by itself monitor certain parameters at composite cell level. It enables global monitoring of cells for adaptation to DBMS changes and fixing of DBMS problems according to defined DBMS policy. A symbolic monitoring functionality distribution is depicted in Figure 4. For diagnostics, we use the state of the cell, and results of data management operations to identify the definite tuning points. For tuning, we use the evolution and evolving cell approach presented in section 4.2.

According to the model, tracing is an important functionality during monitoring. By tracing, we mean collection of cell state information that is needed to diagnose and tune the individual cell as well as complete Cellular DBMS. For each join-point, a before advice should be used for tracing whereas an after advice

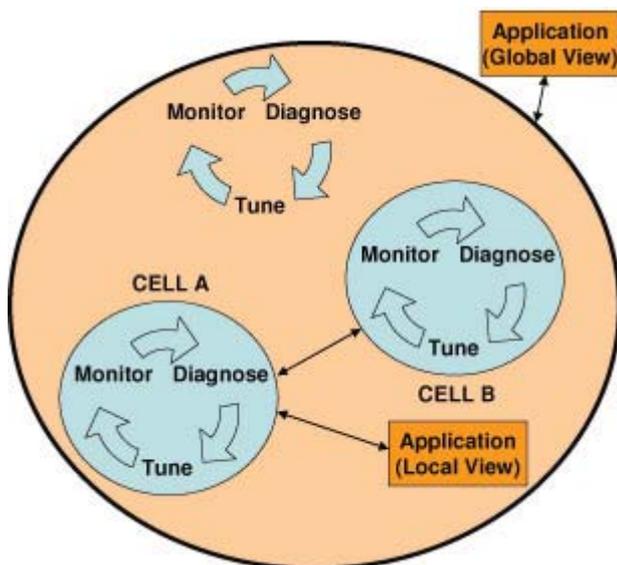


Figure 4. Monitoring functionality distribution

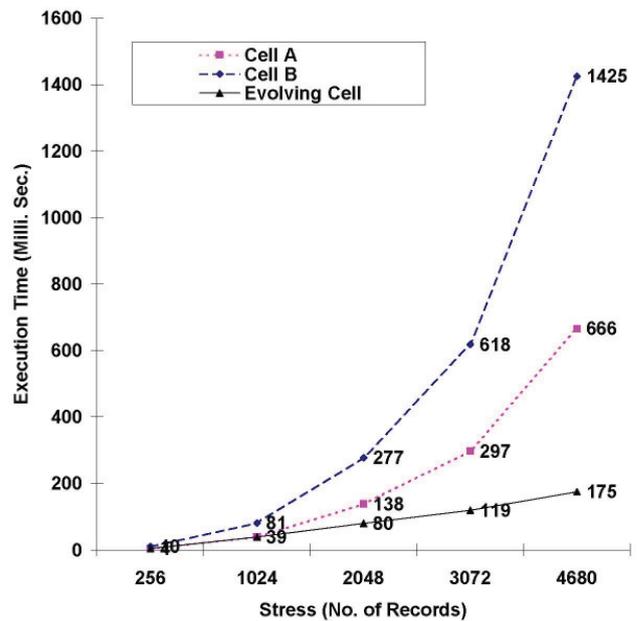


Figure 5. Average execution time graph for stress test in millisecond for different Cellular DBMS cells

should be used to diagnose the abnormality. If any abnormality is detected during diagnostics, tuning should be executed to counter the abnormality.

To explain the concepts in more detail, we describe a scenario. We compose a Cellular DBMS that supports an in-memory data management cell and an in-memory data management composite cell, i.e., a SortedList. We term in-memory data management cell as Cell A and in-memory data management composite cell as Cell B. Cell A stores data in a single memory chunk whereas Cell B is composed from multiple Cell A. It is also shown in Figure 3. Both cells store definite/limited amount of data, however, capacity of data storage in Cell B is larger. In contrast, the complexity and main-memory requirement of Cell A is relatively low. To differentiate the behavior of two cells we presented the average execution time in millisecond of stress test on both cells in Table 1⁶ and Figure 5. We executed test with different stress values, i.e., number of records that are inserted, retrieved, and deleted. For Cell A, we kept memory allocation large enough to accommodate all test data into a single cell. For Cell B, we kept memory allocation of each Cell A small enough so that multiple cells can be used to demonstrate the change in behavior. It can be observed that Cell A performs much faster than Cell B, because of reduced execution complexity. Cell A also consumes less main-memory, because of simple data management structure. Based on the results, we argue that cell complexity should only be increased with the data growth. For example, we should use the Cell A as long as the data is small enough for it to handle. As data grows to exceed the limit of Cell A capacity, we bring the concept of evolving cell to evolve cell from type A to type B, i.e., Cell A becomes part of Cell B and evolved cell has relatively larger data management capability. In Cellular DBMS, we can evolve cells to higher level, e.g., compose Cell C based on multiple B cells and so on. Autonomy should be kept at the fine-grained level of Cell A to ensure highly predictable and tunable behavior at the smallest data management unit.

⁶Average execution time is used to demonstrate the concept and may vary in future work.

| Cell Type | Stress (No. of Records) | | | | |
|------------------|-------------------------|------|------|------|------|
| | 256 | 1024 | 2048 | 3072 | 4680 |
| Composite Cell A | 4 | 39 | 138 | 297 | 666 |
| Composite Cell B | 10 | 81 | 277 | 618 | 1425 |
| Evolving Cell | 4 | 39 | 80 | 119 | 175 |

Table 2. Average execution time for stress test in milliseconds for different Cellular DBMS cells

To generate better results, we first analyzed the optimal memory allocation of Cell A that resulted in the fastest execution time for stress test using Cell B. We observed that for our sample stress data, both, i.e., too small as well as too large memory allocation was found to be inefficient. Once we identified the optimal memory allocation for Cell A, our evolving cell implementation uses Cell A until its data management limit is reached. Monitoring component keeps monitoring the Cell A based on join-point specification and keeps trace of the required information. As soon as our diagnostic implementation detects that Cell A is out of memory, it executes the tuning implementation, which evolves Cell from type A to B by injecting Cell A in Cell B. From end-user and application point of view, it is kept transparent when evolution occurs. By using this approach, we ensure that complexity of data management implementation should only be increased as the amount of data is increased.

4.6 Cellular DBMS Storage Model

Customization capability of Cellular DBMS gives us provision to use any of the available storage models. Cellular DBMS does not restrict the usage of any specific storage model; however, we recommend one in this section based on the Cellular DBMS goals. Cellular DBMS stores data using Decomposed Storage Model (DSM) [5] also known as Column-oriented Storage (COS) [10]. Based on the discussion in the related concepts section, we found COS most appropriate for implementing atomic and autonomous cells. Use of COS enables simple cell design and gives more control over data. We envision achieving all benefits from COS as discussed in related concepts section. In Cellular DBMS, each column is a separate cell. A column data can be stored using a simple as well as composite cell. COS in Cellular DBMS is shown in Figure 6.

5. Cellular DBMS Implementation

Existing Cellular DBMS implementation is an extension of FAME-DBMS, a highly customizable embedded database

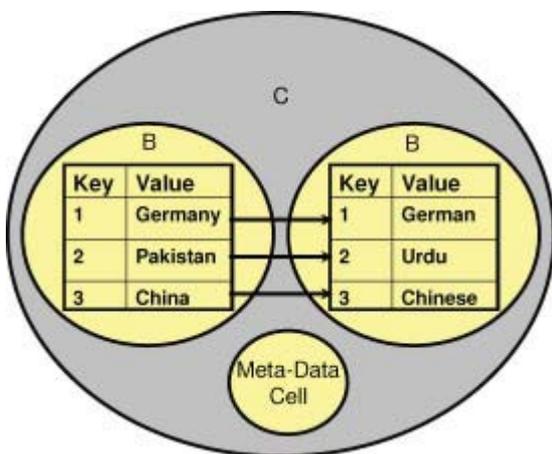


Figure 6. Column-oriented storage using composite cells

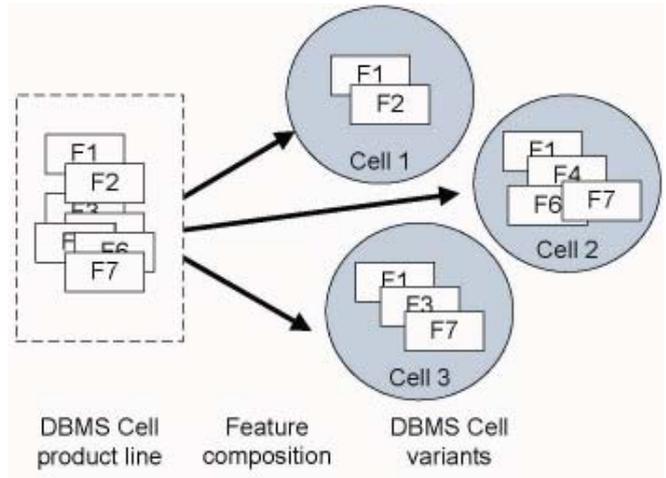


Figure 7. Generating different DBMS cells by composing features (F1--F7) of a DBMS cell product line [39]

management software product line developed for deeply embedded systems [20, 36]. We used FAME-DBMS to generate the cells; however, any customizable embedded database can be used to generate cells. FAME-DBMS is implemented using feature-oriented programming. It untangles and modularizes DBMS functionalities as features. A decomposition of DBMS into features, i.e., the functionalities individual DBMS differ in, allows a developer to generate a tailor-made DBMS variants based on the selection of required features [22, 58]. These different variants are built from the same code base as depicted in Figure 7 [39]. Based on such an SPL, multiple heterogeneous DBMS cells can be generated [20].

The feature model of Cellular-DBMS, as shown in Figure 8⁷, is based on the FAME-DBMS feature model originally published in [20]. A feature model describes the features of an SPL and

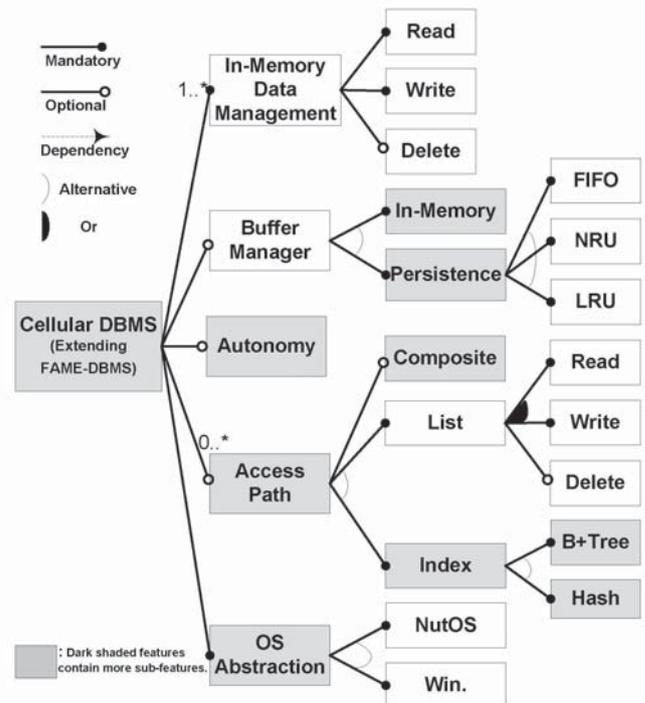


Figure 8. Cellular DBMS feature model

⁷Shown is an excerpt of the feature model with only modified features/concepts required for discussion.

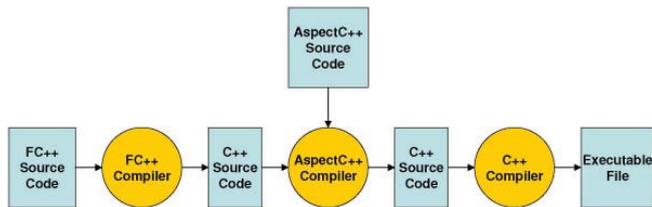


Figure 9. Source code transformation

their relationships [32]. As depicted, the implementation of Cellular DBMS consists of five main features, i.e., In-Memory Data Management, Buffer Manager, Access Path, Autonomy, and OS Abstraction. Every functionality can be implemented differently to achieve benefits, e.g., better performance, and can be described as alternative features. For example, feature Index provides two variants for effective data access, i.e., B+Tree and Hash. It enables us to generate specialized cells by selecting one feature or the other.

Functionality for storing data is provided by feature In-Memory Data Management. This feature contains the functionality of an in-memory embedded database. It can alone be used to construct a simple DBMS cell. It performs data management operations in an in-memory environment and does not have any unneeded persistence functionality, resulting in good performance in terms of fast operations [59]. For example, in the sensor network scenario, most sensor nodes are equipped with storage memory that can be used to store data persistently. To scale a Cellular DBMS cell for such nodes feature Persistence can be used.

The simplest cell, consisting only of feature In-Memory Data Management, supports exactly one column or a fragment of column (if used as a part of high-level composite cell). For multiple columns, we can clone the In-Memory Data Management feature. Cloning a feature means to create multiple instances of it [60]. For example, to support two columns we have to create two instances of the In-Memory Data Management feature and each instance handles one of the columns. Whether a feature can be cloned is depicted with cardinalities in Figure 8. For example, there has to be at least one instance of the In-Memory Data Management feature, but an arbitrary number of instances are allowed.

We bring autonomy to each cell by using an AOP based model. We utilized AspectC++⁸ [61] for using AOP constructs. FeatureC++ also supports AOP extensions as discussed in [62, 63, 64], however, AspectC++ is used independently to control AOP constructs more easily. The code transformation model is shown in Figure 9 that we used in our implementation.

6. Discussion

We envision that Cellular DBMS architecture is scalable for use in embedded systems to enterprise systems. For explanation, we discuss two assumed sample scenarios for Cellular DBMS use in sensor networks and enterprise data management.

6.1 Sensor Networks

Sensor networks are important data-centric systems with hardware and software heterogeneity as depicted in Figure 10. Hardware in sensor networks may vary from 8-bit motes to 32-bit microservers with the program memory that can vary from 48 kB to 512 kB, whereas the data memory may vary from 4 kB to

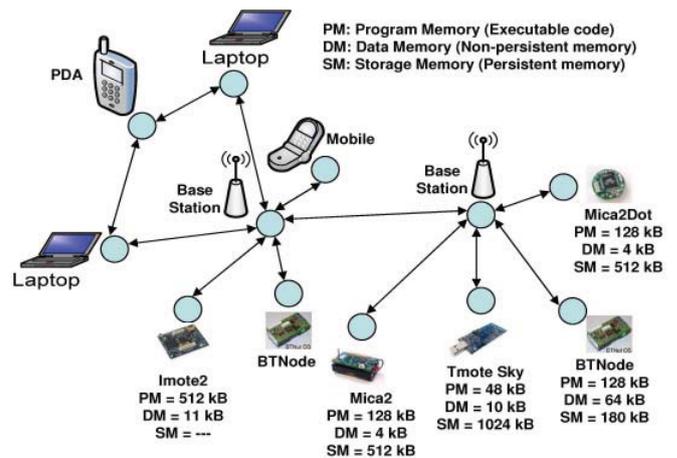


Figure 10. Sensor network scenario [39]

64 kB [65]. Each node varies in terms of the processing power and the memory configuration. Considering extreme resource scarcity and high hardware and software heterogeneity as discussed above, one of the requirements of sensor networks is to make the best use of available resources and exploit the hardware heterogeneity for efficient data management. For deployment on sensor networks, each cell should be customized based on the resources and kind of data that cell handles on the deployment node.

For discussion of our proposed architecture, we consider storage memory and program memory as parameters of interest. To explain the idea, how specialized DBMS cells can be beneficial for data-centric embedded systems, four types of DBMS cells are generated based on different feature selections using FAME-DBMS prototype as shown in Table 3⁹. For each cell, the binary size is different and depends on the selected features of FAME-DBMS prototype. Each cell is a candidate for a different type of node based on the available program and storage memory as well as type of data it handles. Cell A is suitable for nodes without any storage memory, e.g., Imote node. Cell D is suitable for nodes with relatively large data and storage memory, e.g., BTnode rev3. A sample deployment of these customized cells based on a node's resources is shown in Figure 11.

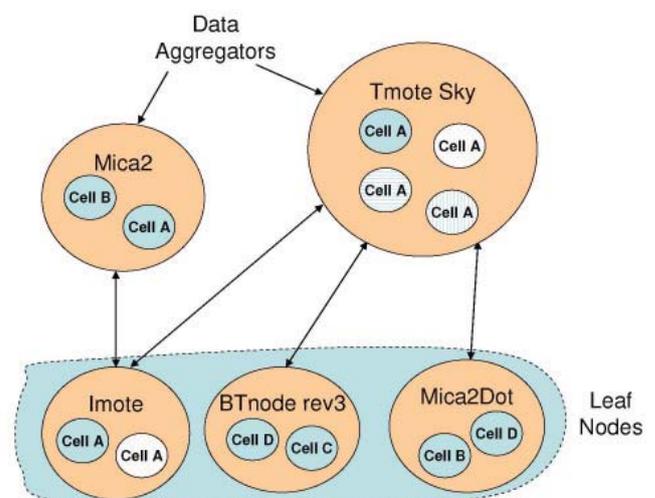


Figure 11. Sample deployment of different Cellular DBMS cells [39]

⁸ "AspectC++", <http://www.aspectc.org/>

⁹ Binary size contains additional overhead of dependencies and may vary in future work.

| Features | In-Memory Data Management | | | Persistence | List | | | Index | Binary Size (.ELF) Compiler: avr-g++ |
|----------|---------------------------|-------|--------|-------------|------|-------|--------|--------|--------------------------------------|
| | Read | Write | Delete | LRU | Read | Write | Delete | B+Tree | |
| Cell A | x | x | x | | | | | | 32 KB |
| Cell B | x | x | x | x | x | | | | 42 KB |
| Cell C | x | x | x | x | x | x | x | | 45 KB |
| Cell D | x | x | x | x | | | | x | 48 KB |

Table 3. Binary size for different Cellular DBMS cells [39]

In the sample deployment, Tmote sky contains the smallest program memory and can only handle small cells like Cell A. However, it also contains largest storage memory making Cell B, C, and D a good candidate for storing relatively large data on storage memory. In contrast, Imote contains the largest program memory, but lacks storage memory, again making Cell A best candidate for deployment. BTnode rev3, Mica2, and Mica2Dot all contain moderate program and storage memory. We argue that, in the demonstrated sample deployment, the Cellular DBMS is a promising solution. The Cell implementation is light-weight and allows for deployment of multiple heterogeneous cells on a single node, enabling specialized handling of data based on available resources and the nature of data.

6.2 A Distributed Virtual Reality Scenario

Virtual Reality (VR) is intensively used in planning and development of new complex products, e.g., in the automotive domain. This development process is also called virtual engineering. However, there is no restriction of the co-operation of virtual and real products over the complete product life cycle. Furthermore, different applications, such as simulation, computation, and visualization, are used here. An application of the interaction of VR and embedded devices is given by Köppen et. al [66]. Cellular DBMS can improve the scenario by customization and adaptation on the hardware, software, and application requirements. This is respected by the fact that data has to be stored and accessed differently for different types of data.

A simulation application for instance uses data as a parameter input on the one hand. On the other hand, the new (simulated) data has to be stored for further usage. In the contrast, visualizing the data has only read access. For both different parts of the VR environment, a type of cell is responsible for the state of predictability. Moreover, in a distributed VR scenario all data has to be accessible at nearly the same time in different workplaces. In such a scenario, a domain expert can view the regarding information in the VR, e.g., an automobile designer views the car body, whereas the electrical engineer makes changes of the wiring at the same time. The car data has to be merged on the one hand to have a consistent state, on the other only domain or application relevant data has to be processed. Using Cellular DBMS enhances the use due to more efficient handling of information and context-adapted functionalities of data provision. It can allow managing different data schema, formats, and storage systems that exist in VR environment with a single DBMS view.

6.3 Enterprise Data Management

Industries with enterprise data management needs are quite satisfied with existing DBMSs in term of their performance. They have many solutions to choose from based on different criteria, e.g., cost, performance, etc. However, maintenance cost of most of the existing DBMSs is high. We argue that Cellular DBMS architecture with its goals to achieve highly predictable, customizable, autonomous DBMS will be able to reduce the maintenance cost.

The data classification we provided in Table 2 is also applicable for enterprise data management. Cellular DBMS gives an end-

user the provision for data management customization at many different levels. An end-user can specify, what functionalities DBMS should have, how these functionalities should be used, and how DBMS should be tailored based on application and data [20, 22, 25, 27, 36, 58, 67]. For example, consider the case for setup data. The frequency of update in setup data is low whereas frequency of retrieval is high. Furthermore, setup data (except some exceptions) is not too large. In an enterprise application, we normally have many setup tables. Existing DBMS handles all tables similarly. Either we have large data or not, we cannot customize the internal implementation to optimize it for handling small data. The same column implementation is used for handling a column with only five values as well as for a column with many MBs of data. Cellular DBMS approach is different. It provides the provision for customization at the fine-grained level of cell. It is possible to compose different cells based on the kind of data, i.e., we can compose four types of cells for all four kinds of data we presented in Table 2. This approach is similar to using four small databases customized for their task instead of using a single large DBMS customized for nothing. Another important aspect is that existing DBMS uses complex internal structures irrespective of existing data size. In contrast, in Cellular DBMS data management complexity only increases with the data growth, i.e., utilizing the resources only when they are needed.

7. Conclusion and Future Work

We proposed a novel DBMS architecture based on composition of multiple cells that are atomic and autonomic customized embedded databases. As explained, these cells provide restricted data management functionality and collaborate to constitute one large Cellular DBMS. This cellbased approach ensures predictable behavior and efficient utilization of resources by keeping the cells simple.

We argue that Cellular DBMS architecture reduces DBMS complexity and when blended with autonomy, it can be used to develop highly predictable autonomous DBMS. In this work, we also presented an AOP based model for implementing autonomy at cell level in Cellular DBMS. We also explained the idea how evolving cells can be used to self-tune data management with data growth. Our presented implementation ensures that initially for small amount of data, simpler data management functionality is used. We evolve the functionality with the data growth maintaining consistent performance. In our proposed architecture, we argue that we can develop highly customizable autonomous DBMS that can scale from requirement of small embedded systems to large-scale enterprise systems.

As a future work in Cellular DBMS, we found many opportunities that are listed below:

- Many concepts, such as hybrid cell, cell mobility, resource balancing, self-* [68] (e.g., self-tuning, selfmanaging, self-adaptation, etc.) capabilities, etc., that we presented in this paper need implementation and performance comparison with existing approaches.
- In the era of multi-core processors, we want to enable Cellular DBMS to exploit parallelism.
- Using differently composed cells simultaneously while minimizing code replication is an important open issue. A software engineering based solution is needed to solve this issue.
- Monitoring is an overhead for high-end embedded system. For implementation of Cellular DBMS in such systems, we want to investigate mechanisms to reduce this overhead.

- Current implementation of cell evolution is explicitly programmed. An important future direction is to enable implicit learning in Cellular DBMS for self-* capabilities.
- Query processing is a mandatory feature for all existing DBMS. For Cellular DBMS architecture, we need specialized mechanism for efficient query processing.

Acknowledgments

We thank Marko Rosenmüller, Sven Apel, Norbert Siegmund, Christian Kästner, Ingolf Geist, Azeem Lodhi, Qaisar Hayat, Ateeq Lodhi, and Sagar Sunkle for their feedback and contribution. Syed Saif ur Rahman is funded by Higher Education Commission of Pakistan and NESCOM, Pakistan. Veit Köppen is funded by the German Ministry of Education and Science (BMBF) within the ViERforES project (no. 01IM08003C).

References

- [1] Agrawal, R., Ailamaki, A., Bernstein, P. A., Brewer, E. A., Carey, M. J., Chaudhuri, S., Doan, A., Florescu, D., Franklin, M. J., Garcia-Molina, H., Gehrke, J., Gruenwald, L., Haas, L. M., Halevy, A. Y., Hellerstein, J. M., Ioannidis, Y. E., Korth, H. F., Kossmann, D., Madden, S., Magoulas, R., Ooi, B. C., O'Reilly, T., Ramakrishnan, R., Sarawagi, S., Stonebraker, M., Szalay, A. S., Weikum, G. (2009). The Claremont report on database research, *Commun. ACM*, v. 52, pp. 56-65.
- [2] Weikum, G., Moenkeberg, A., Hasse, C., and Zabback, P. (2002). Self-tuning database technology and information services: from wishful thinking to viable engineering, *VLDB '02: In: Proceedings of the 28th international conference on Very Large Data Bases*, VLDB Endowment, pp. 20-31.
- [3] Chaudhuri, S., Weikum, G. (2000). Rethinking Database System Architecture: Towards a Self-Tuning RISC-Style Database System, *VLDB '00: In: Proceedings of the 26th International Conference on Very Large Data Bases*, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., pp. 1-10.
- [4] Copeland, G. P., Khoshafian, S. N. (1985). A decomposition storage model, *SIGMOD '85: In: Proceedings of the 1985 ACM SIGMOD international conference on Management of data*, New York, NY, USA: ACM, pp. 268-279.
- [5] Batory, D. S. (1979). On searching transposed files, *ACM Trans. Database Syst.* v. 4, pp. 531-544.
- [6] Karasalo, I., Svensson, P. (1983). An overview of cantor: a new system for data analysis, *SSDBM'83: In: Proceedings of the 2nd international workshop on Statistical Database Management*, Berkeley, CA, US: Lawrence Berkeley Laboratory, pp. 315324.
- [7] Karasalo, I., Svensson, P. (1986). The design of cantor: a new system for data analysis, *SSDBM'86: In: Proceedings of the 3rd international workshop on Statistical and scientific database management*, Berkeley, CA, US: Lawrence Berkeley Laboratory, pp. 224-244.
- [8] Svensson, P. (2008). The Evolution of Vertical Database Architectures ---A Historical Review (Keynote Talk), *SSDBM '08: In: Proceedings of the 20th International conference on Scientific and Statistical Database Management*, Berlin, Heidelberg: Springer-Verlag, pp. 3-5.
- [9] Patterson, D. A., Ditzel, D. R. (1980). The case for the reduced instruction set computer, *SIGARCH Comput. Archit. News*, vl. 8, pp. 25-33.
- [10] Stonebraker, M., Abadi, D. J., Batkin, A., Chen, X., Cherniack, M., Ferreira, M., Lau, E., Lin, A., Madden, S., O'Neil, E., O'Neil, P., Rasin, A., Tran, N., Zdonik, S. (2005). C-store: a column-oriented DBMS, *VLDB '05: In: Proceedings of the 31st international conference on Very large data bases*, VLDB Endowment, pp. 553-564.
- [11] de Vries, A. P., Mamoulis, N., Nes, N., Kersten, M. L. (2001). *Efficient image retrieval by exploiting vertical fragmentation*, Amsterdam, The Netherlands.
- [12] Abadi, D.J., Marcus, A., Madden, S. R., Hollenbach, K. (2007). Scalable semantic web data management using vertical partitioning, *VLDB '07: In: Proceedings of the 33rd international conference on Very large data bases*, VLDB Endowment, pp. 411-422.
- [13] Holloway A. L., DeWitt, D. J. (2008). Read-optimized databases, in depth, *Proc. VLDB Endow.*, v. 1, pp. 502-513.
- [14] Zukowski, M., Boncz, P. A., Nes, N., Heman, S. (2005). MonetDB/X100 -A DBMS In The CPU Cache, *IEEE Data Engineering Bulletin*, v. 28, pp. 17-22.
- [15] Zukowski, M., Nes, N., Boncz, P. (2008). DSM vs. NSM: CPU performance tradeoffs in block-oriented query processing, *DaMoN '08: In: Proceedings of the 4th international workshop on Data management on new hardware*, New York, NY, USA: ACM, pp. 4754.
- [16] Ailamaki, A., DeWitt, D. J., Hill, M. D. (2002). Data page layouts for relational databases on deep memory hierarchies, *The VLDB Journal*, v. 11, pp. 198-215.
- [17] Olson, M. A. (2000). Selecting and Implementing an Embedded Database System, *Computer*, v. 33, pp. 27-34.
- [18] Nori, A. (2007). Mobile and embedded databases, *SIGMOD '07: In: Proceedings of the 2007 ACM SIGMOD international conference on Management of data*, New York, NY, USA: ACM, pp. 11751177.
- [19] Olson, M. A., Bostic, K., Seltzer, M. (1999). Berkeley DB, *ATEC '99: In: Proceedings of the annual conference on USENIX Annual Technical Conference*, Berkeley, CA, USA: USENIX Association, p. 43
- [20] Rosenmüller, M., Siegmund, N., Schirmeier, H., Sincero, J., Apel, S., Leich, T., Spinczyk, O., Saake, G. (2008). FAME-DBMS: tailor-made data management solutions for embedded systems, *SETMDM '08: In: Proceedings of the 2008 EDBT workshop on Software engineering for tailor-made data management*, New York, NY, USA: ACM, pp. 1-6.
- [21] Seltzer, M. I., Olson, M. A. (1999). Challenges in embedded database system administration, *WOES'99: In: Proceedings of the Workshop on Embedded Systems*, Berkeley, CA, USA: USENIX Association, p. 11.
- [22] Leich, T., Apel, S., Saake, G. (2005). Using Step-Wise Refinement to Build a Flexible Lightweight Storage Manager, *ADBIS '05: In: Proceedings of the 9th East-European Conference on Advances in Databases and Information Systems*, Springer Verlag, pp. 324-337.
- [23] Özsu, M. T., Yao, B. (2001). Building component database systems using CORBA, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- [24] Tesanovic, A., Te anovic, R., Nyström, D., J.r. Hansson., Norström, C. (2003). Towards Aspectual Component-Based Development of Real-Time Systems, *In: Proceedings of the 9th International Conference on Real-Time and Embedded Computing Systems and Applications*, Springer-Verlag, pp. 278-298.

- [25] Saake, G., Rosenmüller, M., Siegmund, N., Kästner, C., Leich, T. (2009). Downsizing Data Management for Embedded Systems, *Egyptian Computer Science Journal (ECS)*, v. 31, pp. 1-13.
- [26] Pukall, M., Leich, T., Kuhlemann, M., Rosenmüller, M. (2007). Highly Configurable Transaction Management for Embedded Systems, *{AOSD} Workshop on Aspects, Components, and Patterns for Infrastructure Software*.
- [27] Rosenmüller, M., Leich, T., Apel, S., Saake, G. (2007). Von Mini-über Micro-bis zu Nano-DBMS: Datenhaltung in eingebetteten Systemen, *Datenbank Spektrum*, v. 7, pp. 33-43.
- [28] Pohl, K., Böckle, G., Van der, F. J. (2005). *Software Product Line Engineering: Foundations, Principles and Techniques*, Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- [29] Clements, P., Northrop, L., Northrop, L. M. (2001). *Software Product Lines : Practices and Patterns*, Addison-Wesley Professional.
- [30] Rosenmüller, M., Apel, S., Leich, T., Saake, G. (2009). Tailor-Made Data Management for Embedded Systems: A Case Study on Berkeley DB, *Data and Knowledge Engineering (DKE)*.
- [31] Batory, D., Sarvela, J. N., Rauschmayer, A. (2003). Scaling step-wise refinement, *ICSE '03: In: Proceedings of the 25th International Conference on Software Engineering*, Washington, DC, USA: IEEE Computer Society, pp. 187-197.
- [32] Kang, K. C., Cohen, S. G., Hess, J. A., Novak, W. E., Peterson, A. S. (1990). *Feature-Oriented Domain Analysis (FODA) Feasibility Study*.
- [33] Kiczales, G., Lamping, J., Mendhekar, A., Maeda, C., Lopes, C. V., Loingtier, J., Irwin, J. (1997). Aspect-Oriented Programming, *ECOOP '97: In: Proceedings of the 11th European Conference on Object-Oriented Programming*, Jyväskylä, Finland: Springer-Verlag, pp. 220-242.
- [34] Kiczales, G., Hilsdale, E., Hugunin, J., Kersten, M., Palm, J., Griswold, W. G. (2001). An Overview of AspectJ, *ECOOP '01: In: Proceedings of the 15th European Conference on Object-Oriented Programming*, London, UK: Springer-Verlag, pp. 327-353.
- [35] Chang, F., Dean, J., Ghemawat, S., Hsieh, W. C., Wallach, D. A., Burrows, M., Chandra, T., Fikes, A., Gruber, R. E. (2006). Bigtable: a distributed storage system for structured data, *OSDI '06: In: Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation*, Berkeley, CA, USA: USENIX Association, p. 15.
- [36] Siegmund, N., Rosenmüller, M., Moritz, G., Saake, G., Timmermann, D. (2009). Towards Robust Data Storage in Wireless Sensor Networks, *DAIT '09: In: Proceedings of Workshop on Database Architectures for the Internet of Things*, Springer.
- [37] Sabry, J. H., Adams, C. L., Vaisberg, E. A., Crompton, A. (2003). *Database system for predictive cellular bioinformatics*, United States Patent 6631331.
- [38] Ur Rahman, S. S., Lodhi, A., Saake, G. (2009). Cellular DBMS -Architecture for Biologically-Inspired Customizable Autonomous DBMS, *NDT '09: In: Proceedings of the First International Conference on the Networked Digital Technologies*, Washington, DC, USA: IEEE Computer Society, pp. 310315.
- [39] Ur Rahman, S. S., Rosenmüller, M., Siegmund, N., Sunkle, S., Saake, G., Apel, S. (2009). Data Management for Embedded Systems: A Cell-based Approach, *EDACS '09: 20th International Workshop on Database and Expert Systems Application (DEXA 2009)*, IEEE Computer Society, pp. 9-13.
- [40] Ur Rahman, S. S., Saake, G. (2009). Cellular DBMS: An Attempt Towards Biologically-Inspired Data Management.
- [41] Kersten, M. L. (1998). A Cellular Database System for the 21st Century, *ARTDB 97: In: Proceedings of the Second International Workshop on Active, Real-Time, and Temporal Database Systems*, London, UK: Springer-Verlag, pp. 39-50.
- [42] Kersten, M., Weikum, G., Franklin, M., Keim, D., Buchmann, A., Chaudhuri, S. (2003). A database striptease or how to manage your personal databases, *VLDB 2003: In: Proceedings of the 29th international conference on Very large data bases*, VLDB Endowment, pp. 1043-1044.
- [43] Verhaegh, W., Aarts, E., Korst, J. (2009). *Intelligent Algorithms in Ambient and Biomedical Computing (Philips Research Book Series)*, Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- [44] Verroca, F., Eynard, C., Ghinamo, G., Gentile, G., Arizio, R., D'Andria, M. (1999). A Centralised Cellular Database to Support Network Management Process, *ER '98: In: Proceedings of the Workshops on Data Warehousing and Data Mining*, London, UK: Springer-Verlag, pp. 311-322.
- [45] Kodama, T., Kunii, T. (2002). Development of new DBMS based on the cellular model-from the viewpoint of a data input., *IEIC Technical Report (Institute of Electronics, Information and Communication Engineers)*, v. 102, pp. 97102.
- [46] Kodama, T., Kunii, T., Seki, Y. (2004). A Development of a Cellular DBMS Based on an Incrementally Modular Abstraction Hierarchy, *Joho Shori Gakkai Kenkyu Hokoku*, v. 2004, pp. 43-50.
- [47] Greenwood, P., Blair, L. (2004). Using Dynamic Aspect-Oriented Programming to Implement an Autonomic System, *DAW '04: In: Proceedings of the 2004 Dynamic Aspects Workshop*, pp. 76-88.
- [48] Truyen, E., Joosen, W. (2008). Towards an aspect-oriented architecture for self-adaptive frameworks, *ACP4IS '08: In: Proceedings of the 2008 AOSD workshop on Aspects, components, and patterns for infrastructure software*, New York, NY, USA: ACM, pp. 1-8.
- [49] Von Neumann, J. (1958). *The computer and the brain*, New Haven.
- [50] Von Neumann, J. (1963). *The General and Logical Theory of Automata. John von Neumann--Collected Works, 5. AH Taub*, Macmillan, New York.
- [51] C.S. Calude, C. S., Păun, G. (2000). Computing with cells and atoms in a nutshells, *Complex.*, v. 6, pp. 3848.
- [52] *Applications of Membrane Computing*, (2006). Springer.
- [53] Păun, G. (2002). *Membrane Computing: An Introduction*, Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- [54] Păun, G., Rozenberg, G., Salomaa, A. (2000). Membrane computing with external output, *Fundam. Inf.* v. 41, pp. 313-340.
- [55] Angert, E. R. (2005). Alternatives to binary fission in bacteria, *Nature Reviews Microbiology*, v. 3, pp. 214-224.
- [56] Lightstone, S. S., Lohman, G., Zilio, D. (2002). Toward autonomic computing with DB2 universal database, *SIGMOD Rec.* v. 31, pp. 55-61.
- [57] Chaudhuri, S., Narasayya, V. (2007). Self-tuning database systems: a decade of progress, *VLDB '07: In: Proceedings of*

the 33rd international conference on Very large data bases, VLDB Endowment, pp. 3-14.

[58] Rosenmüller, M., Kästner, C., Siegmund, N., Sunkle, S., Apel, S., Leich, T., Saake, G. (2009). SQL à la Carte --Toward Tailor-made Data Management, *13. GI-Fachtagung Datenbanksysteme für Business, Technologie und Web (BTW)*.

[59] Garcia-Molina, H., Salem, K. (1992). Main Memory Database Systems: An Overview, *IEEE Trans. on Knowl. and Data Eng.*, v. 4, pp. 509-516.

[60] Czarnecki, K., Helsen, S., Eisenecker, U. W. (2004). Staged Configuration Using Feature Models., *SPLC '04: In: Proceedings of the 3rd software product line conference*, Springer, pp. 266-283.

[61] Spinczyk, O., Lohmann, D., Urban, M. (2005). AspectC++: an AOP Extension for C++, *Software Developer's Journal*, pp. 68-74.

[62] Apel, S., Leich, T., Rosenmüller, M., Saake, G. (2005). *FeatureC++: Feature-Oriented and Aspect-Oriented Programming in C++*.

[63] Apel, S., Leich, T., Rosenmüller, M., Saake, G. (2005). *FeatureC++: On the Symbiosis of Feature-Oriented and Aspect-Oriented Programming*, *GPCE '05: In: Proceedings of*

the International Conference on Generative Programming and Component Engineering, Springer, pp. 125-140.

[64] Apel, S., Leich, T., Saake, G. (2008). Aspectual Feature Modules, *IEEE Transactions on Software Engineering*, v. 34, pp. 162-180.

[65] Elson, J., Bien, S., Busek, N., Bychkovskiy, V., Cerpa, A., Ganesan, D., Girod, L., Greenstein, B., Schoellhammer, T., Stathopoulos, T., Estrin, D. (2003). *EmStar: An Environment for Developing Wireless Embedded Systems Software*, CENS Technical Report 9.

[66] Köppen, V., Siegmund, N., Soffner, M., Saake, G. (2009). An Architecture for Interoperability of Embedded Systems and Virtual Reality, *IETE Technical Review*, v. 26, pp. 350-356.

[67] Tesanovic, A., Sheng, K., Hansson, J. (2004). Application-Tailored Database Systems: A Case of Aspects in an Embedded Database, *IDEAS '04: In: Proceedings of the International Database Engineering and Applications Symposium*, Washington, DC, USA: IEEE Computer Society, pp. 291-301.

[68] Sterritt, R., Parashar, M., Tianfield, H., Unland, R. (2005). A concise introduction to autonomic computing, *Advanced Engineering Informatics*, v. 19, pp. 181-187. 10. *Informatics*, v. 19, 2005, pp. 181-187. 10.

Author Biographies

Syed Saif ur Rahman is a member of the Database Group at the Otto-von-Guericke University Magdeburg, Germany. His research interests include DBMS architectures, embedded databases, and biologically-inspired customizable autonomous DBMS development.

Veit Köppen is a member of the Database Group at the Otto-von-Guericke University Magdeburg, Germany. Currently he is the project coordinator in the ViERforES project funded by the German Ministry of Education and Research. His research interests include Business Intelligence, data quality, interoperability aspects of embedded devices, and process management.

Gunter Saake is fulltime professor for the area of "Databases and Information systems" at the Otto-von-Guericke University, Magdeburg. His research interests include database integration, tailor-made data management, object-oriented information systems, and information fusion.