

Mälardalen University Licentiate Thesis
No.9

COMET: A Component-Based Real-Time Database for Vehicle Control-Systems

Dag Nyström

May 2003



MÄLARDALEN UNIVERSITY

Department of Computer Science and Engineering
Mälardalen University
Västerås, Sweden

Copyright © Dag Nyström, 2003
ISBN 91-88834-46-8
Printed by Arkitektkopia, Västerås, Sweden
Distribution: Mälardalen University Press

Abstract

Vehicle control-systems have evolved from small isolated controllers to complex distributed computer-systems. These systems include nodes spanning from simple 8-bit micro-controllers with a minimum of memory to complex 32-bit processors with vast resources. The main motivation for this evolution is the need for increased functionality in motor vehicles. Examples of such functionality include momentary fuel consumption measurements, anti-spin systems, and computerized diagnostics of vehicle-status. The control of the increased functionality requires the handling and maintenance of larger volumes of data, and has created a need for a uniform and efficient way to access and maintain this data. A real-time database management system could satisfy this requirement but an extensive survey of commercial and experimental database management systems has shown that there is currently no database system suitable for vehicle control systems available. In today's systems, data management is performed in an ad-hoc fashion at a low level of abstraction, using internal data-structures, e.g., shared variables and structures. This approach requires that the consistency of the data is maintained by the application, by, for example, resolving data access conflicts through the use of semaphores.

This thesis presents a flexible and configurable database management system designated COMET, suitable for embedded systems and in particular, vehicle control-systems. To be able to handle the varying requirements imposed by different systems, COMET emphasizes configurability and tailorability, by adopting a component-based architecture.

The result of this research is the implementation of COMET BaseLine, which is an instance of COMET suited to a particular vehicle control-system. The required behaviour of this database is based on requirements gathered from a case study performed at Volvo Construction Equipment Components AB in Eskilstuna. To fulfill these requirements, a concept called database pointers has been introduced and implemented. Database pointers provide controlled direct access to individual data elements in the database, efficiently and temporally deterministic, providing at the same time a high level of abstraction.

Acknowledgements

This thesis is presented as a partial fulfillment for the degree of Licentiate in Engineering, at the department of computer science and engineering, Mälardalen university, Sweden. The research presented in this thesis has been performed within the COMET (component-based real-time embedded database) project, jointly executed by Mälardalen university and Linköping university.

I sincerely thank my supervisors Christer Norström from Torshälla, and Jörgen Hansson for their invaluable support and the feedback they supplied on my work.

I wouldn't be here without the support and cooperation of my colleague and co-researcher, Aleksandra Tesanovic, at Linköping university.

This work has been funded by ARTES (A network for real-time research and graduate Education in Sweden). I would like especially to thank Hans Hansson, programme director of ARTES.

The cooperation of industry has been a vital part of this project and I would like to thank the staff at Volvo Construction Equipment Components AB, in Eskilstuna, Sweden, for the two successful weeks that we spent there. A special thanks to Nils-Erik Bånkestad for his support during the entire project. Furthermore, many thanks to Bengt Gunne at Upright Database Technology, Uppsala Sweden, for sharing his deep knowledge of embedded databases with us.

Also thanks to all my colleagues at the department for making these years so memorable. I would like to direct my special gratitude to the following: Thomas Nolte, my dear friend throughout our time at the university, Harriet Ekwall, for solving all those insolvable problems, Jukka Mäki-Turja, colleague and fellow skydiver, Joel, Daniel, Anders P., Mic, Anders "Fnurkelmannen" W., and Kristian.

Many thanks to Susanne Eriksson, Ingela Hedman, and Gisèle Mwepu for their great work implementing the COMET BaseLine. I wish you the best.

Finally a warm hug to my wife Jenny and my daughter Liv, for supporting me and loving me during these years.

So, buckle up folks, here we go!

Dag Nyström, Västerås in May 2003

Contents

1	Introduction	1
1.1	Motivation	2
1.2	Contribution	3
1.3	The Research Work and Method	3
1.4	Conclusions and Future Work	5
1.5	Thesis Outline	6
2	Paper A: Real-Time and Embedded Databases in Practice	9
2.1	Introduction	11
2.1.1	An example application	14
2.2	Database Systems	18
2.2.1	Traditional Database Systems	18
2.2.2	Embedded Database Systems	18
2.2.3	Commercial Embedded DBMS: a Survey	21
2.2.4	Current State-of-the-art From Research Point of View	44
2.2.5	Real-Time Properties	45
2.2.6	Distribution	49
2.2.7	Transaction Workload Characteristics	53
2.2.8	Active Databases	57
2.2.9	Observations	62
2.3	Summary	63
2.3.1	Conclusions	63
2.3.2	Future Work	63
3	Paper B: Data Management Issues in Vehicle Control Systems: a Case Study	71
3.1	Introduction	73

3.2	The Case Study	74
3.2.1	Rubus	76
3.2.2	VECU	76
3.2.3	IECU	78
3.2.4	Data Management Requirements	79
3.3	Modeling the System to Support a RTDB	82
3.3.1	Data Management Implications	83
3.3.2	DBMS Design Implications	85
3.3.3	Mapping Data Requirements to Existing Database Platforms	87
3.4	Conclusions	88
4	Paper C: Database Pointers: a Predictable Way of Manipulating Hot Data in Hard Real-Time Systems	91
4.1	Introduction	93
4.2	Background and Related Work	94
4.2.1	Relational Query Processing	95
4.2.2	Tuple Identifiers	96
4.2.3	Related Work	96
4.3	Database Pointers	97
4.3.1	The DBPointer Data Type	98
4.3.2	The Database Pointer Table	98
4.3.3	The Database Pointer Interface	99
4.3.4	The Database Pointer Flag	100
4.3.5	Application Example	101
4.4	Concept Evaluation	103
4.5	Conclusions and Future Work	104
5	Paper D: The COMET BaseLine Database Management System	109
5.1	Introduction	111
5.2	Aspect-Oriented Software Development	112
5.3	ACCORD	114
5.4	RTCOM: Real-Time Component-Model	114
5.5	The Architecture of COMET	116
5.6	The Components in COMET	118
5.6.1	User Interface Component	118
5.6.2	Scheduling Management Component	119
5.6.3	Transaction Management Component	119
5.6.4	Index Management Component	121

5.6.5	Lock Management Component	121
5.6.6	Memory Management Component	122
5.6.7	Checkpointing and Recovery Component	122
5.7	The Aspects in COMET	123
5.7.1	Concurrency-Control Aspect	123
5.7.2	Logging and Recovery Aspect	123
5.8	COMET BaseLine	124
5.8.1	Relational Processing	124
5.8.2	Database Pointer Processing	126
5.8.3	COMET BaseLine Component-Model	126
5.9	COMET BaseLine Architecture	129
5.9.1	Relational User Interface Component	129
5.9.2	Relational Transaction Management Component . . .	130
5.9.3	Database Pointer User Interface Component	131
5.9.4	Database Pointer Transaction Management Component	132
5.9.5	Index Management Component	132
5.9.6	Memory Management Component	133
5.10	Conclusions and Future Work	133

List of Publications

Publications Included in Thesis

Paper A

Aleksandra Tešanović, Dag Nyström, Jörgen Hansson & Christer Norström
Embedded Databases for Embedded Real-Time Systems: A Component-Based Approach
MRTC Report ISSN 1404-3041 ISRN MDH-MRTC-43/2002-1-SE
Mälardalen Real-Time Research Centre, Mälardalen University, January 2002.

This report is a survey that covers the following areas: (i) Real-time databases (ii) Embedded databases (iii) Component-based databases (iv) Component-based software engineering. It is a background study for the field of research for the entire COMET project, i.e., a state of the arts.

The surveys on real-time and embedded databases were written by Dag, and the surveys over component-based databases and component-based software engineering were written by Aleksandra

Paper B

Dag Nyström, Aleksandra Tešanović, Christer Norström, Jörgen Hansson & Nils-Erik Bånkestad
Data Management Issues in Vehicle Control Systems: a Case Study
In proceedings of the 14th Euromicro Conference on Real-Time Systems, pages: 249 -256, IEEE, Vienna, Austria, June 2002.

This paper presents a case-study on data management for a vehicle control-system developed at Volvo Construction Equipment Components AB, Eskilstuna, Sweden. Furthermore, it elaborates on how to redesign the system to incorporate a real-time database management system into it. Finally it discusses how to design a real-time database management system that suits the control-system.

Dag and Aleksandra jointly focussed on the data management issues in the current system. Dag also investigated the database requirements and its integration with the control-system. Nils-Erik provided valuable knowledge of the application domain.

Paper C

Dag Nyström, Aleksandra Tešanović, Christer Norström & Jörgen Hansson
Database Pointers: a Predictable Way of Manipulating Hot Data in Hard Real-Time Systems

In proceedings of the 9th International Conference on Real-Time and Embedded Computing systems and Applications, pages: 623 -634, Tainan, Taiwan, February 2003.

This paper proposes a new concept of interfacing a database. This concept, which resemble of pointer operations, writes data directly to the memory location inside the database where the data is stored, thus the index lookup routine is bypassed. This concept is beneficial for real-time control systems.

Responsible for the general ideas and main author and driving force of the paper during the writing process was Dag. Aleksandra aided in the writing process, as well as giving feedback and improvements on the general ideas.

Paper D

Dag Nyström, Aleksandra Tešanović, Christer Norström & Jörgen Hansson
The COMET BaseLine Database Management System

MRTC Report 98/2003, ISSN 1404-3041 ISRN MDH-MRTC-98/2003-1-SE
Mälardalen Real-Time Research Centre, Mälardalen University, April 2003.

This report presents the general ideas of the experimental component-based real-time database management system COMET. Furthermore, it describes the COMET BaseLine which is an instance of COMET suited for a particular vehicular control-system. Finally, it briefly explains the design method and component-model used during the development on COMET.

All authors are responsible for different parts of the main ideas presented in this report. The driving author of this report was Dag.

Publications by the Author, Not Included in Thesis

Aleksandra Tešanović, Dag Nyström, Jörgen Hansson & Christer Norström
Integrating Symbolic Worst-Case Execution Time Analysis with Aspect-Oriented System Development

In proceedings of OOPSLA 2002 Workshop on Tools for Aspect-Oriented Software Development, Seattle, USA, November 2002.

This workshop paper proposes a method to calculate worst case execution time for aspect oriented software.

Aleksandra Tešanović, Dag Nyström, Jörgen Hansson & Christer Norström
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, pages: 279 -298, Tainan, Taiwan, February 2003.

This conference paper presents the COMET component model which can be used in combination with aspect oriented programming. The paper also presents an idea how to implement the COMET DBMS using these components

Aleksandra Tešanović, Dag Nyström, Jörgen Hansson & Christer Norström
Aspect-Level Worst-Case Execution Time Analysis of Real-Time Systems Composed Using Aspects and Components

Proceedings of the 27th IFAC/IFIP/IEEE Workshop on Real-Time Programming, Poland, Elsevier, May 2003. To appear.

This workshop paper proposes a method to calculate worst case execution time for aspect oriented real-time software.

Chapter 1

Introduction

Most functionality in a modern car is in one way or another controlled by a computer. For example, the speedometer on the dash board may, instead of being connected to the engine by a rotating wire, be connected to an onboard computer feeding it with the current speed of the vehicle.

The advantages of computer control of functionality are many. First of all, it reduces the amount of electrical wiring in the car. Instead of having an individual wire for each subsystem, different subsystems can share a single digital cable, called a data bus. Secondly, by supplying a central control-system with all vehicle data, such as speed, engine temperature, fluid levels, etc, this information is made available to different subsystems. The vehicle speed, for example, is supplied to the speedometer, but it can also be an important input to the anti-lock braking system (ABS). A third reason is that new functionality can be easily implemented in the car, requiring only computer software for its control. Examples of new functionality are, anti-skid systems to keep the vehicle steerable during emergency braking, air bags, and on-the-fly diagnostics and statistics about the conditions in different subsystems in the vehicle.

An investigation made by Volvo suggests that the use of computer systems in vehicle control is increasing, and is expected to continue to increase, by 7-10% annually [1]. Adding functionality to the car increases the volume of data to be handled by the control-system. This data requires more and more complex maintenance.

To further complicate matters, vehicular computer-systems are often distributed to several loosely coupled nodes, called electronic control units (ECUs), throughout the vehicle. These nodes vary, both in computational power and

the amount of resources available. Some nodes may be 8-bit micro-controllers with a small on-board memory while other nodes may be powerful processors with vast resources. These nodes need to be connected in order to share common information, i.e., the data distributed among the nodes.

In today's systems, data is handled in an ad hoc fashion, using internal data structures, i.e., shared variables. This means that the data must be kept consistent by the application. There are many types of consistency to handle, such as (i) concurrency consistency, i.e., data accesses must be serialized, and this is enforced using mechanisms such as semaphores. Further, (ii) data in real-time systems has a limited temporal validity, both absolute and relative [2]. Finally, (iii) data must be kept consistent among nodes in the distributed system. It is essential to consider all these issues while designing, maintaining and extending the application and to be aware of the increasing risk of software errors.

The current data management approach is becoming increasingly insufficient as systems become increasingly complex and a need for data management on a higher level of abstraction has emerged. One way of providing such a data management would be to incorporate a real-time database management system (DBMS) into the system. DBMSs provide a generic and uniform way of accessing data. Furthermore, they can automatically enforce consistency mechanisms, including temporal validity control and distribution control.

1.1 Motivation

A number of database management systems intended for embedded systems are currently available on the market. These include Polyhedra [3], Pervasive.SQL [4], Berkeley DB [5], and TimesTen [6]. Some of these systems are small enough to fit in an environment such as a vehicular control system. None, however, incorporates any real-time or deterministic behaviour. This means that they are not behaviourally analyzable, which is one requirement for deployment in a vehicular system.

On the other hand, there are a number of database management systems that do incorporate real-time algorithms, such as DeeDS [7], RODAIN [8], STRIP [9], and BeeHIVE [10]. Some of these systems, such as the DeeDS system, are intended for hard real-time systems, and might therefore be predictable and analyzable enough to be deployed in such a system. However, none is intended for embedded environments and none could, in its current form, fit into a vehicle control system.

Further, the varying requirements of data management in different vehicle control systems, require different behaviours of the database management system used. In other words, a DBMS suitable for use in one vehicle might not be suitable for use in a different vehicle. In fact, nodes within the same vehicle might be so heterogeneous that a DBMS for one type of node might not fit another.

One solution to this would be to develop an “in house” database manager to suit one particular vehicular system. However implementing, and even more importantly validating, a database management system for a vehicle control-system is not a trivial task, and thus this solution is not time-efficient.

Our approach is to develop a tailorable database management system using a component-based approach. This approach would enable the database to be configurable to suit a number of different types of vehicular control-systems.

1.2 Contribution

In this thesis we will identify the data management requirements of a specific industrial vehicle control-system. We show how the control-system can be redesigned to incorporate a real-time database management system. Further, we have designed and implemented a real-time database suited to this control-system. To permit this implementation, we have introduced a new data access concept for database management systems called database pointers.

The database, named COMET BaseLine, is an instance of the COMET DBMS, a component-based database management system, suited to embedded systems, in particular vehicle control-systems.

1.3 The Research Work and Method

One central aspect of this research has been close interaction with industry. This has enabled us to work on solving research questions relevant to industrial systems in practice. To achieve an understanding of both the academic problems surrounding real-time data management, and data management requirements from vehicle control-systems, extensive background work has been performed.

The research began by investigating the current state-of-the-art in real-time and embedded database management systems [11]. We performed an extensive survey of commercially available embedded DBMSs, investigating a handful of systems on the basis of a number of criteria. In addition to this survey, a second

survey was performed, this time, of experimental real-time DBMSs, developed in academia. The latter survey also provided us with documentation of basic database management systems theory, which was simultaneously investigated.

Equipped with this new knowledge, an industrial stay was arranged. During a period of two weeks, we visited Volvo Construction Equipment Components AB (Volvo CE), in Eskilstuna, Sweden. Volvo CE has developed a vehicular control-system used in construction vehicles such as wheel loaders, articulated haulers, and excavators. In this case study [12], we investigated the data management currently used in these systems, and discussed the possibilities of redesigning the system to incorporate a real-time DBMS. Furthermore, we considered how such a DBMS should be designed to fulfil the requirements of the control system.

Reviewing the earlier survey, we found no existing DBMS suitable for a system such as this. The commercial DBMS systems were, in some cases, small enough to fit in the nodes of the control-system but they had no hard real-time database properties and were therefore insufficiently reliable for integration with such a system. Some real-time DBMS systems on the other hand, have mechanisms which support a predictable behaviour but these are not intended for use in resource-constrained embedded systems.

This realization provided further motivation for the development of a tailorable DBMS that could be configured to suit a particular system. Initial work on a suitable architecture for COMET was begun. A component-based approach was used to enable parts of the database to be easily exchanged, to permit the necessary functionality to be plugged in and unnecessary functionality to be replaced. One problem, however, was that some functionality cross-cut several components, inhibiting the modularity we aimed at. To solve this problem, an aspect-oriented approach was used. Aspect-oriented software development (AOSD) enables certain functional properties (aspects) to be extracted and treated as separate artifacts [13]. These aspects could then be woven into the system at compile-time. An architectural framework for COMET was developed, this consisting of a set of components and a set of aspects.

The next step was to implement an instance of COMET suitable for the system investigated in the case study at Volvo CE. One problem still to be solved was how to allow the control-tasks to access the database in a predictable way without too much overhead. A database access method, called database pointers, was developed [14]. Database pointers are used to access individual data elements in the same way as shared variables are accessed but in this case the data elements reside inside a database controlled by a DBMS. To ensure database consistency, these accesses must be performed in a controlled way.

Finally the implementation of the first instance of COMET could begin. This instance, called COMET BaseLine, implements a main-memory DBMS. This DBMS can be accessed using the relational interface, which provides a subset of the database query language SQL. Furthermore, COMET can be accessed using database pointers. The database, which is suited to the control-system investigated at Volvo CE, is now fully functional and has a memory footprint of around 20kb.

1.4 Conclusions and Future Work

In this thesis, we have presented a real-time database management system, called COMET, suited to embedded real-time systems, in particular vehicular systems. COMET has a component-based architecture that supports aspects. This enables functionality to be tailored in many different ways, such as using parameterization, defining aspects, and replacing components. This is important because the requirements of different embedded applications vary. A first instance of COMET, called COMET BaseLine has been presented. This instance is specifically suited to a vehicular system developed at Volvo Construction Equipment Components AB (Volvo CE) in Eskilstuna, Sweden. The case study of the data management requirements of this system has been presented in the thesis. A concept called database pointers has been introduced to permit the hard real-time controlling system to efficiently access the database. This concept has been successfully implemented in COMET BaseLine. This thesis leaves certain questions unanswered and therefore much work in this area remains to be performed. Examples of incomplete or pending research include:

- Integration, verification and testing of COMET BaseLine together with the Volvo CE control-system.
- Development of a concurrency control algorithm suited to database pointers. Ideally, hard database pointer transactions would be allowed to run without blocking or, at least, with bounded blocking times.
- Verification of the COMET architecture for other types of control-systems to determine the flexibility and configurability of COMET.
- Further investigation of the benefit of using aspects to tailor the behavior of a COMET instance.

- Development of configuration and analysis tools for COMET, to aid the developer of a control-system in choosing suitable components for COMET.

1.5 Thesis Outline

The thesis is outlined as follows: In paper A, we survey commercial embedded DBMSs as well as experimental research DBMSs, using a number of criteria. We continue in paper B with a case study of the data management in a vehicle control-system used in practice. In this paper we also elaborate on how to redesign the system to incorporate a real-time database. The demands on the real-time database are also presented. In paper C we introduce a new concept called database pointers that is used to access individual data elements in a real-time database in an efficient and temporally predictable way. Finally, in paper D we present the concepts of our experimental database management system, COMET. The architecture of COMET is described, as well as the individual components and aspects identified in COMET. Furthermore, the first implemented instance of COMET, the COMET BaseLine, is presented. COMET BaseLine implements database pointers and is tailored to fit the vehicle control-system described in paper B.

Bibliography

- [1] L. Casparsson, A. Rajnak, K. Tindell, and P. Malmberg. Volcano - a Revolution in On-Board Communications. Technical report, Volvo Technology Report, 1998.
- [2] K. Ramamritham. Real-Time Databases. *International Journal of distributed and Parallel Databases*, 1(2):199–226, 1993.
- [3] Polyhedra Plc. <http://www.polyhedra.com>.
- [4] Pervasive Software Inc. <http://www.pervasive.com>.
- [5] Sleepycat Software Inc. <http://www.sleepycat.com>.
- [6] TimesTen Performance Software. <http://www.timesten.com>.
- [7] S. F. Andler, J. Hansson, J. Eriksson, J. Mellin, M. Berndtsson, and B. Efrting. DeeDS Towards a Distributed and Active Real-Time Database System. *ACM SIGMOD Record*, 25(1):38–40, 1996.
- [8] J. Lindstrom, T. Niklander, P. Porkka, and K. Raatikainen. A Distributed Real-Time Main-Memory Database for Telecommunication. In *Proceedings of the Workshop on Databases in Telecommunications*, pages 158–173. Springer, September 1999.
- [9] B. Adelberg, B. Kao, and H. Garcia-Molina. Overview of the STanford Real-time Information Processor (STRIP). *SIGMOD Record*, 25(1):34–37, 1996.
- [10] J. A. Stankovic, S. H. Son, and J. Liebeherr. *Real-Time Databases and Information Systems*, chapter BeeHive: Global Multimedia Database Support for Dependable, Real-Time Applications, pages 409–422. Kluwer Academic Publishers, 1997.

- [11] Aleksandra Tešanović, Dag Nyström, Jörgen Hansson, and Christer Norström. Embedded Databases for Embedded Real-Time Systems: A Component-Based Approach. Technical Report MRTC Report ISSN 1404-3041 ISRN MDH-MRTC-43/2002-1-SE, Dept. of Computer Engineering, Mälardalen University, January 2002.
- [12] Dag Nyström, Aleksandra Tešanović, Christer Norström, Jörgen Hansson, and Nils-Erik Bänkestad. Data Management Issues in Vehicle Control Systems: a Case Study. In *Proceedings of the 14th Euromicro Conference on Real-Time Systems*, pages 249–256. IEEE Computer Society, June 2002.
- [13] Aleksandra Tešanović, Dag Nyström, Jörgen Hansson, and Christer Norström. 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*, pages 278–298, February 2003.
- [14] Dag Nyström, Aleksandra Tešanović, Christer Norström, and Jörgen Hansson. Database Pointers: a Predictable Way of Manipulating Hot Data in Hard Real-Time Systems. In *Proceedings of the 9th International Conference on Real-Time and Embedded Computing Systems and Applications*, pages 623–634, February 2003.