ARTES project 0005-18
2000-05-23

Pre-Implementation Analysis of Distributed Control Systems - PICADOR

Dr. Martin Törngren, Tekn. Lic. Ola Redell, Prof. Jan Wikander (martin, ola, jan@md.kth.se),
Division of Mechatronics (DAMEK), Department of Machine Design,
Royal Institute of Technology, SE-100 44 Stockholm, Sweden

Summary

The PICADOR project is closely related to the AIDA project at DAMEK, KTH. While the AIDA project has a wide frontier towards implementation of distributed, embedded real-time control systems, the research in PICADOR will focus specifically on static pre-implementation timing analysis. Methods for timing analysis of distributed real-time control implementations will be developed to support control design with tight estimates of timing behaviour such as jitter and control delays. Funding is sought for two years to an amount of 600 KSEK/year. The project will be carried out in cooperation with SAAB Automobile AB.

Problem Statement

The machine industry is faced with an increasing amount of functionality that is being implemented in software where the infrastructure is based on embedded distributed computer systems. In such systems, the networks, forming a backbone for information exchange, constitute one important enabler for the introduction of new functionality. Examples include diagnostics, sensor or state sharing, new functions and their coordination. This implies a necessary shift from node oriented to functionality oriented design. There are a number of reasons for this including

- the fact that a new function typically uses inputs (states) and outputs (states) related to existing functions and different nodes.
- the need to verify functions through modelling, analysis and synthesis prior to the implementation (it is too late to find bugs during system integration)
- the need to enable solution exploration, alternative implementations and reuse of functionality.

Given the separate development of functions there is a strong need to gradually take implementation related effects into account as the development proceeds. One field where this is common practice is control engineering, where after functional analysis and simulation, implementation effects like actuator saturation and quantization due to limited word length or insufficient computational accuracy can be modelled and included in the functional models. Such an approach provides a step-wise progress towards the implementation.

In a similar fashion, the hypothesis in PICADOR, is that implementation effects caused by limited computational resources and transient errors can and should be modelled at the functional level. Examples of such implementation effects and how they in turn can affect the functions include;

- Jitter, due to scheduling and varying execution times, can for example cause time-varying end-to-end delays resulting in reduced performance and even instability in feedback control. A related effect is that of vacant sampling which can be caused by jitter in conjunction with the communication policies used.
- Transient errors, due to for example temporary processor and/or communication failure, can have a similar effect as jitter, but can also lead to periods of open-loop control. Estimating the length of temporary broken data- and or control-flows is important for designing robust functions and proper error handling.

Concurrently with the functional development, there is a need to

- develop or extend the computer system architecture, based on a synthesis of all functional and non functional requirements and to
- select the mapping of functions (including allocation and scheduling) to the physical components and verifying it.

The verification part of the second case is the one primarily considered in PICADOR.

Main Ideas

The goal of this project is to extract timing information from a potential control system implementation in order to feed back the information to the control design. The main idea is to analyse well defined models of an implementation using algorithms that take scheduling, clock synchronisation, communication principles and task synchronisation through shared resources into account. In the best case the development of analysis algorithms will lead to tightly bounded estimates of jitter and synchronisation skews (between e.g. sampling activities on different nodes in a distributed system). Traditionally control designers have relied on pessimistic worst case estimates and iterative tests for extracting this information. The analysis as defined here will rely on a solid base of models describing the control architecture and its timing requirements together with the target computer architecture, its timing behaviour and the system software.

Typically in today's research on real-time systems and their timing behaviour, the focus is set on response time analysis and design of systems for meeting requirements on response times, deadlines and to some extent jitter. For control implementations however, a very important design parameter is the degree of satisfaction of the requirements of precise periodicity that are posed by e.g. sampling and actuation activities (Törngren, 1995). The idea is to study and develop methods and algorithms for static analysis of timing behaviour, given a description of the underlying hardware and system software (e.g. RTOS). The methods should not only be applicable to systems designed with rate monotonic techniques or similar, but also to more general implementations such as existing in many industrial systems today, with different combinations and hierarchies of schedulers for different resources.

In order to successfully meet the goals of the project, general modelling techniques need to be used that allow unambiguous description of real-time implementations. The models should specifically support the description of systems with multiple resources such as processors, communication buses and memories. Another important feature of the models must be the ability to describe various communication media and the arbitration (scheduling) of competing message transmissions over communication resources that are shared between processors. A modelling framework like the one described here has already, in most parts, been developed within the AIDA project (Redell, 1998). Some further development and refinement of the modelling framework will however be necessary to fully support the envisioned methods for analysis. These refinements include e.g. detailed descriptions of general communication resources and a general way to specify schedulers and hierarchies of schedulers.

Expected Results and Impact

As an outcome of the described research a method and algorithms for pre-implementation temporal analysis of control implementations will be available. The methodology will assist a control system designer in refining the control system to account for timing properties of the real-time implementation. This means that it will become simpler to design stable high performance control systems with standard real-time implementation techniques. And in the end it may be possible to prove (to some level of accuracy) that a real-time implementation of a control system is feasible and that it satisfies the requirements posed by the control engineer. Analysis such as the one described here will also help in speeding up the development process by limiting time consuming development iterations and avoiding late discovery of timing related design errors.

A case study will be used to demonstrate the use of the proposed modelling technique and the result of the analysis. The developed algorithms will be implemented in the AIDA research tool-set that will be developed as part of the AIDA, DICOSMOS and PICADOR projects. The tool-set, that will assist in taking the step from a control design to a real-time implementation in a distributed architecture, has been planned for some time and is described in more detail in (Redell, 1998). The specification of the modelling framework in AIDA was the first step towards and the basis for such a tool-set.

Project Plan

June 2000 - November 2000. Development of a basic analysis method, handling a subset of the intended target systems. Updating of the AIDA models, with focus on the resource models.
August 2000 - December 2000. First implementation of the AIDA tool-set and inclusion of proposed analysis method.
November 2000 - November 2001. Development of a more generalized analysis method spanning a larger set of the systems that may be modelled with the AIDA framework. Further work will be done on the refinement of the AIDA models and on the design and implementation of the AIDA tool-set.
June 2000 - April 2001. Case study performed in cooperation with some industrial partner.
August 2000 - December 2001. Ph.D. level courses of at least 15 credits.
February 2002. PhD Dissertation.

Preliminary budget

Funding is sought for one PhD student for two years. The student, Ola Redell, received his licentiate thesis in 1998, and recently returned to complete his PhD after two years of industrial consultancy work. He will therefore work close to 100% in this research project aiming to complete his PhD within two years. Therefore, 600KSEK per year is sought from ARTES. Additional funding for mobility with the industrial partner is sought from ARTES separately.

Related Research

Related research on static analysis of timing behaviour has been done by e.g. (Harbour et al., 1994 and Sun et al., 1997). Harbour et al. developed a method for bounding response times of tasks on a uniprocessor that where scheduled using a fixed priority scheduler. The tasks were comprised of a number of sequential sub-tasks that each could have its own priority. Sun et al. developed similar methods on the same theme however with a bit better performance (tighter bounds). Other work on timing analysis include the development of the RMA (Rate Monotonic Analysis) methodologies during the last decade started by Liu and Layland in 1973 and summarized by Klein et al. (1993). This research of course gives a good and very powerful base for timing analysis but has some limitations. First, the interest here is analysis in general for a large set of (potentially combined) scheduling techniques and hence not constrained to fixed priority scheduling. Second, RMA in general generates quite pessimistic bounds that are not completely satisfactory for control design. Third, RMA is just like most other analysis methods focused on bounding of response times using worst case analysis and the support for analysis of jitter and synchronised events is not as well developed. Furthermore we are interested in the behaviour and performance in a distributed system that includes combinations of multiple schedulers on both processors and communication resources. This is not well covered by the RMA theory.

Another area of real-time research that is related to the one described here is the end-to-end deadline analysis for distributed systems. But also in this case the main focus has been on response time analysis and schedulability in general and not in the meeting of requirements on jitter and synchronisation. Work on scheduling with end-to-end deadline requirements include e.g. (Bettati and Liu, 1992, Saksena and Hong, 1996).

A very interesting effort on modelling and tool-based analysis is the SEW tool-set and the functionality included therein (Chatterjee et al., 1997). Another related research tool-set for timing analysis is PERTS that has now been commercialised under the name Rapid RMA (Liu et al., 1993).

Relation to the profile, as defined in this CFP

The PICADOR project is strongly related to the specification and design of heterogeneous real-time systems, early analysis of architectural choices, grow and change aspects, resource handling, distributed systems, implementation of control systems, and methodological implications.

Industrial relevance

See the problem statement and expected results and impact.

Relation to other SSF programmes.

DAMEK is a partner of the SSF funded Centre for Autonomous Systems (http://www.cas.kth.se) where DAMEK is responsible for the Difficult Terrain Demonstrator, the WARP walking robot. The problems studied in this project are highly relevant for a large class of embedded control systems including autonomous robots.

Context

The research group. The project is managed by a project management group including Associate Prof. Martin Törngren (project leader), Prof. Jan Wikander, and Kennet H. Lind (Manager-Electrical Architecture Technical Development, Saab Automobile AB).

Complementary activities and funding. Other related projects carried out at the division of Mechatronics in the area of embedded real-time systems are (see http://www.md.kth.se/~martin/ for www links):

- Interdisciplinary techniques and methods for design of distributed control systems - the DICOSMOS project (NUTEK funded). In this project timing issues are studied from a control perspective, rather than from the real-time analysis perspective as is the case in PICADOR. Hence the research in PICADOR will complement the one in DICOSMOS.
- Models and tools - the AIDA project (funded by ARTES, NUTEK and BGM - Branchgruppen Mekatronik)
- Dependability in terms of functional interference and integration - the newly started FINE project (ARTES and CAS funded)
- The MARCH project - Mechatronics software architecture (ARTES).

These projects together bring a critical mass enabling for example implementation of a common modelling framework and tool prototype.

Research cooperation. Within the above mentioned projects, cooperation exists with computer engineering at CTH, control engineering at LTH and computer engineering at MDH. It is also the intention to increase the direct international cooperation.

Industrial cooperation. Kennet Lind, kenneth.h.lind@saab.com, tel. no: +46-(0)520-483464, Saab Automobile AB, A1-3 TVEBA, SE-461 80 Trollhättan.

Appendices

Short CV for Martin Törngren and supporting letter from SAAB Automobile AB.

References

Bettati, R. and Liu, J.W.-S., (1992), End-to-End Scheduling to Meet Deadlines in Distributed Systems, Distributed Computing Systems, Proceedings, 1992 , Page(s): 452 -459

Chatterjee, S., et al., (1997), SEW: a toolset for design and analysis of distributed real-time systems, Third IEEE Real-Time Technology and Applications Symposium, 1997

Harbour, M.G. et al., (1994), Timing Analysis for Fixed-Priority Scheduling of Hard Real-Time Systems, Software Engineering, IEEE Transactions on, Vol. 20 1, Jan. 1994, Pages 13-28

Klein, M.H. et al., (1993), A Practitioners Handbook for Real-Time Analysis: Guide to Rate Monotonic Analysis for Real-Time Systems, Kluwer Academic Publishers, ISBN 0-7923-9361-9, 1993

Liu, C.L. and Layland J.W., (1973), Scheduling Algorithms for Multiprogramming in Hard-Real-Time Environment, Journal of the Association of Computing Machinery 20(1):46-61, pp. 174-189

Liu, J.W.S. et al., (1993), PERTS: A prototyping environment for real-time systems, Real-Time Systems Symposium, 1993., Proceedings. , 1993 , Page(s): 184 -188

Redell, O., (1998), Modelling of Distributed Real-Time Control Systems An Approach for Design and Early Analysis, Licentiate Thesis, Dept. of Machine Design, KTH, TRITA-MMK 1998:9, ISSN 1400-1179, ISRN KTH/MMK-98/9-SE

Saksena, M. and Hong, S., (1996), Resource Conscious Design of Distributed Real-Time Systems: An End-to-End Approach, Conference on Engineering Complex Computer Systems, IEEE , pp.306-313, 1996

Sun, J. et al., (1997) Bounding Completion Times of Jobs with Arbitrary Release Times, Variable Execution Times and Resource Sharing, Software Engineering, IEEE Transactions on, Volume: 23 10 , Oct. 1997 , Page(s): 603 -615

Törngren M., (1995), Modelling and Design of Distributed Real-Time Control Applications, Doctoral thesis, Department for Machine Design, Royal Inst. of Technology, Stockholm, TRITA-MMK 1995:7, ISSN 1400-1179, ISRN KTH/MMK--95/7--SE