Project report 1999-08-02 for ARTES project 9811-3.

Identification of Complexity-Reduction Techniques for Optimal Scheduling in Embedded Distributed Real-Time Systems

Jan Jonsson

Department of Computer Engineering
Chalmers University of Technology
S–412 96 Göteborg, Sweden
janjo@ce.chalmers.se

Mikael Strömberg

Mecel AB
Box 140 44
S–400 20 Göteborg, Sweden
mikael.stromberg@mecel.se

August 2, 1999

Summary

The application of optimal search strategies to scheduling for distributed real-time systems is, in general, plagued by an inherent computational complexity. This has effectively prevented the integration of search strategies such as branch-and-bound [1] in scheduling frameworks and tools used in practice today. The integration of an optimal scheduling strategy could lead to, for example, higher application schedulability, better utilization of scarce resources, or higher system reliability, than can be attained for a sub-optimal scheduling technique. Inspired by recent research results, this ARTES project aims at demonstrating that optimal scheduling is, in fact, a viable alternative for many real-time scheduling scenarios. Our approach is based on the hypothesis that, with detailed knowledge about the real-time application and its characteristics, it is possible to make intelligent choices in the configuration of the search algorithm in such a way that the time it takes to generate an optimal assignment of tasks to system resources is reduced to a tractable value.

The project is a cooperation between two ARTES nodes: Department of Computer Engineering, Chalmers, Göteborg and Mecel AB, Göteborg, and one external node: Real-Time Computing Laboratory, University of Michigan, Ann Arbor. One senior researcher and one Licentiate/Ph.D. student from Chalmers will be actively working on the project. In addition, a scientist at Mecel AB will contribute to the project with expertise in design techniques for distributed vehicular real-time applications and by providing examples and case studies.

1 Problem Statement

In recent decades, distributed real-time systems have emerged in many application domains including drive-by-wire automotive systems, intelligent traffic control systems, banking systems, and autonomous robots. Depending on application domain, one may strive for an optimal or a sub-optimal solution to the task assignment and scheduling problem. For embedded, safety-critical real-time systems, where schedules must be generated *a priori*, optimal scheduling is an attractive choice because it leads to (depending on scheduling objective) higher application schedulability, better utilization of scarce resources, or higher reliability, than does a sub-optimal scheduling technique. For example, a high utilization of the system hardware is extremely important in high-volume real-time systems (*e.g.*, vehicles) where the hardware cost is a dominating factor in the total product cost. In fact, one can foresee increasingly higher demands on the computing services that will be offered by future high-volume embedded systems. However, the requirements on such systems often result in conflicting design goals. In order to deliver high performance, it is essential to use as powerful hardware as possible; on the other hand, hardware cost and power dissipation must be kept at a minimum to allow for a large sales market and high mobility. This kind of multifaceted optimization criteria calls for effective scheduling techniques that can keep resource utilization at a maximum while still meeting other optimization criteria. Unfortunately, the state-of-the-art in industry has not reached this level of maturity, but often rely on strategies based primarily on intuition and tradition. For future distributed systems with more tasks and resources to handle, human intuition no longer suffices to find the most cost-effective solution for a given optimization criterion. It is, thus, necessary to use automated tools with intelligent heuristics to find good design solutions.

Among the optimal task assignment and scheduling techniques, the branch-and-bound (B&B) strategy [1] has been successfully used within different application fields, such as real-time control systems [2], digital signal processing systems [3], and fault-tolerant systems [4]. Unfortunately, scheduling of real-time applications with the objective of finding an optimal solution is, in the general case, hampered by an inherent computational complexity[1]. This has effectively prevented an adaption of optimal strategies in scheduling frameworks and tools used in practice today. To investigate whether a practical use of optimal search strategies, such as the B&B algorithm, indeed exists, it is important to assess the complexity of these strategies for real-world applications. Recent research [5, 6, 7] has indicated the viability of finding techniques that can offer significant complexity reduction[2] under the assumption of certain optimization criteria. While these techniques duly account for general application characteristics such as task execution times, task deadlines or the number of processors used in the system, they are not cognizant of constraints that represent other important characteristics of the application and its environment. Examples of such constraints are those that stem from demands on high system reliability (task replication constraints) or a low total system cost (minimize the number of communication links).

This project addresses the problem of scheduling embedded distributed real-time systems using algorithms that are cognizant of the application and its operating environment.

2 Main Ideas

Our approach is based on the hypothesis that, with detailed knowledge about the real-time application and its environment, it is possible to make intelligent choices in the configuration of the B&B algorithm in such a way that the time it takes to find an optimal schedule is reduced to a tractable value. Examples of

[1] In the case of the B&B algorithm, the search for an optimal solution is organized by means of a *search tree*, where the vertices in the tree represents partial and complete solutions to the problem. To measure the computational complexity of the algorithm, we use the *time complexity* of the B&B algorithm, that is, the number of vertices explored during the search for an optimal solution.

[2] Since the addressed search problem is known to be NP-hard, the intent of the project is to find techniques for complexity reduction *on the average*, using heuristics for embedded distributed systems that make optimal assignment techniques practically useful for a large fraction of the applications considered. In the case that the computational complexity for a specific application should be intractable, the designer can always resort to near-optimal heuristics with acceptable complexity. The development of such techniques is also part of the project. In particular, one of our goals is to find rules of thumb that can aid in choosing the correct technique for each specific application.

such configuration choices are those that determine the order in which search-tree vertices are traversed [5], or those that identify and avoid redundant branches in the search tree [6]. Other configurations deliver tractable performance by making explicit assumptions about the underlying task-scheduling policy [2].

Unfortunately, the techniques proposed to date in the open research literature all suffer from the same major drawback, namely that they assume a general application model with very limited semantics for describing specialized application constraints typically found in embedded computer systems. In such systems, the application and its behavior is well-known which allows for a large degree of fine-tuning as regards the scheduling algorithm used. Furthermore, the design optimization criteria used for such systems differ relatively much from that of general computer systems. Some examples of specific assumptions that can be made for an embedded system are listed below:

Locality constraints: Many functions in an embedded system encompass handling external sensors and actuators. To avoid unnecessary communication network traffic, the tasks implementing these functions need to be located close to the corresponding sensors or actuators.

Vendor-specific constraints: In many embedded systems, certain tasks are handled by computer nodes delivered by an external vendor. In these cases, the designer may not have many options in the assignment of task periods and deadlines or in the choice of task dispatching policy.

The salient feature of the B&B strategy is its ability to avoid inferior and redundant branches in the search tree by means of intelligent heuristics. One of the more powerful heuristics in this context is the *lower-bound function* which calculates optimistic predictions on the quality of all complete solutions emanating from a branch in the search tree. If the predicted quality of a specific branch is inferior to that of the best solution found so far, no solution in that branch will be optimal and, hence, the branch does not have to be explored at all. Much research has been done in finding powerful lower-bound functions and a plethora of such functions exist for various optimization criteria. A somewhat overlooked technique for reducing the B&B search complexity is to use *dominance relations* [8, 9], that is, rules for determining whether one partial solution is inferior to another. With the aid of dominance relations, it is possible to avoid the exploration of even more branches in the search tree without any loss of optimality. The usefulness of the dominance concept in general real-time computing has been demonstrated in, for example [2]. Surprisingly enough, however, no studies have been made to evaluate and compare the actual impact of different dominance relations on the complexity of the B&B algorithm for realistic applications. This is what we intend to pursue in this project.

This project aims at discovering techniques (lower-bound functions and dominance relations) that will significantly reduce the search complexity when detailed information is given about the application and its operating environment.

3 Expected results and impact

The expected research contributions of this project are as follows:

C1. Support for multiple concurrent optimization criteria in the B&B framework. Current work only considers a single optimization criteria.

C2. Heuristic lower-bound functions and dominance relations tailored for embedded distributed real-time systems. Current work has focused on finding techniques that are generally applicable but that are not cognizant of application-specific constraints.

By applying our techniques to applications provided to us from Mecel, the research will demonstrate practical techniques to achieve high resource utilization and meet other design criteria for interesting test cases. Since embedded distributed real-time systems is a rapidly-emerging application domain, the research is expected to generate new design-support tools for these applications. Furthermore, the task assignment problem is solved predominantly "by hand" in industry today. With the aid of a well-functioning automated framework — such as the B&B algorithm — the task-assignment process will be less error prone and will lead to more cost-effective[3] design solutions.

4 Project plan

Experimental plan. The first phase of the project will be concerned with the analysis of distributed real-time applications and the identification of their optimization criteria. This investigation is carried out by the project members at Chalmers in cooperation with application designers at Mecel. The investigation will be presented in a technical paper intended for publication at an international conference. A preliminary title of the paper is "Timing, Physical and Economical Constraints in Embedded Real-Time Systems: Where Do They Come From and Where Do They Go?".

The second phase is concerned with developing techniques within the B&B framework that will aid in reducing the complexity associated with finding optimal schedules. The impact of the new techniques will be evaluated using real applications as well as randomly-generated applications. Initially, an evaluation platform, called **GAST** [10, 11], and developed at Chalmers will be an important experimental resource for this part of the project. One of the salient features of **GAST** is an implementation of a parametrized B&B framework [12, 5] that can be used as a test-bed for implementing new techniques for complexity reduction and evaluating their impact on scheduling performance. Later on in the project, our B&B framework will be based on the SICStus Prolog [13] interpreter from the Swedish Institute of Computer Science. SICStus Prolog has been found suitable for our purposes as it includes a constraint solver for finite domains which offers some predefined constraints specific for scheduling problems. A user interface to the framework will be developed in cooperation with Mecel AB to enable importing and exporting data between SICStus Prolog and Mecel's **Butler** tool [14]. By means of this user interface, it will be possible to get example data from Mecel as well as vizualize schedules made by our B&B framework.

To demonstrate the applicability of the techniques expected as a result from this projects, we will assume realistic real-time architectures. Our main approach for modeling the real-time applications will follow the guidelines of **BASEMENT** [15], a distributed architecture for real-time control applications resulting from a joint work by Mecel AB and Uppsala University. Another target systems used in the project will be **EMERALDS** [16], a real-time kernel for embedded applications. **EMERALDS** has been developed at University of Michigan and is available as a software product that can be run on many existing hardware platforms, for example, the Digital Alpha or the Motorola 68340 microprocessor.

[3]By the term "cost effective", we primarily mean resource effective in terms of, *e.g.*, hardware, power consumption or weight. However, we also use this term refer to shorter turn-around times for correcting erroneous design choices.

Milestones and deliverables. The Ph.D. candidate (Cecilia Ekelin) began working in the project in May 1999. Therefore, the milestones and deliverables apply for the period 990501–040501. This encompasses 2×2 years of research and one year of teaching.

Activity 9905–0001: The first phase of the project. Identification of optimization criteria used in the design of distributed vehicular real-time applications.

Deliverable 0001: State-of-the-art report on B&B algorithms and their application to industrial real-time design constraints.

Activity 0001–0005: The second phase of the project. Development of a B&B framework based on SICStus Prolog. Development of interfaces between the framework and **Butler**.

Deliverable 0005: Demonstration of the B&B framework and its interaction with **Butler**. Refined specifications of the Licentiate project topics and project plans.

Activity 0005–0109: Development of complexity-reduction techniques in the B&B framework.

Deliverable 0109: Licentiate thesis. Specification of the activities in the remaining time period.

Activity 0109–0309: Development of complexity-reduction techniques in the B&B framework.

Deliverable 0309: Integration of developed techniques in academic evaluation tools and industrial design tools. Detailed evaluation of the performance of the developed techniques.

Activity 0309–0405: Writing of Ph.D. thesis.

Deliverable 0405: Ph.D. thesis.

5 Industrial relevance

Optimal scheduling in distributed real-time systems has mainly been reserved for academic purposes. With the introduction of multiple processors in high-volume embedded systems, the scheduling of scarce computing resources can no longer be allowed to be in the hands of an *ad hoc* designer. The techniques developed in the project are expected to be useful for a quite broad range of embedded applications. The general results of the project are expected to be useful, not only for Mecel AB, but also for other industries concerned with the design of embedded real-time systems, such as Ericsson Microwave Systems or Volvo Technical Development.

6 Cooperations and mobility

Similar to some other ARTES projects this project addresses the issues of effective real-time system design. Therefore, we expect the output from this project to be useful in the other related ARTES projects. In particular, we anticipate a close cooperation and synergy effects with the ARTES project "Flexible Reliable Timing Constraints" which is led by Dr. Gerhard Fohler at Mälardalen University. For this reason, we anticipate frequent visit to Dr. Fohler's research group.

Professor Shin's research group at the University of Michigan is one of the few that has made significant work on the B&B algorithm for real-time scheduling purposes. We, therefore, plan to invite Professor Shin or some of his colleagues to Chalmers for scientific input in the project. We also expect to let the Ph.D. candidate visit Professor Shin's group (most likely after the candidate has achieved the Licentiate degree).

7 Context

This research will be carried out in the High-Performance Computer Architecture Group at Chalmers which is led by Professor Per Stenström and consists of three senior researchers (Per Stenström, Fredrik Dahlgren, and Jan Jonsson) with a very strong scientific background. Per Stenström has conducted research in multiprocessor technology since 1984 and his work within memory technology and performance issues in multiprocessor architecture has been internationally recognized. Fredrik Dahlgren is also widely recognized for his contributions in multiprocessor architecture. Jan Jonsson has made significant contributions in the field of design and scheduling of parallel and distributed real-time systems.

Mecel AB has a group consisting of around 10 persons working with the design of distributed real-time systems in vehicular applications. The group is led by Mikael Strömberg.

The Real-Time Computing Laboratory at the University of Michigan is led by Professor Kang G. Shin and consists of seven senior researchers with an outstanding scientific background. Two of the researchers at RTCL (Kang G. Shin and Farnam Jahanian) have conducted internationally recognized research in distributed real-time and fault-tolerant computing systems since 1985. In addition, Kang G. Shin has made significant research contributions in the fields of process control, computer architecture, and computer networking.

8 Publications

The following publications have been not been produced as part of the described project but constitute the foundation for our project plan.

1. J. Jonsson, "Effective Complexity Reduction for Optimal Scheduling of Distributed Real-Time Applications," In *Proceedings of the 19th International Conference on Distributed Computing Systems*, May/June 1999.
2. J. Jonsson, "The Impact of Application and Architecture Properties on Real-Time Multiprocessor Scheduling," PhD Thesis, School of Electrical and Computer Engineering, Chalmers University of Technology, September 1997.
3. J. Jonsson and K. G. Shin, "A Parametrized Branch-and-Bound Strategy for Scheduling Precedence-Constrained Tasks on a Multiprocessor System," In *Proceedings of the 26th International Conference on Parallel Processing*, August 1997.

9 Electronic availability

The current status of the project is continuously updated at the following URL:

`http://www.ce.chalmers.se/~case/hpcag/echidna.html`

References

[1] E. L. Lawler and D. E. Wood, "Branch-and-Bound Methods: A Survey," *Operations Research*, vol. 14, no. 4, pp. 699–719, July/Aug. 1966.

[2] D.-T. Peng and K. G. Shin, "Static Allocation of Periodic Tasks with Precedence Constraints in Distributed Real-Time Systems," *Proc. of the IEEE Int'l Conf. on Distributed Computing Systems*, New Port Beach, California, June 1989, pp. 190–198.

[3] K. Konstantinides, R. T. Kaneshiro, and J. R. Tani, "Task Allocation and Scheduling Models for Multiprocessor Digital Signal Processing," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. 38, no. 12, pp. 2151–2161, Dec. 1990.

[4] S. M. Shatz, J.-P. Wang, and M. Goto, "Task Allocation for Maximizing Reliability of Distributed Computer Systems," *IEEE Trans. on Computers*, vol. 41, no. 9, pp. 1156–1168, Sept. 1992.

[5] J. Jonsson and K. G. Shin, "A Parametrized Branch-and-Bound Strategy for Scheduling Precedence-Constrained Tasks on a Multiprocessor System," *Proc. of the Int'l Conf. on Parallel Processing*, Bloomingdale, Illinois, Aug. 11–15, 1997, pp. 158–165.

[6] I. Ahmad and Y.-K. Kwok, "Optimal and Near-Optimal Allocation of Precedence-Constrained Tasks to Parallel Processors: Defying the High Complexity Using Effective Search Techniques," *Proc. of the Int'l Conf. on Parallel Processing*, Minneapolis, Minnesota, Aug. 10–14, 1998, pp. 424–431.

[7] J. Jonsson, "Effective Complexity Reduction for Optimal Scheduling of Distributed Real-Time Applications," *Proc. of the IEEE Int'l Conf. on Distributed Computing Systems*, Austin, Texas, May 31 –June 5, 1999, pp. 360–369.

[8] T. Ibaraki, "The Power of Dominance Relations in Branch-and-Bound Algorithms," *Journal of the Association for Computing Machinery*, vol. 24, no. 2, pp. 264–279, Apr. 1977.

[9] J. Erschler, G. Fontan, C. Merce, and F. Roubellat, "A New Dominance Concept in Scheduling n Jobs on a Single Machine with Ready Times and Due Dates," *Operations Research*, vol. 31, no. 1, pp. 114–127, Jan./Feb. 1983.

[10] J. Jonsson, *The Impact of Application and Architecture Properties on Real-Time Multiprocessor Scheduling*, Ph.D. thesis, School of Electrical and Computer Engineering, Chalmers University of Technology, Göteborg, Sweden, Sept. 1997.

[11] J. Jonsson, "GAST: A Flexible and Extensible Tool for Evaluating Multiprocessor Assignment and Scheduling Techniques," *Proc. of the Int'l Conf. on Parallel Processing*, Minneapolis, Minnesota, Aug. 10–14, 1998, pp. 441–450.

[12] W. H. Kohler and K. Steiglitz, "Enumerative and Iterative Computational Approaches," *Computer and Job-Shop Scheduling Theory*, E. G. Coffman, Jr., Ed., chapter 6, pp. 229–287. Wiley, New York, 1976.

[13] Swedish Institute of Computer Science, "SICStus Prolog," *http://www.sics.se/isl/sicstus.html*.

[14] Mecel AB, "Butler," *http://www.mecel.se/*.

[15] H. A. Hansson, H. W. Lawson, M. Strömberg, and S. Larsson, "BASEMENT: A Distributed Real-Time Architecture for Vehicle Applications," *Real-Time Systems*, vol. 11, no. 3, pp. 223–244, Nov. 1996.

[16] K. M. Zuberi and K. G. Shin, "EMERALDS: A Microkernel for Embedded Real-Time Systems," *Proc. of the IEEE Real-Time Technology and Applications Symposium*, Brookline, Massachusetts, June 10–12, 1996, pp. 241–249.

A Short CV for Jan Jonsson

Affiliation

Chalmers University of Technology, Department of Computer Engineering, S–412 96 Göteborg, Sweden.
Phone: +46 (31) 772 5220, Fax: +46 (31) 772 3663; Email: janjo@ce.chalmers.se

Permanent positions and degrees

- Assistant professor (*swe.* universitetslektor) of real-time systems, Chalmers, since June 1999.
- Assistant professor (*swe.* forskarassistent) of real-time systems, Chalmers, since May 1998.
- Assistant professor of data communications, Halmstad University (Sweden), 1997.
- Ph.D. degree in computer engineering, Chalmers (Sweden), in 1997.
- M.Sc. degree in computer science and engineering, Chalmers (Sweden), in 1992.

Visiting positions

- Visiting scientist, Real-Time Computing Laboratory, University of Michigan, Ann Arbor (USA), August–December 1996.

Main research interests

- Real-time scheduling strategies for multiprocessor architectures
- Real-time data communications

Selected professional activities

- Member of the program committe for the *7th IEEE Workshop on Parallel and Distributed Real-Time Systems*, San Juan, Puerto Rico, April 1999.