

Methods for Integration of Heterogeneous Real-Time Services into High-Performance Networks

Magnus Jonsson and Bertil Svensson

Centre for Computer Systems Architecture, Halmstad University, Box 823, S-301 18 Halmstad, Sweden
email: {magnus.jonsson, bertil.svensson}@cca.hh.se, Phone: +46 35 16 71 00, Fax: +46 35 12 03 48

Summary

There is an industrial need of new high-performance networks with support for the often heterogeneous real-time requirements in emerging (often embedded) applications like multimedia, radar signal processing, and telecommunication applications. This project aims to provide novel methods to satisfy this need. It can, for example, be the support of services for both soft and hard real-time requirements, and with different quality of service demands (e.g., probability of meeting soft deadlines). The integration of services into both existing and novel high-performance networks with special architectural features/support will be considered. To guarantee the industrial relevance of the project we have chosen a partner with great experience of industrial real-time applications. The main research efforts will be carried out at Halmstad University but our industrial partner (Carlstedt Research & Technology), will take active part in the project, e.g., in the choice of relevant case studies and evaluation criteria. Evaluation of commercial exploitation possibilities is specifically addressed in the project plan.

1. Problem statement

Parallel and distributed computing systems become more and more powerful and hence place increasingly higher demands on the networks that interconnect their processors or processing nodes. Many of the applications running on such systems, especially embedded systems applications, have real-time requirements and, with increasing application demands, high-performance networks are the hearts of these systems. A new view of real-time computing is therefore enforced, focusing on methods and features in high-performance real-time networks. Especially, methods of integrating different kinds of real-time services (or quality of service, QoS) in these networks must be developed. The research efforts on such real-time communication systems are very limited today and emerging standards, like Gigabit Ethernet, do not support real-time communication.

Our focus is on SAN (System Area Network), typically incorporated into embedded systems, and LAN (Local Area Network). Application examples are future radar signal processing systems, distributed multimedia systems, clusters of multimedia servers, telecommunication systems, and image processing applications. A typical example is the radar signal processing system described in [1] [2]. We have developed a high-bandwidth network which takes advantage of typical new features in novel fiber-optic network architectures, but *methods and protocols with good integrated support for diverse real-time services and QoS must be developed.*

2. Main ideas

Fiber-optics is a promising design alternative for future interconnection networks. Numerous configurations with different degrees of optics, optoelectronics, and electronics have been proposed. Our hypothesis is that the high bandwidth offered by fiber-optic networks, or possibly by emerging electrical techniques, offers new possibilities when developing methods or protocols for real-time communication. The main reasons are: (i) the high transport rate of packets which makes it possible to have a significant amount of control traffic, and (ii) the short packet-delay for such traffic.

In addition to the increased bandwidth, novel fiber-optic network architectures [3] can bring other useful features. This is the case in our proposed ring network made up of fiber-ribbon point-to-point links [4]. The special feature of the network is a high-speed control-channel ring which is formed by one fiber in each link. The other fibers in the fiber-ribbons together form a ring dedicated for data traffic. The control ring improves performance of the network by sending medium access control information immediately before the data transmissions. High throughputs can be achieved in the network due to pipelining (spatial bandwidth reuse) [5], i.e., several packets can be traveling through the network simultaneously but in different segments of the ring. Fiber-ribbon links offering an aggregated

bandwidth of several Gb/s have already reached the market [6]. The increasingly good price/performance ratio for fiber-ribbon links indicates a great success potential for the proposed kind of network.

We have proposed to use the control ring for coordination of different kinds of services where slot-reserving combined with efficient slot-reuse is the main real-time service [7]. We believe, however, that good methods for the integration of both soft and hard real-time communication at different levels of QoS can be developed. Specifically, we have the idea of using the dedicated control ring, together with a distributed algorithm (implementing EDF, earliest deadline first, or similar), to globally "optimize" traffic with soft real-time requirements [8] and, at the same time, offer guaranteeing services for traffic with hard real-time requirements.

Often, a lot of the communication in distributed systems is group communication and process synchronization (e.g., semaphores and barrier synchronization). Therefore, real-time features of such mechanisms will also be emphasized in the project. In addition, special features of new network architectures can bring possibilities like the control ring for coordination purposes.

Low-level support, where some mechanisms might be implemented in hardware, can further increase performance and possibilities of implementing new methods. The project will therefore contain experimental phases including hardware design. Some hardware is currently being developed to demonstrate our ring network but without real-time features. Fiber-ribbon links for the demonstrator have been sponsored by Motorola. We believe this demonstrator is a good platform for further developments and experiments.

As stated above we have specific ideas and earlier work to rely on in further developments. However, it is important to stress that we have the ambition of finding methods with general applicability in the area of multi-service high-speed real-time networking.

Earlier work by our research group includes case studies of communication in radar signal processing applications. In the proposed project, our ambition is to evaluate the methods to be developed by using a broader spectrum of applications. To be able to perform relevant case studies, we have therefore chosen an industrial partner with knowledge from a wide range of applications. Also, case studies from a broader spectrum of applications help the evaluation and continuous development of services for heterogeneous real-time requirements. Our industrial partner is CR&T (Carlstedt Research and Technology) which, together with Halmstad University, makes two ARTES nodes included in the project.

3. Expected results and impact

The main expected results are:

- Methods for integration of soft and hard real-time communication services, at different levels of QoS (e.g., error and deadline-violation probability), into high-performance networks.
- General outlines of how to take advantage of architectural features in high-speed networks and novel optical network architectures when designing distributed real-time systems.
- Performance analysis of proposed methods.
- Case studies where the industrial applicability of the proposed methods is confirmed.
- A Ph.D. who has good understanding of high-speed real-time networking together with its industrial applicability.
- Commercial opportunities in high-speed networks for embedded real-time applications as well as methods for application mapping onto such networks.

4. Project plan

The tasks and deliverables below are specified for the period 991001-041001, which includes 2×2 years of research and one year of teaching for one Ph.D. student. The project work can be described as two parallel tracks that run concurrently through the whole project: (i) development and general analysis of methods and architectures, and (ii) case studies, where developed methods are evaluated in terms of applicability in industrial applications, and other application oriented work. The main contributions from our industrial partner will be in the second track.

Task A (9910-0003): Identify emerging industrial applications with high communication demands and heterogeneous real-time requirements. The industrial partner, with its broad area of competence, will play an important active role in this selection of relevant applications. The selection results in both criteria for the survey (see Task B), motivations/outlines for the development of new methods, and in representative case studies to be performed when methods have been developed later on during the project. Examples of demands, which in ongoing projects have been stated for radar signal processing systems, are real-time services for group communication and process-synchronization.

Deliverable 1 (0003): Technical report from Task A.

Task B (9910-0006): Survey of real-time communication and high-performance networks (focus on real-time even though the amount of reported work is limited) including both novel optical network technologies and existing/emerging standards. Especially, network architectures are evaluated in the scope of their suitability for integration of different real-time services.

Deliverable 2 (0006): State-of-the-art report from Task B.

Task C (0003-0011): Develop methods/protocols/hardware-specifications based on the idea about taking advantage of architectural support to implement efficient global "optimization" of real-time traffic in a high-performance network.

Deliverable 3 (0010): First conference paper, mainly based on work in Task C.

Task D (0008-0103): Implement, verify, and measure performance of the results from Task C on the demonstrator platform (including hardware update).

Deliverable 4 (0103): Technical report and/or conference paper, mainly based on work in Task D.

Task E (0012-0108): Develop methods of integrating different real-time services and QoS (e.g., services for messages with soft and hard deadline services, and for real-time group-communication) into a high-performance network.

Task F (0105-0112): Implement, verify, and measure performance of the results from Task E on the demonstrator platform (including hardware update).

Deliverable 5 (0112): Conference paper, mainly based on work in Task E and F.

Task G (0111-0203): Case study where appropriate application(s) are chosen in co-operation with CR&T for use in suitability tests and performance analysis of developed methods.

Task H (ongoing refinement during the whole project): Develop general outlines of how to take advantage of architectural features in high-speed networks and novel optical network architectures when designing distributed real-time systems.

Deliverable 6 (0204): Licentiate thesis. Specification of the remaining part of the project. Possibly conference papers and/or journal papers.

Second half of the project (0204-0410): Evaluate industrial relevance of work done so far and consider commercial exploitation possibilities. Further on, proposed methods should be evaluated in the scope of how they can be used in a more general framework. This includes looking at possibilities of integrating the methods into development tools or similar. In addition, more case studies should be done in order to get further performance measures and refinements of the methods, both to increase the industrial influence and to get better understanding of support for heterogeneous real-time requirements.

Deliverable 7 (0410): Ph.D. thesis.

5. Preliminary budget

We ask for funding of one Ph.D. student. The budget below is specified for a period of 4 years. For a five year period (80 % and 20 % for Ph.D. student and supervisor/project-leader respectively), the annual cost is 549 kkr for the first year. Assuming that the cost grows by 4 % per year, the cost for year 2 through 5 is 571, 594, 617, and 642 kkr, respectively. Thus we ask for a total funding of 2973 kkr. Carl Bergenhem, currently with a project employment at Halmstad University, is an appropriate candidate for employment as Ph.D. student, while Magnus Jonsson will supervise and lead the project. Main thesis advisor/examiner is Professor Bertil Svensson.

Annual budget (kkkr)

Salary (incl. Social costs)	
PhD student (100%)	313
Senior researcher (25%) (M. Jonsson)	109
Department cost, incl. computer cost (10 % of above)	42
Travel	19
Project specific equipment	28
Localities (8.7% of the above)	44
University administration (13.64% of the above)	76
TOTAL (first year) excl VAT	631
University VAT 8% (add 8.7%)	55
TOTAL (first year) incl. VAT	686

6. Related research

There are rather many high-performance networks, both commercially available and reported in the research literature. To our knowledge, however, they all have limitations when it comes to their suitability of adoption into emerging embedded distributed real-time systems. Most of them fail in offering real-time services while the real-time services in networks which really have support for it do not offer enough heterogeneous ranges of services. For example, Mercury's RACEway only has support for priority-stamped messages, while the real-time protocol for WDM star networks described in [9] only has support for soft real-time requirements. The TD-TWDM network [10] and the CC-FPR network [4] have no support for global "optimization" of best-effort traffic (only EDF-queues locally in the nodes). A lot of work on real-time features in packet-switched networks have been carried out [11], but the focus has been on wide area networks, e.g., the function of each separate switch and the handling of each separate logical connection. Research focus must be placed on methods to offer services for heterogeneous real-time requirements on a system-wide integration/optimization basis. For the correct function of an embedded system, the real-time requirements must be considered for the system as a whole, not just trying to meet the requirements in each node separately.

7. Industrial relevance

In order to cope with more heterogeneous real-time requirements and higher bandwidth requirements in emerging embedded distributed systems applications (see application examples above), companies like our industrial partner must adopt new methods to get the desired real-time communication services. Our industrial partner works in a wide area of applications together with several companies. Therefore, we foresee a good variety of application examples to be used in case studies, together with good chances for the research results to be useful and used by the industry. Case studies of several applications and implementation of vital parts on a hardware platform will verify the applicability of developed methods in real industrial environments. This will be a continuous iteration to secure the industrial relevance of the research.

We judge the exploitation possibilities as good. The project might, for example, result in new (or stronger) business areas for involved companies, a new company which can offer special competence in the field of high-performance real-time networks, and/or a new company with products in the field (software and/or hardware).

8. Relation to the profile and other SSF programs

Other SSF programs which might be related from a broad point of view (e.g., applications where a high-performance real-time network might be needed) are: Computer Science and Systems Engineering (ECSEL), Telecommunications (PCC), The Interactive Institute, Photonics, and Autonomous Systems. From a more narrow point of view we find many relations to ongoing ARTES research (research on, e.g., distributed real-time systems and multiprocessors, scheduling, and heterogeneous real-time systems) and we find our project very suitable for the current ARTES call. Some of the keywords found in the call, and which we find especially noteworthy as related to our project, are: multimedia, avionics, defense applications, telecommunications, predictability, reliability, *heterogeneous real-time systems*, *distributed systems*, results applicable to industrial problems, case studies, and exploitation possibilities.

9. Context

The research will be carried out at the Centre for Computer Systems Architecture (CCA), Halmstad University. The research center, and its leader Professor Bertil Svensson, has long experience in, e.g., embedded real-time computing, distributed real-time systems, and parallel computing in general. The project is to be led by Magnus Jonsson who has made significant contributions in the field of high-performance real-time networks and who is planned to defend his Ph.D. thesis this fall.

CR&T is a company working in the field of, e.g., embedded systems, real-time systems, data communications, parallel and distributed processing, and dependable and fault-tolerant systems. Real-time communication is relevant in several of CR&T's areas, both in pure communication systems and in important sub-systems of embedded/distributed systems. In addition to our industrial partner taking active part in this project, we have a large network including both academic, industrial, national, and international contacts.

The proposed project is complementary to two projects (funded by KKS and NUTEK) which have been running for a while. The other two projects are more oriented towards parallel processing, especially embedded signal processing, including research on fiber-optic network architectures and protocols for these networks. As indicated above, some support for real-time traffic exists in the networks, but the proposed project will have a clear orientation of real-time communication which complements the former/ongoing work well. The co-operation with Ericsson Microwave Systems in the two projects gives us really demanding application examples useful in the proposed project.

References

- [1] M. Taveniku, A. Ahlander, M. Jonsson, and B. Svensson, "The VEGA moderately parallel MIMD, moderately parallel SIMD, architecture for high performance array signal processing," *Proc. 12th International Parallel Processing Symposium & 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP'98)*, Orlando, FL, USA, Mar. 30 - Apr. 3, 1998, pp. 226-232.
- [2] M. Jonsson, A. Ahlander, M. Taveniku, and B. Svensson, "Time-deterministic WDM star network for massively parallel computing in radar systems," *Proc. Massively Parallel Processing using Optical Interconnections, MPPOI'96*, Lahaina, HI, USA, Oct. 27-29, 1996, pp. 85-93.
- [3] M. Jonsson, "Fiber-optic interconnection networks for signal processing applications," *4th International Workshop on Embedded HPC Systems and Applications (EHPC'99)*, San Juan, Puerto Rico, Apr. 16, 1999. Published in *Lecture Notes in Computer Science*. vol. 1586, Springer Verlag, pp. 1374-1385, 1999, ISBN 3-540-65831-9.
- [4] M. Jonsson, "Two fiber-ribbon ring networks for parallel and distributed computing systems," *Optical Engineering*, vol. 37, no. 12, pp. 3196-3204, Dec. 1998.
- [5] P. C. Wong and T.-S. P. Yum, "Design and analysis of a pipeline ring protocol," *IEEE Transactions on communications*, vol. 42, no. 2/3/4, pp. 1153-1161, Feb./Mar./Apr. 1994.
- [6] D. Bursky, "Parallel optical links move data at 3 Gbits/s," *Electronic Design*, vol. 42, no. 24, pp. 79-82, Nov. 21, 1994.
- [7] M. Jonsson, C. Bergenheim, and J. Olsson, "Fiber-ribbon ring network with services for parallel processing and distributed real-time systems," *Submitted for reviewing*, Feb. 1999.
- [8] C. Bergenheim and J. Olsson, "Protocol suite and demonstrator for a high performance real-time network," *Master thesis, Centre for Computer Architecture (CCA), Halmstad University, Sweden*, Jan. 1999.
- [9] A. Yan, A. Ganz, and C. M. Krishna, A distributed adaptive protocol providing real-time services on WDM-based LAN's, *Journal of Lightwave Technology*, vol. 14, no. 6, pp. 1245-1254, 1996.
- [10] M. Jonsson, K. Börjesson, and M. Legardt, "Dynamic time-deterministic traffic in a fiber-optic WDM star network," *Proc. 9th Euromicro Workshop on Real Time Systems*, Toledo, Spain, June 11-13, 1997, pp. 25-33.
- [11] H. Zhang, "Service disciplines for guaranteed performance in packet-switching networks," *Proceedings of the IEEE*, vol. 83, no. 10, pp. 1374-1396, Oct. 1995.

Appendix A: Short CV of Magnus Jonsson

Magnus Jonsson, Lic. Tech., Acting Associate Professor

Present position: Acting Associate Professor in Data Communications, Halmstad University, Halmstad, Sweden.

Education:

- Ph.D. planned for October 1999 with the thesis title "High-performance fiber-optic communication networks for distributed real-time computing systems".
- Lic. Tech. in Computer Systems Engineering with the thesis "Fiber-optic interconnections in high-performance real-time computer systems" at Chalmers University of Technology, Göteborg, Sweden, 1997.
- M.Sc. in Computer Systems Engineering, Halmstad University, Halmstad, Sweden, 1994.

Publications: More than ten international refereed papers, most of them in the area of real-time and fiber-optic communication.

Research orientation: Computer communications, embedded real-time systems, parallel and distributed computing, optical network architectures. Former research activities also include robotics and artificial neural networks.

Current and recent Teaching:

Course: Data Communications (both Electrical Engineering and Computer Engineering programs), course responsible and examiner.

Coordinator and examiner of all Bachelor degree projects in the program of Computer Systems Engineering, 1998.

Course: Parallel Computer Architecture, course responsible and examiner.

Course: Communication, course responsible and examiner (focused on presentation techniques but report writing was also included).

Advisor of Master and Bachelor degree projects in the programs of Computer Systems Engineering and Electrical Engineering at Halmstad University.

Other commissions:

Member of the Research Committee (Forskningsnämnden) at Halmstad University.

Member of the Committee of Computer and Information Technology (dIT-nämnden) at Halmstad University.

Recent publications relevant for the project:

Jonsson, M., C. Bergenhem, and J. Olsson, "Fiber-ribbon ring network with services for parallel processing and distributed real-time systems," *Submitted for reviewing*, Feb. 1999.

Jonsson, M., "Fiber-optic interconnection networks for signal processing applications," *4th International Workshop on Embedded HPC Systems and Applications (EHPC'99)*, San Juan, Puerto Rico, Apr. 16, 1999. Published in *Lecture Notes in Computer Science*. vol. 1586, Springer Verlag, pp. 1374-1385, 1999, ISBN 3-540-65831-9.

Jonsson, M., "Two fiber-ribbon ring networks for parallel and distributed computing systems," *Optical Engineering*, vol. 37, no. 12, pp. 3196-3204, Dec. 1998.

Jonsson, M., "Control-channel based fiber-ribbon pipeline ring network," *Proc. Massively Parallel Processing using Optical Interconnections (MPPOI'98)*, Las Vegas, NV, USA, June 15-17, 1998, pp. 158-165.

Taveniku, M., A. Ahlander, M. Jonsson, and B. Svensson, "The VEGA moderately parallel MIMD, moderately parallel SIMD, architecture for high performance array signal processing," *Proc. 12th International Parallel Processing Symposium & 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP'98)*, Orlando, FL, USA, Mar. 30 - Apr. 3, 1998, pp. 226-232.

Jonsson, M., B. Svensson, M. Taveniku, and A. Åhlander, "Fiber-ribbon pipeline ring network for high-performance distributed computing systems," *Proc. International Symposium on Parallel Architectures, Algorithms and Networks (I-SPAN'97)*, Taipei, Taiwan, Dec. 18-20, 1997, pp. 138-143.

Jonsson, M., K. Börjesson, and M. Legardt, "Dynamic time-deterministic traffic in a fiber-optic WDM star network," *Proc. 9th Euromicro Workshop on Real Time Systems*, Toledo, Spain, June 11-13, 1997, pp. 25-33.

Jonsson, M. and B. Svensson, "On inter-cluster communication in a time-deterministic WDM star network," *Proc. 2nd Workshop on Optics and Computer Science (WOCS)*, Geneva, Switzerland, Apr. 1, 1997.

Jonsson, M., A. Åhlander, M. Taveniku, and B. Svensson, "Time-deterministic WDM star network for massively parallel computing in radar systems," *Proc. Massively Parallel Processing using Optical Interconnections, MPPOI'96*, Lahaina, HI, USA, Oct. 27-29, 1996, pp. 85-93.

Appendix B: Short CV of Bertil Svensson

Bertil Svensson, Ph.D., Professor

Date of birth: February 15, 1948, Eldsberga, Sweden.

Present position: Professor of Computer Systems Engineering, Halmstad University and Chalmers University of Technology, Halmstad and Göteborg, Sweden.

Qualifications:

- Ph.D. with the thesis "LUCAS Associative Processor Array - Design and Applications" at Lund University, Lund, Sweden, 1983.
- M.Sc. in Electrical Engineering at Lund University, Lund, Sweden, 1970.

Experience:

- December, 1998 to present: Professor of Computer Systems Engineering, Halmstad University and Chalmers University of Technology, Halmstad and Göteborg, Sweden.
- December, 1991 to April, 1999: Professor of Computer Systems Engineering, Chalmers University of Technology, Göteborg, Sweden.
- May, 1983 to December, 1991: Associate Professor in Computer Engineering, Halmstad University, Sweden.
- 1974 to 1983: Periodically Lecturer, periodically Research Assistant, Dept. of Computer Engineering, University of Lund, Sweden.
- 1977 to 1983: Co-founder and co-owner of Synthese, a company for development and production of development tools for microprocessors, Lund, Sweden.
- 1970-1971, 1972-1974: Teaching assistant, Dept. of Computer Engineering, University of Lund, Sweden.

Appointments:

- July, 1993 to April, 1999: Vice-Dean of the School of Electrical and Computer Engineering at Chalmers University of Technology and, as such, responsible for the undergraduate study programme (for the Master's degree) in Computer Science and Engineering
- November, 1987 to present: Head of the Centre for Computer Systems Architecture (CCA), Halmstad University, Sweden. (First called the Centre for Computer Science).
- May, 1986 to June, 1988: Initiator and co-leader of the Centre for Image Processing and Computer Graphics
- October, 1983 to December, 1991: Vice President of Halmstad University, Sweden.
- July, 1989 to October, 1991: Acting Professor of Computer Systems Engineering, Luleå University of Technology, Sweden.

Publications: More than 60 papers published in international journals and conference proceedings. Co-author of two books.

Research supervision: 13 Licentiate theses, 5 Ph.D. theses

Research orientation: Massively parallel computer architectures, real-time parallel computer systems, artificial neural networks, autonomous robot navigation.

Appendix C: Letter of intent from CR&T

Göteborg 990517

Intresseförklaring för projektet Methods for Integration of Soft and Hard Real-Time Communication into High-Performance Networks

CR&T har under flera år samlat en omfattande kompetens inom området distribuerade realtidssystem genom att i företaget samla forskare och civilingenjörer med avancerad teknisk bakgrund i ett flertal datavetenskapliga och datortekniska områden, samt genom en stor mängd utrednings- och utvecklingsuppdrag för industriföretag.

All erfarenhet inom detta ämnesområde pekar på att hanteringen av krav på realtidsegenskaper i distribuerade system får en allt större betydelse och fordrar utveckling av nya metoder för att möta snabbt ökande krav på prestanda, analyserbar Quality of Service, heterogenitet, och tillämpning av nya tekniker för beräkning och kommunikation.

Mot denna bakgrund anser vi att Magnus Jonssons och Bertil Svenssons projektförslag tar upp mycket centrala frågeställningar, och kan leda till resultat vilka har tillämpning inom flera sammanhang där CR&T är verksamt. Inte minst är problemet med att kombinera mjuka och hårda realtidskrav av särskild betydelse. Kombinationen med kommunikationsnätverk med mycket hög bandbredd är också betydelsefull.

Sammanfattningsvis har det föreslagna forskningsprojektet hög relevans inom CR&Ts verksamhetsområden. Vi är därför beredda att ingå i en referensgrupp för projektet, där vår roll bland annat blir att bidra med vår samlade erfarenhet av den aktuella typen av system.

Kontaktperson för projektet hos CR&T är Jonas Vasell (e-post vasell@crt.se, telefon 031-701 42 00).

Jesper Vasell
vice VD, Carlstedt Research & Technology AB