A Running Leap for Embedded Signal Processing to Future Parallel Platforms

Bertil Svensson¹, Zain Ul-Abdin¹, Per M. Ericsson², Anders Åhlander², Hoai Hoang Bengtsson³, Jerker Bengtsson², Verónica Gaspes¹, Tomas Nordström¹

1. Centre for Research on Embedded Systems, Halmstad University, PO Box 823, SE 301 18 Halmstad, Sweden

2. Saab AB (EDS), Gothenburg, Sweden

3. Viktoria Swedish ICT, Gothenburg, Sweden.

E-mail: {bertil.svensson, zain-ul-abdin, veronica.gaspes, tomas.nordstrom}@hh.se; {per.m.ericsson, anders.ahlander, jerker.bengtsson}@saabgroup.se; hoai.bengtsson@viktoria.se Telephone: +46 35 167131 (Bertil Svensson)

ABSTRACT

This paper highlights the collaboration between industry and academia in research. It describes more than two decades of intensive development and research of new hardware and software platforms to support innovative, high-performance sensor systems with extremely high demands on embedded signal processing capability. The joint research can be seen as the run before a necessary jump to a new kind of computational platform based on parallelism. The collaboration has had several phases, starting with a focus on hardware, then on efficiency, later on software development, and finally on taking the jump and understanding the expected future. In the first part of the paper, these phases and their respective challenges and results are described. Then, in the second part, we reflect upon the motivation for collaboration between company and university, the roles of the partners, the experiences gained and the long-term effects on both sides.

Categories and Subject Descriptors

C.1.4 [Parallel Architectures]: Distributed architectures, Mobile processors

D.1.3 [Concurrent Programming]: Distributed programming, Parallel programming

Keywords

Industry-academia collaboration; Embedded signal processing; Parallel computing platforms; Software development.

1. INTRODUCTION

In advanced industrial products, embedded software and hardware are core ingredients that are renewed between successive product

WISE'14, September 16, 2014, Vasteras, Sweden.

Copyright © 2014 ACM 978-1-4503-3045-9/14/09...\$15.00. http://dx.doi.org/10.1145/2647648.2647653

generations. In the case when specialized hardware is needed there is a natural focus on developing hardware. In cases when the hardware is based on commercially available platforms, the focus is more on methods for software development. Software is often a major investment that needs to be transferred from one hardware platform to the next. The platform development has traditionally been incremental, making the software transfer problem possible to handle with traditional software development tools (compilers, simulators, profiling tools, etc.). Facing the current switch in platform technology from single processor systems via multicore systems to manycore platforms, new questions arise: what do the new platforms look like? Is there a predictable evolution in future platforms and, if so, do we know how? Similar questions apply to the applications; do we know how these develop (driven by needs and possibilities, and by non-functional properties, etc.)?

Such questions are in focus in the current phase of a long-lasting research cooperation between university and industry that is described and reflected over in this paper. By looking at the development over time, we see that the question of jumping to radically new platforms can be seen as a process over several years - even decades. This is what is behind the title of this paper - we don't see the leap to future parallel platforms as something that comes suddenly, rather it has been prepared by the running for several years - this will hopefully make the jump more controlled and powerful. The intention of the paper is not so much to describe the specific research problems, approaches, solutions, results and technical conclusions; rather the focus is on the different phases in this long-term collaboration, on the different roles of the partners, on the experiences gained, and on the longterm effects on the two organizations.

The company side of the collaboration is the radar development part of Saab, a Sweden-based security and defense company. In the early phases of the cooperation this division was part of Ericsson under the company names Ericsson Radar Electronics and Ericsson Microwave Systems. The university side is the Centre for Research on Embedded Systems (CERES) at Halmstad University. In the early phases of the cooperation also Chalmers University of Technology was involved. In this paper we will refer to the cooperating partners as "the company" and "the university", respectively.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2. CO-PRODUCTION RESEARCH BETWEEN INDUSTRY AND ACADEMIA

Hardware and software platforms for high-end embedded signal processing has been the theme of a long-term cooperation between the company and the university for more than twenty years. The research cooperation started in the early 1990's. The university had a research interest in new, massively parallel computer architectures and was looking for challenges in industrial applications that could be met with this kind of architectures. The company, on the other hand, was in the initial phase of adopting the new, array based, radar technology that relied on active electronically scanned array (AESA) antennas in combination with highly computation demanding signal processing techniques such as space-time adaptive processing (STAP). Similar challenges were found in the synthetic aperture radar based radar techniques (SAR).

2.1 Phases

Signal processing in array-based radar is a true performance challenge. During the first years of cooperation, the focus was on possible hardware structures that could – by specialized design – deliver the required performance, while still being programmable, albeit in a limited sense. We call this PHASE 1, which thus had a focus on hardware.

With the need for larger flexibility and higher degrees of programmability, the focus then gradually switched towards fully programmable hardware structures that could support both efficiency in application programming and efficiency in energy usage. This was PHASE 2, which thus had a focus on efficiency.

A few years into the current century, commercial manycore architectures and coarse-grained reconfigurable array processors started to appear [1]. It became clear that there might be suitable compute modules available for future needs in, e.g., radar signal processing. The architecture of complete systems, as well as the software development, will still be great challenges, however. This shifted the focus of the collaboration towards system and software development. This is the recently finished PHASE 3.

Now when manycore processors have arrived and can be used in applications, we study how they will develop during the next decade so that we can adapt the software development methods to future architectures, and not only to the architectures we see today. Heterogeneity in manycores is one of the foci in the ongoing PHASE 4.

This evolution can be seen in relation to the different projects carried out involving the company and the university over the last decade and a half. Of course, this gradual shift of focus has characterized also other collaborations and was linked into the needs of other companies and research interests of other universities, which has happened, e.g., through EU projects.

2.2 Co-Production Projects

The phases and projects in the collaboration are summarized in Figure 1.

2.2.1 Phase 1 – Focus on Hardware

The first project in the collaboration was not actually set up as a collaboration project, rather it contained activities that triggered the cooperation and resulted in the first joint work.

1989-1996 REMAP – massively parallel architectures meet applications that need them

The goal of the REMAP project (Real-time, Embedded, Modular, Adaptive, Parallel processor project) [2, 3, 4], was to gain new knowledge about the design and use of massively parallel computer architectures in embedded real-time systems. The prototype system, REMAB-Beta, built in the project was designed with massively parallel artificial neural network computations as the main target application area, building on the belief that learning and adaptation is an important part of embedded, actionoriented systems interacting with real-world environments. However, the project also caught interest from industry developing advanced sensor systems and thus marked a starting point of a long-term collaboration between the company and the university. As one of the first efforts related to this application area, a phased array multi-channel radar signal processing chain was used as a case study for the REMAP-Gamma architecture that was developed towards the end of the project [5].

An important activity at this stage was an advanced course given by the company's radar experts, targeting all researchers at the universities (Chalmers and Halmstad), including PhD students, junior researchers and professors from different areas of electrical and computer engineering. The common understanding of the techniques used and the challenges coming up was invaluable for the development of a fruitful cooperation.

1997-2000 PARAD/HSSP – modular massively parallel architectures

The PARAD/HSSP project (Parallel architectures for high-speed signal processing applications) was the first major collaboration project between the company and the university. The goal was to develop a scalable and flexible system architecture with enough performance for multi-channel signal processing. This required developing (1) new processor-array modules with enough processing performance, (2) a flexible, high performance communication architecture between modules, and (3) an environment for application software development [6, 7]. As a final, extensive part of the project the company made a major investment in the development of a complete system proposal, using a VLSI technology that would allow productification approximately in 2002 [8].

2000-2003 HiSPOT – optics-based massively parallel architectures

The HiSPOT project (High speed optoelectronics for optical interconnects) aimed at the development of technology and architectures for high throughput (>100 Gbps) optical inter-chip interconnect networks for future communication and radar systems. On the academic side, Halmstad University joined forces with Chalmers University of Technology, and on the company side three companies participated. In the project, basic academic research – investigating what can be achieved in terms of new high-speed components – met industrial application demands in the foreseeable future. In this meeting place, the elaboration of different architectures played an important intermediary role.

The achievements of the project on the system level included the identification of the planar free-space optical interconnect as a suitable technology for real-time embedded signal processing



Figure 1. Projects landscape corresponding to various phases.

systems and telecommunication systems using hypercube architectures [9].

2.2.2 Phase 2 – Focus on Efficiency

With the development of new architectural concepts and concrete modules for their realization, which characterized the first phase of the joint research, came the understanding that the nonfunctional aspects, in terms of energy efficiency and programming efficiency, also needed attention. Since these are things related to the possible productification, these aspects were particularly underlined by the company.

2000-2003 EEE – energy- and engineer efficiency

The EEE project (Energy- and engineering efficiency in parallel architectures and application development for future embedded signal processing systems) emphasized not only the need for performance but – even more – the need to develop applications in an economically efficient way. From the energy efficiency point of view, the primary purpose was to gain knowledge about which components will be best to build systems on and what system architectures should be used. The assumption was made that, in a future product development phase, similar components will probably be available commercially.

Concerning engineer efficiency, a platform approach was suggested [10]. The platform should offer an understandable and stable development model, and at the same time give the possibility to take advantage of the rapid technology development, including the use of new parallel architectures.

Thus it must support multiple hardware targets, and the development model should decouple application development from mapping aspects. The GEPARD platform, developed at the company, was presented as an illustrating example of the approach, applied to STAP and SAR radar signal processing.

2003-2006 RSA – reconfigurable system architectures

The RSA project introduced multi-level reconfigurability as a design principle for meeting the great need for embedded highperformance real-time computing in applications where mainstream computer system approaches are not sufficient. The purpose was to find ways of reaching the performance and characteristics of special-purpose processors without developing systems from scratch for each application. Reconfigurability and a modular approach were seen as a key to success [11]. One of the sub-goals was to evaluate the efficiency of some emerging, reconfigurable, parallel processing paradigms of varying granularity in selected embedded applications.

The RSA project marked the first activities related to new programming models, languages and tools. For example, a domain-specific stream processing model was developed [12], and a minimal parallel processing language targeted for reconfigurable architectures was evaluated through selected implementation experiments [13].

2005-2007 MISPA – memory intensive signal processing architectures

The real-time image forming in future, high-end synthetic aperture radar systems is an example that puts new demands on computer architectures. In the MISPA project the question was asked if it was at all possible to meet the demands with state-of-the-art technology or foreseeable new technology. The computational flow, with its associated memory, data bandwidth and processing demands, was analyzed in detail [14]. A critical issue was the execution of the hotloops, in which contributing data elements are combined to new data elements, after complicated index calculations and data interpolations. The chosen interpolation method affects the total performance and memory demands significantly; realizing this in a systolic or SIMD fashion was therefore suggested.



Figure 2. Simplified block diagram of AESA signal processing chain.

2.2.3 Phase 3 – Focus on Software Development

In this phase of cooperation, our focus moved towards developing methods for high performance signal processing software that is portable to several parallel platforms. We realized that several of the demands of these complex applications are common with other industrial domains such as baseband processing in telecommunication systems, which led to expanding our collaboration with European partners.

2007-2010 EPC – software tools for efficient parallel computing

The EPC (Embedded Parallel Computing) project had an emphasis on understanding which overall (hardware and software) architectures that are best suited for high-performance embedded applications. The project highlighted the usefulness of domain-specific programming techniques [15], specialized models of computation [16], and their corresponding mapping strategies for efficient execution on parallel and reconfigurable computing systems [17].

2010-2013 SMECY – complete software tool chains

It is evident that complex industrial applications require a holistic approach for software development, which requires development of complete software tool-chains. This was addressed in the EU/Artemis funded SMECY (Smart Multicore Embedded Systems) project [18] in which both the company and the university participated. The tool-chain developed in the project included aspects of application modeling, execution platform modeling, functional simulation, optimization and verification, parallelisation and design space exploration [19].

2.2.4 Phase 4 – Focus on the JUMP to Future Manvcore Platforms

Now that manycore processors are emerging as a viable processing platform in the future, in this ongoing phase we are investigating how manycore technologies align with the performance and cost requirements for real-time signal processing in the multi-functional systems and whether existing software approaches are sufficient to meet the engineering efficiency and portability demands.

2011-2013 JUMP – software tool chains for future

manycore architectures

The overall goal of the JUMP (JUmp to Manycore Platforms) project was to investigate and develop programming and development technologies enabling the efficient design and

programming of highly integrated embedded manycore systems. In particular the project dealt with embedded signal processing systems where traditional generic processor architectures and their corresponding programming models are not efficient enough to meet the required performances and/or other non-functional system properties. Manycore processor architectures offer scalable parallelism and the performance needed for implementation of the functionality required in high-end embedded sensor- and communication systems. We developed software tool support for targeting the emerging manycore architectures [20, 21, 22].

2014- ESCHER – meeting portability and programmability

from the SW side, and power efficiency and performance from the HW side

The main challenge ahead for industry is how to efficiently perform the technology shift from current application hardware platforms to future generations of platforms including heterogeneous manycore technology. The scientific challenge is to reconcile conflicting needs: development productivity and software portability from software perspective and meeting the performance and power efficiency from the hardware side. Software portability, program correctness and programmer productivity all demand that programs are written in a highlevel language capable of expressing application-level parallelism while abstracting away platform dependent physical parallelism. In contrast, performance efficiency is traditionally better achieved by developing code tuned for the specific target hardware. The ESCHER project (Embedded Streaming Computations on Heterogeneous Energy-Efficient Architectures) aims to bridge this gap between software development and the targeted heterogeneous hardware platforms [23, 24].

2.3 Industrial Applications as Challenges AESA

An Active Electronically Scanned Array (AESA) is a type of phased array radar whose transmitter and receiver functions are composed of numerous small solid-state transmit/receive modules. It may consist of hundreds, or thousands, of antenna elements. An AESA radar application is composed by a sequence of stages each one computing a different phase. In this document, a simplified radar signal processing chain is used to illustrate the computation of interest. The radar chain is composed by: a Digital Beam Forming (DBF) module, a Doppler Filter Bank (DFB) module, a Pulse Compression (PC) module, a Constant False Alarm Ratio (CFAR) module and finally an Integration (INT) module, as shown in Figure 2.

The main bottleneck of this application is the composition of the DFB and the PC modules. These parts of the processing have significant demands on computational performance (up to 73728 Million Operations Per Seconds for the DFB module) and memory bandwidth (up to 4GB/s), as well as different data access patterns between calculation steps. In addition, there are tight implementation constraints in terms of size, power consumption and cost. Furthermore, the entire signal processing chain has to be completed within a few tenths of milliseconds, so minimizing the latency of each computational step is another important requirement.

SAR

Synthetic-Aperture Radar (SAR) systems produce high resolution images of the ground by transmitting, and then receiving echoes of, radar pulses along a linear flight path, as shown in Figure 3. The amount of raw radar data depends on the size of the covered area, the resolution, and the processing mode. The back-projection integration technique has been applied in SAR systems, enabling processing of the image in the time domain, which makes it possible to compensate for nonlinear flight tracks. However, the cost is typically a high computational burden. The Fast Factorized Back-projection (FFBP) is a computationally efficient algorithm for image forming in the time domain. It reduces the performance requirements significantly relative to those for conventional global back-projection techniques. In FFBP, the whole aperture initially consists of a large number of small subapertures with low angular resolution. These subapertures are iteratively merged into larger ones with higher angular resolution, until the full aperture with full angular resolution is obtained. In reality, the flight path is not perfectly linear; this of course influences the image quality negatively. The autofocus algorithm is applied to compensate for the non-linearity in the flight path. The autofocus calculations use the image data itself and are done before each subaperture merge. The criterion assumed in this study is maximization of correlation of image data. As the criterion calculations are carried out many times for each merge (to test several hypotheses, each requiring substantial interpolation), it is important that these are done efficiently.





The integration time may be several minutes. The computational performance demands are tens or hundreds of GFLOPS. Thus, the large data sets that need to be processed by the SAR signal processor make it hard to meet the high performance that is required for real-time image creation, i.e. when the images are created during the flight. Another related challenge is to cope with the increased computational demands within a limited power budget. Thus there is a dire need to come up with solutions for on-board computing hardware that can meet the performance and energy requirements of real-time image creation in next generation radar systems.

3. MOTIVATIONS FOR COOPERATION

For the company, the motivation for cooperation is to obtain a better understanding of future opportunities related to advancement of technology, and thus build a more solid base for its decision-making. The company can benefit from collaborating with the researchers in early-stage design-space exploration of technology choices and evaluate their possible impact on future products and services. The university is also seen as a resource for dialogue, for active learning and for knowledge sharing related to ongoing research also in areas in which the company is not actively involved. This has the potential to lead to new avenues of cooperation in future research. The company also acknowledges the advantage of such a cooperation to identify the gaps in skills of its existing environment, such as the understanding of new models of computation that can be the basis of future software tools, and to bridge this knowledge gap in time. Doctoral and Master students working in the projects may also be a source of qualified human resources for the company. In particular, the opportunity to hire people with a PhD degree in an area perfectly matching the need of the company is seen as an advantage.

From the university perspective, the motivating factors include the identification of industrial challenges in order to motivate research directions. Close cooperation like the one we see in this case can result in a more accurate quantification of such challenges. The university gets an opportunity to get a complete view of the needs of the industry that can then be realized in future technology solutions. This is an aspect that is important for PhD and MSc students to get in contact with. We have also seen that complex industrial problems often require multidisciplinary research, which initiates the cooperation between different academic disciplines. The university also gets insights in the business aspects of the technology and application domains that they are working on, which in turn has a potential to result in more commercial spinoffs from the research.

The collaboration also provides benefits for education: students get the opportunity to do qualified master theses, teachers involved in the collaboration can use examples and guest lecturers from industry. Courses on the PhD level that are relevant for both academia and insustry can be given jointly, and the participation from both sides improves the courses.

4. ROLES OF THE PARTNERS

The primary role of the academic partner is to develop new knowledge and disseminate it through publications and teaching and use this knowledge for continued research. The academic partner also contributes by studying trends and analyzing the state-of-the-art, by conducting surveys of emerging technologies and developing prototypes and case studies based on industrial problems. The analysis of trends and state-of-the-art has led to suggestions for new platforms and design methods for the industry.

The company, on the other hand, has the role of using the knowledge developed in collaboration with the academic partner for improved products and services and spread the use of this knowledge through its business and marketing activities. The company also contributes by providing an understanding of future demands in the form of requirements and use cases. The company can actively push the research to provide concrete results and deliverables that are applicable in product development, either immediately or in the near future.

The company can also provide valuable feedback to the university through, e.g., evaluating the research from an industrial perspective or connecting results to products and business requirements and opportunities. The company also uses the outcomes of the cooperation to develop new platforms and utilize the development methodologies, thereby setting out new directions for product development.

A common role of both the collaborating partners is to educate each other, such as the knowledge transfer of models of computation from the university and the know-how of state-ofthe-art radar technology from the company.

5. EXPERIENCES

An interesting experience from our long-term collaboration is that research related to architecture has the potential to serve as "intermediator" between research on new basic technologies on the one side, and the industrial applications and their demands on the other. The clearest example is the HiSPOT project that bridged from new surface emitting lasers and diffractive optics components on the basic technology side to future radar systems on the application side. Similarly, fundamental computer science research related to programming languages meets industrial applications via demands in the form of efficiencyelevating domain specific languages.

On the more organizational side, our experiences tell that there is a lot to be gained if the partners - and all the participating staff – are well aware of the fact that their goals and objectives may differ. Stating the goals, objectives and tasks clearly, and making them known to everybody, facilitates that everybody strives in the same direction. An excellent way of achieving this mutual understanding is to let PhD students and researchers from the university spend time in industry. The partners in this cooperation have experience of even doing this in the form of strategic recruitment to fill a joint position. This means that a person at the junior researcher level is recruited for employment at the university but is partly financed by the industry and expected to spend half of the research time at the company. In our case this arrangement facilitated the cooperation in the JUMP project as well as the active joint participation in the EU/Artemis project SMECY.

A good way to make goals and tasks concrete is to develop demonstrators or prototypes. Doing so will put focus on specifications/discussions of requirements as well as on time planning. The integration work in the final stages of demonstrator/prototype implementation also serves as an integration of the human resources working in the project.

During the many years of collaboration we have also seen the value of distinct and clear – yet sensitive – project leadership. In general, academia has a lot to learn from industry when it comes

to project leadership and administration. Follow-up of progress and adapting to necessary changes is an important part.

Collaboration between only two partners is easier and more efficient than having many partners involved. However, the latter case has other advantages such as building larger networks, benefitting from crosstalk between projects and finding un-expected new opportunities. In general, a tight twoparty collaboration in combination with joint involvement in multiple-partner projects is an excellent combination.

Another challenge is timing. When industry has identified a problem that needs a solution it wants to start immediately, while university projects that require external financing might need up to a year from initial discussions to project start. It is our experience that flexibility within the ongoing projects, combined with a trustful cooperation with continuous discussions about what the next joint project should be about, can solve these problems.

6. LONG-TERM EFFECTS

A solid relation in the form of long-term research collaboration has clear long-term effects. First of all, it facilitates the joint work, since there is a common understanding to build upon when it comes to deciding what problems or opportunities that need to be studied. University-based research becomes part of the company's long-term planning and strategic working methods; likewise the university gets a better understanding also of the long-term perspectives in the company and can adjust its strategies to meet industrial needs.

On the company side, new employees are often recruited directly after a PhD, Licentiate or Master degree. In the described, 20 years long collaboration there are many examples of this. Some of the employees choose to – over the years – go back and forth between industry and academia, thus contributing further to tightening the cooperation. Of particular interest is the possibility of "synchronizing" an industrial career as a specialist/expert with an academic career, for example in the form of adjunct lecturer or adjunct professor.

On the university side the benefits of having staff moving back and forth between academia and industry are also apparent. Within the collaboration described in this paper there are some good examples of this. It is also interesting to see how this longterm research collaboration - together with other industry collaborations - has influenced the development of research orientation. CERES is now an internationally recognized research center that is particularly well known for its research in co-production with industry. This also strengthens the entire university, in particular in its profile as "the innovationpromoting university". Content wise, it is interesting to see how the continuous gradual shift in orientation of CERES' research has accompanied the development of the industry collaboration described in this paper: to involve more of, e.g., models of computation. programming languages and software development tools. This is to the benefit of several other industrial and academic collaborations.

7. CONCLUSIONS

This paper has summarized an example of industry-academia collaboration in different phases over a period of more than two decades and highlighted the key benefits, the challenges, and the impact that such cooperation has on the collaborating partners in the long term.

In terms of benefits, we conclude that our collaboration leads to a better understanding of future opportunities by the industry based on the advancement of technology, and for the academia it enables them to motivate their future research directions based on industrial challenges. Academia also gets the benefit of insights into the business aspects of the technology. Although the goals of the collaborating partners may – and should – differ, defining clear objectives and tasks enables the participating staff to strive in a common direction, such as joint development of prototypes and demonstrators.

The collaboration has also experienced challenges of various kinds, particularly in relation to the opportunities about developing hardware platforms and software development methods. Finding productive and rewarding ways of working together is a challenge in itself.

We also observe that the long-lasting collaboration has shifted focus over the years, from hardware platform development to the development of software methodologies and tools, and we have been able to align the scientific development with the industrial development.

8. ACKNOWLEDGMENTS

We acknowledge the devoted work of the – well over a hundred – project participants from all levels of the university and the company, as well as from other universities and other companies. We are grateful to other academic and industrial partners contributing in joint projects. Finally, we thank our financiers, in particular The Knowledge Foundation but also Vinnova (and its predecessor Nutek) and SSF.

9. REFERENCES

- Z. Ul-Abdin and B. Svensson, "Evolution in Architectures and Programming Methodologies of Coarse-Grained Reconfigurable Computing." *Microprocessors and Microsystems*, Vol. 33(3), 2009, pp. 161-178.
- [2] K. Nilsson, B. Svensson, and P. A.Wiberg, "A modular, massively parallel computer architecture for trainable realtime control systems", *Control Engineering Practice*, vol. 1, no. 4, 1993, pp. 655-661.
- [3] L. Bengtsson, A. Linde, B. Svensson, M. Taveniku ,and A. Åhlander, "The REMAP massively parallel computer platform for neural computations", *Proceedings of the Third International Conference on Microelectronics for Neural Networks (MicroNeuro '93)*, Edinburgh, Scotland, UK, 1993, pp. 47-62.
- [4] T. Nordström, B. Svensson, L. Bengtsson, M. Taveniku, and A. Linde, "The REMAP reconfigurable architecture: A retrospective" Chapter in Omondi, A. R. and J.C. Rajapakse (eds.) FPGA Implementations of Neural Networks. Springer, 2006.
- [5] L. Bengtsson and B. Svensson, "A globally asynchronous, locally synchronous SIMD processor", *Proceedings of MPCS'98: Third International Conference on Massively Parallel Computing Systems*, Colorado Springs, Colorado, USA, April 2-5, 1998.
- [6] M. Taveniku, A. Åhlander, M. Jonsson, and B. Svensson, "The VEGA moderately parallel MIMD, moderately parallel SIMD, architecture for high-performance signal processing", *Proceedings* 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.

- [7] M. Jonsson, B. Svensson, M. Taveniku, and A. Åhlander, "Fiber-ribbon pipeline ring network for high-performance distributed computing systems", *Proceedings of International Symposium on Parallel Architectures, Algorithms and Networks (I-SPAN '97),* Taipei, Taiwan, Dec. 18-20, 1997, pp. 138-143.
- [8] A. Åhlander and A. Åström, "A high speed signal processing system", *Sixth Annual High Performance Embedded Computing Workshop*, Lincoln Laboratory, Lexington, MA, USA, 24–26 September 2002.
- [9] H. Forsberg, M. Jonsson, and B. Svensson, "Embedded signal processing using free-space optical hypercube interconnects" *SPIE Optical Networks Magazine*, vol. 4, No. 4, pp. 35-49, July/August 2003.
- [10] A. Åhlander, A. Åström, B. Svensson, and M. Taveniku, "Meeting engineer efficiency requirements in highly parallel signal processing by using platforms" *Proceedings* of IASTED Parallel and Distributed Computing and Systems 2005, Phoenix, AZ, USA, Nov. 14-16, 2005.
- [11] D. Johnsson, A. Åhlander, and B. Svensson, "Analyzing the advantages of run-time reconfiguration in radar signal processing", *Proceedings of 17th IASTED Int'l Conf. on Parallel and Distributed Computing and Systems (PDCS* 2005), Phoenix, AZ, USA, Nov. 14-16, 2005.
- [12] J. Bengtsson and B. Svensson, "A configurable framework for stream programming exploration in baseband applications", 11th Int'l Workshop on High-level Parallel Programming Models and Supportive Environments (HIPS 2006), Proceedings of IPDPS 2006, The Int'l Parallel & Distributed Processing Symposium, Rhodes, Greece, April 25-29, 2006.
- [13] Z. Ul-Abdin and B. Svensson, "A study of design efficiency with a high-level language for FPGAs", Proceedings of 14th Reconfigurable Architectures Workshop RAW 2007, held in conjunction with the IEEE International Parallel and Distributed Processing Symposium (IPDPS 2007), Long Beach California, USA, March 26-27, 2007.
- [14] A. Åhlander, H. Hellsten, K. Lind, J. Lindgren, and B. Svensson, "Architectural challenges in memory-intensive, real-time image forming", *Proceedings of 36th International Conference on Parallel Processing (ICPP)*, Xi'an, China, Sept. 10-14, 2007.
- [15] J. Bengtsson and B. Svensson, "A domain-specific approach for software development on multicore platforms", *Computer Architecture News*, December 2008.
- [16] Z. Ul-Abdin and B. Svensson, "Using a CSP based programming model for reconfigurable processor arrays", *Proceedings of 2008 International Conference on ReConFigurable Computing and FPGAs (ReConFig'08)*, Cancun, Mexico, Dec. 3-5, 2008.
- [17] J. Bengtsson and B. Svensson, "Manycore performance analysis using timed configuration graphs", *Proceedings of IEEE International Conference on Embedded Computer Systems: Architectures, Modelling, and Simulation* (SAMOS IX), Samos, Greece, July 20-23, 2009.

- [18] "Smart Multicore Embedded Systems", Editors: M. Torquati, K. Bertels, S. Karlsson, F. Pacull, Springer, ISBN: 978-1-4614-8800, 2013.
- [19] Z. Ul-Abdin, E. Gebrewahid, and B. Svensson. "Managing Dynamic Reconfiguration for Fault-tolerance on a Manycore Architecture", *Proceedings of 19th International Reconfigurable Architectures Workshop (RAW'12) in conjunction with International Parallel and Distributed Processing Symposium (IPDPS'12)*, Shanghai, China, May 2012.
- [20] Z. Ul-Abdin, A. Åhlander, and B. Svensson, "Real-time Radar Signal Processing on Massively Parallel Processor Arrays", *Proceedings of 47th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove-CA, USA, November 2013.
- [21] Z. Ul-Abdin, A. Åhlander, and B. Svensson, "Energy Efficient Synthetic-Aperture Radar Processing on a Manycore Architecture", *Proceedings of 42nd International*

Conference on Parallel Processing (ICPP), Lyon, France, October 2013.

- [22] Z. Ul-Abdin and B. Svensson. "Occam-pi for Programming of Massively Parallel Reconfigurable Architectures", *International Journal of Reconfigurable Computing*, Vol. 2012, Article ID 504815, 2012.
- [23] S. Savas, E. Gebrewahid, Z. Ul-Abdin, T. Nordström, M. Yang, "An Evaluation of Code Generation of Dataflow Languages on Manycore Architectures", *Proceedings of* the 20th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2014), Chongqing, China, Aug. 20-22, 2014.
- [24] E. Gebrewahid, M. Yang, G. Cedersjö, Z. Ul-Abdin, V. Gaspes, J. W. Janneck, B. Svensson, "Realizing Efficient Execution of Dataflow Actors on Manycores", *Proceedings* of the the 12th International Conference on Embedded and Ubiquitous Computing (EUC 2014), Milan, Italy, Aug. 26-28, 2014.