

# Applications of wait/lock-free protocols to real-time systems

A8-9805; Håkan Sundell; 99-04-01; 80%; 00-06-30

## 1st year Progress Report

June 2000

### 1 Project plan

Our plan was to have two graduate students doing their doctorate research within the project, one at Mälardalens University at Västerås and one at Chalmers University at Göteborg.

More precisely the theme of the work that the two students had to get involved in was to explore research advances in the area of lock-free synchronisation and non-blocking concurrent data structures implementation and apply them for gaining in efficiency in the OS kernel level; the target is to develop an OS prototype, with a lock-free (non-blocking) kernel. The plan of the proposed work involves the steps below:

- First we planned to study some existing real-time operating systems and the architectures for which they are applicable for. The purpose of this study is twofold: (a) to identify the data structures whose implementation is a significant factor for the performance of the respective real-time OS (b) to identify the synchronisation capabilities (via the wait-free agreement protocols that they support) of the architectures on top of which these OS are running. These will lead to the identification of the most efficient feasible ways for lock-free implementation of these basic kernel data structures for the respective architectures.
- In the second step of this project we plan to gradually incorporate the new data structure implementations in the existing kernels and use simulation platforms for their testing and evaluation. We believe that it is very important not to build a new OS from scratch but modify already existing ones; in this way we will be able first to concentrate on the wait-free/lock-free aspects on real time OS and, second, to measure exactly at the end the improvements as we expect that the new techniques will give to the systems; third but not least the expected benefits of the OS will thus be immediately available to the community that already is using these OS.
- The last phase of the plan will be to test all aspects of the performance of the new kernel by running the kernel on the targeted real architecture with real data.
- The study and development will throughout the project be guided by case-studies. Initially we will concentrate on using wait-free snapshots in an automotive diagnostics system

On a more concrete level, the individual roles of the two (cooperating) graduate students will roughly be the following:

- The MdH student will focus on aspects related to the implementation of wait-free techniques in Real-Time kernels. This will, in addition to the actual implementation work, include evaluation and analysis of implementations and applications.
- The Chalmers student will focus on wait-free techniques and algorithms. This includes study of applications and existing kernels, development and analysis of algorithms, as well as evaluation of the use of algorithms.

## 2 Results

During the first year of the life of the project we tried to follow the basic lines of our plan. As a result of this effort:

- We managed to recruit two good students: B Alvin and H. Sundell. Unfortunately B. Alvin left the project some months after the project started. So the project had to run with only one student in its first year course.
- After initiating contacts with ENEA OSE systems and Northern Real-Time group we have received from ENEA OSE systems the sources of their RTOS and from Northern Real-Time group the TPK-sources.
- We have studied the above mentioned operating systems thoroughly, focusing at the synchronisation data structures that they use and the ways that these data structures are used inside the OS. We have identified a number of synchronisation data structures used by the above OS kernels whose implementations plays a significant role on the performance as well as the behaviour of the respective OS under faults. Our study was in-line with our expectations and more significantly gave us new interesting material for our research.
- We have looked at a number of shared data structures that are widely used in modern commercial OS (the ones that we got from our industrial partners included) aiming at developing efficient non-blocking implementations for them. We were quiet successful and we have proposed efficient non-blocking implementations for a number of fundamental shared data structures like buffers, queues and snapshot objects that outperform the best known non-blocking and lock-based algorithms in the literature. The performance of most of these new implementations was tested experimentally and computed analytically. Implementations of these shared data structures have run on top of existing multiprocessor systems that the ENEA OS can support. These results are described in our publications.

## 3 Changes

The big change in the project was that we had to run the project with one student instead of two. This did not change the basic plan of the project, we are still following the same basic steps described at section 1. But, we changed slightly the milestones of the first year basically to save some time since we are running with one student less. So we:

1. Have shown via our research papers the applicability of wait-free techniques in the real-time systems but we did not have time to write a paper to summarise these.

2. We have a number of implementations for the snapshot mechanism that run on top of existing multiprocessor systems (SUN ENTERPRISE10000).
3. We did not have time to use the wait-free monitoring in a diagnostic systems, but, we have used it as a debugging tool for debugging our parallel implementations.
4. We did not tried to find a wait-free algorithm for agreement in real-time systems because of lack of time (or one student) and because it was not necessary at this part of our project.

As we have mentioned above there were no changes at the basic steps of our plan, some small modifications at the milestones were made in-order to save time. We do not plan to change the plan for the next year; we hope that a new student is going to be recruited soon so that the project can start running in full scale.

## 4 Publication

1. B. ALLVIN, H. HANSSON, A. ERMEDAHL, M. PAPATRIANTAFILOU, H. SUNDELL, P. TSIGAS. Evaluating the Performance of Wait-Free Snapshots in Real-Time Systems. In Proc of the 1999 Konferens om Realtidssystem (SNART '99), Linköping, 24-25 August, 1999.
2. P. TSIGAS, Y. ZHANG. Non-blocking Data Sharing in Multiprocessor Real-Time System. In Proceedings of the 6th International Conference on Real-Time Computing Systems and Applications (RTCSA '99), part of the federated 1999 International Computer Congress (ICC '99), IEEE press, pp. 247-254.
3. P. TSIGAS AND Y. ZHANG. The Effect of Faults on the Performance of Non-blocking Shared Data Objects in Multiprocessor Systems. Fast Abstract in the Proceedings of the 1999 Pacific Rim International Symposium On Dependable Computing (PRDC 1999) part of the federated 1999 International Computer Congress (ICC'99).
4. H. SUNDELL, P. TSIGAS, AND Y. ZHANG. Bounding Wait-Free Snapshot Algorithm Buffers Using Timing Information. (Will be submitted)
5. P. TSIGAS AND Y. ZHANG. A fast, scalable non-blocking, concurrent FIFO queue and its performance evaluation on the SGI Origin2000 and SUN Enterprise10000. (Under Submission.)

## 5 Industrial cooperation

ENEAS OSE systems has offered to us the sources of their OS. Please see the attach letters from ENEAS OSE systems.

X-From\_: tsigas@cs.chalmers.se Tue Jun 13 17:17:18 2000  
Sender: tsigas@cs.chalmers.se  
Date: Tue, 13 Jun 2000 17:16:13 +0200  
From: Philippas Tsigas <tsigas@cs.chalmers.se>  
Organization: Chalmers Technical University  
X-Accept-Language: el, en  
To: artes@docs.uu.se  
CC: tsigas@cs.chalmers.se  
Subject: Our Report for the first year

Hi Roland,

Please find attached our first year report.

I am also sending to you by fax two letters from ENEA OSE systems that have to be attached to our report.

If something is not clear please do not hesitate to contact me.

Best regards,

Philippas

PS Thanks for your help with the deadline.



## Letter of Support

ENEAS OSE Systems AB expresses by this letter its support for the project:

***"Applications of wait/lock-free protocols to real-time systems"***

proposed to be performed within the framework of ARTES by Prof. Hans Hansson at Mälardalen University together with Drs. Marina Papatriantafillou and Philippos Tsigas at Chalmers University.

The project aims at developing and evaluating new synchronization primitives which have a potential of improving performance and predictability of Real-Time Operating Systems (such as our OSE family of Operating Systems). Even though these techniques are not yet mature enough to be included in commercial Operating Systems, we have a strong interest in following their development and if they prove successful consider including them in our products.

We intend to actively support the project by engaging in discussions, participation in its industrial reference group, and by (for experimental purposes) giving the project access to a relevant version of our OSE-platform.

A handwritten signature in black ink, appearing to read 'Lars Österberg', is written over a dotted line.

Lars Österberg  
President

ENEAS OSE Systems AB

Täby, Nov 30, 1998

MDH, Kopie

# CONFIDENTIAL INFORMATION NON-DISCLOSURE AGREEMENT

ko115e/OSE-99:253 R1.0

MÄLARDALENS HÖGSKOLA	
IDt	
990816	
12/99/2	
Handl.	Kopie
H.H.	P.T.

Between

Enea OSE Systems AB, P.O Box 232, S-182 23 TÄBY, Sweden  
(hereinafter referred to as "Sender")

and

Mälardalens Högskola, Chalmers Universitet & Uppsala Universitet (The Research Project called WARP), persons involved are: Philppas Tsigas, Per Håkan Sundell, Hans Hansson, Björn Allvin, Andreas Ermdahl and Marina Papatriantafilou.  
Representing the project is Hans Hansson, MRTC, Idt, Box 883, 721 23 VÄSTERÅS (e-mail han@idt.mdh.se).  
(hereinafter referred to as "Receiver")

the following agreement has been entered.

## 1. PREAMBLE

Whereas, Receiver has requested to receive certain and other information of a confidential or proprietary nature from Sender; and

Whereas, Sender considers this information confidential but is willing to provide such information on a confidential basis;

Now, therefor, the parties agree as follows:

## 2. DEFINITION OF "CONFIDENTIAL INFORMATION"

"Confidential information" means any and all unpublished information owned or controlled by Sender that relates to the technical, manufacturing, marketing, sales, or financial operations of Sender and that is not generally disclosed by Sender to the public, including without limitation proprietary processes and designs, trade secrets, know-how, inventions (whether or not patentable), formulas, technical drawings and data, research subjects, methods and results, proprietary computer software, unpublished product specifications and characteristics, manufacturing process, production techniques, plans for future products, business and marketing plans and strategy, products development plans, pricing policies, cost and profit information, customer lists, supplier identities and the like, whether disclosed orally, in writing or by inspection.

**Non-Disclosure Agreement ko115e/OSE-95:0 05 R1.0**

2 (3)

**3. NON-DISCLOSURE AND USE**

Receiver agrees:

- a) to use Confidential Information only for the purpose of evaluating possible future business arrangements between Receiver and Sender,
- b) to retain in confidence and all such Confidential Information, and
- c) not to disclose any such Confidential Information to anyone except employees or consultants of Receiver who are authorised to receive Confidential Information. Receiver shall, upon request by Sender, provide Sender with a list of persons who have access to Confidential Information from Sender.

**4. LIMITATIONS**

The obligations imposed by this Agreement on Receiver shall not apply to any Confidential Information that:

- a) was in the public domain at the time it was disclosed to Receiver;
- b) was known to Receiver at the time of disclosure as evidenced by documentation bearing a date prior to the date of such disclosure;
- c) is disclosed with the prior written approval of Sender;
- d) was independently developed by Sender without any use of the Confidential Information;
- e) becomes publicly known to Receiver from a source other than Sender without break of this Agreement.

**5. TERM**

The obligations of Receiver under this Agreement shall continue for two (2) years after the date of disclosure to Receiver or by the effective date of this Agreement, whichever is later in time.

**6. RETURN OF CONFIDENTIAL INFORMATION**

Upon request by Sender or upon completion of the business dealings relating to the Confidential Information, Receiver shall promptly return to Sender, or if requested by Sender, destroy all tangible material and copies thereof that disclose or relate to any of the Confidential Information and provide written confirmation of this by an officer of Receiver.

**7. GOVERNING LAW**

This Agreement shall be governed by the laws of Sweden. All disputes arising in connection with this Agreement shall be finally settled in Stockholm, Sweden. The official language in all proceedings concerning this Agreement shall be the Swedish language.

**8. ENTIRE AGREEMENT**

This Agreement contains the entire agreement between the parties as to the subject matter thereof and supersedes any previous or contemporaneous understandings or agreements, oral or written, as to the subject matter thereof.

IN WITNESS HEREOF, this Agreement has been duly signed by parties, hereto in duplicate.

Agreed to by:

Agreed to by:

Täby.....990624.....  
(Place and date)

.....990804.....  
(Place and date)

**SENDER**

**RECEIVER**

.....  
(Signature)

.....  
(Signature)

Enea OSE Systems AB  
Lars Österberg

Mälardalens Högskola  
Hans Hansson

Checked.....