Project no. 248828

ADVANCE

Strategic Research Partnership (STREP)
ASYNCHRONOUS AND **D**YNAMIC **V**IRTUALISATION THROUGH PERFORMANCE
ANALYSIS TO SUPPORT **C**ONCURRENCY **E**NGINEERING

First Market Analysis, Dissemination and Exploitation plan

D9

Due date of deliverable: January 31, 2011
Actual submission date: January 31, 2011(revised version February 16, 2011)

Start date of project: February 1^{st}, 2010

Type: Deliverable
WP number: WP9
Task number: WP9a-c

Responsible institution: SAP
Editor & and editor's address: Bernd Scheuermann
SAP Research
Research Center Karlsruhe
Vincenz-Prießnitz-Str. 1
76131 Karlsruhe, Germany

Version 2.0 / Last edited by Bernd Scheuermann / February 16, 2011

Project co-funded by the European Commission within the Seventh Framework Programme		
Dissemination Level		
PU	Public	√
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Revision history:

Version	Date	Authors	Institution	Section affected, comments
0.1	22/12/2010	Bernd Scheuermann	SAP	Initial version
1.0	31/01/2011	Raimund Kirner, Bernd Scheuermann	HERTS, SAP	Included market analysis, dissemination plan, exploitation plan
1.1	04/02/2011	Clemens Grelck, Heinz Hertlein, Jan Kuper, Raphael Poss, Bernd Scheuermann, Sven-Bodo Scholz, Marc Schrijver, Volkmar Wieser	BioID, HERTS, Philips, SAP, SCCH, Twente, UvA	Updated dissemination and exploitation plan
2.0	16/02/2011	Rob Albers, Koby Crammer, Clemens Grelck, Kevin Hammond, Chris Jesshope, Jan Kuper, Raphael Poss, Bernd Scheuermann, Sven-Bodo Scholz, Marc Schrijver, Volkmar Wieser	HERTS, Philips, SAP, SCCH, Technion, Twente, USTAN, UvA	Updates included

Tasks related to this deliverable:

Task No.	Task description	Partners involved°
WP9a	Dissemination activities	HERTS*, SAP, USTAN, TWENTE, UvA, PHILIPS, TECHNION, UCI, BioID, SCCH
WP9b	Market analysis, monitoring and strategy	HERTS, SAP*, USTAN, TWENTE, UvA, PHILIPS, TECHNION, UCI, BioID, SCCH
WP9c	Definition of exploitation plan for partners within the consortium	HERTS, SAP*, USTAN, TWENTE, UvA, PHILIPS, TECHNION, UCI, BioID, SCCH

°This task list may not be equivalent to the list of partners contributing as authors to the deliverable
*Task leader

Executive Summary

This document is the common result of the planning phases of the tasks WP9a, WP9b and WP9c on market analysis, dissemination and project exploitation. The market analysis identifies the relevant target market segment, describes market trends and carries out a market study comparing ADVANCE with other potentially competing products. On the dissemination side, we have identified the key communication messages that the project needs to bring across. We analyzed our communication target groups and the communication means that we intend to use and made a decision on which combinations of communication media and target groups we will focus. Our analysis shows that the web site is suitable for serving almost all information stakeholders. Therefore we concentrate our dissemination effort on this medium. Complementary to that, in order to also influence and educate the scientific community, we will publish results both at conferences and in scientific journals. An in-depth list of opportunities for this is presented. Additionally, we will engage in in-person events and other communication efforts that complement our dissemination strategy. Especially the industrial partners will also invest into internal communication, a prerequisite for efficient transfer of results. Regarding exploitation we clearly enumerate the foreseeable project results and explain how, both commonly as a consortium and individually by every partner, we intend to exploit them, creating impact to deliver a competitive advantage for the European Union. The academic partners can exploit the results for various internal projects, related and future research projects, and teaching. They plan to provide their part of the implementation as open source, allowing for easy and broad adoption by the community.

Contents

Chapter 1

Introduction

Work package 9 has the following overall objectives:

- to systematically analyse and to continuously monitor the market situation and trends
- to define a market strategy as well as a rollout plan
- to define individual exploitaton plans
- to create feedback and to recommend adjustements to the project plan if necessary
- to raise public awareness of the ADVANCE project.

This deliverable contains the first market analysis, dissemination and exploitation plan which are the outcome of the project planning phase. The deliverable is structured as follows. Chapter 2 presents the results of the first market analysis and the plan for forthcoming activities. In Chapter 3, we describe the general strategy of ADVANCE towards dissemination and talk about the concrete tools to fulfill this strategy. Chapter 4 describes our exploitation approach, introducing our general approach to exploitation, a brief description of common measures, and concrete partner strategies and already identifies some concrete transfer results and targets.

Chapter 2

First Market Analysis

This chapter presents the results of the first market analysis undertaken for project ADVANCE. First the goals are defined. Then the target market segment is identified, market trends are described and an in-depth market study is presented.

2.1 Goals

The goals of the project are re-iterated as they represent the foundation for all planning activities involved in the market analysis, dissemination and exploitation. Subsequently, the goals of the market analysis are explained.

2.1.1 Project Goals

The overall *aims* of the **ADVANCE** project are:

Aim 1: to produce a new software development approach that will allow multi-core hardware to be programmed easily and effectively, taking into account extra-functional requirements, such as the throughput of, or the average delays in, software components, subsystems and the entire system.

Aim 2: to apply *statistically-valid* techniques in order to better configure and adapt the software components and hardware platforms that comprise the computing system, under the guidance of these extra-functional requirements;

Aim 3: to develop methods and tools to support this approach, that are capable of taking into account both information provided by the applications programmer and information derived from hardware;

Aim 4: to use *hardware virtualisation* to apply this approach *generically* to a variety of multi-core/many-core architectures;

Aim 5: to demonstrate the validity of this new approach against a number of realistic applications.

The concrete *objectives* of the **ADVANCE** project are:

Objective 1: to develop language extensions that capture extra-functional requirements and expectations statistically;

Objective 2: to develop new cost-directed analyses and compiler transformations that statistically match programmer intentions and requirements with the hardware capabilities that are available;

Objective 3: to adopt new hardware virtualisation techniques that will allow multi-core/many-core hardware capabilities to be modelled at a high level of abstraction, but where the generated code can be specialised to execute efficiently on a full range of hardware platforms;

Objective 4: to develop new placement and mapping technology that will realize dynamic placement on the available multi-core/many-core hardware, based on the results of statistical cost-directed analysis;

Objective 5: to develop new feedback mechanisms that will inform the analysis, compilation and hardware virtualisation phases about actual application performance on the available hardware;

Objective 6: to validate the overall approach using a representative set of industrially-derived high-performance applications;

Objective 7: in order to facilitate the uptake of the incremental technologies developed by the project, to create a new user community with the help of which the new insights, techniques and experimental tools developed in the ADVANCE project can be disseminated effectively.

2.1.2 Goals of the Market Analysis

Consortium partners plan the application of the relevant results of this project in their activities or use them in order to enhance their worldwide competitiveness, to strengthen their product portfolios and to broaden their network of customers. The market analysis supports this plan and helps understanding the areas in which ADVANCE products and services can compete in the marketplace.

An ongoing monitoring of the market throughout the project lifetime shall guarantee that evolving market trends and changing customer requirements are considered. The first market analysis in this deliverable shall provide the foundation for these further market observations.

2.2 Target Market Segment

The market identified for the ADVANCE project results is the target segment of tools for multicore software development. Products in this segment are emerging,

and we anticipate that technologies in this area will evolve significantly. Currently, tool support for multicore software development lags behind the progress made in commercially available and forthcoming multicore system architectures. This gap needs to be closed to make efficient use of the capabilities of parallel multicore systems. In particular, ADVANCE targets the development of software for heterogeneous multicore systems which may also consider resources for high performance computing on GP GPUs and FPGAs. Providing tools for such systems is even more challenging as the product needs to be able to deal with the co-development of applications for completely different structures. The complexity involved in developing applications for heterogeneous devices should however be hidden through appropriate virtualization techniques and programming interfaces.

2.3 Market Trends

According to Moore's law the number of transistors on CPUs doubles every two years. Generally, an increasing number of transistors means higher compute power. Today this trends is still valid even though the clock speed of latest processors has experienced stagnation. Further increasing clock speed comes with a trade-off: Higher clock rates imply higher power consumption. Also cooling such devices requires increasingly more energy. The ratio between compute power per watt does not scale in a linear manner. This phenomenon is also referred to as the "power wall". Shortly, the additional amount of compute power is too expensive with respect to the electric power needed to achieve it.

A second trend can be observed in modern chip production technology which yields steadily smaller chip components through miniaturization. As a consequence more processors can be placed on the chip area which gave raise to the emergence of current multicore devices. The current strategy is to accommodate more small processors on one chip instead of increasing the speed of a single processor. Quad-core and Octa-core processors become standard in current desktops and laptops. Additionally, techniques like hyperthreading further increase the potential degree of parallelism as each core may schedule additional hardware-scheduled threads.

The implications for software developers are tremendous: They can no longer rely on clock speed evolution to accelerate the speed of their applications for single CPUs. Moreover, programming for a single processor was based on a deterministic computational model such that software could be understood and be debugged in a straight-forward manner. This has changed significantly. Typically, software programmed for single CPUs does not "automagically" adapt itself to the higher number of cores available. Developers themselves have to adapt their code, to parallelize it and to distribute its workload efficiently. Software developers need to handle the existence parallelism in order to implement new compute-intensive features and to prevent degrading performance.

The challenges involved with these market and technology trends have also

been analyzed by Gartner identifying the Seven Grand Challenges Facing IT published in 2008. Two of these challenges directly address the problems involved in parallel multicore software development. Gartner reports that *Parallel Programming* was required to make efficient use of multitude of cores available. The challenge, however, relied in dividing a problem into smaller sub-problems which were then to be executed on the individual processors. Consequently, Gartner outlines that the key issues were to effectively break up processes into specific subprocesses, to achieve an efficient multicore task scheduling and to design the architecture of the parallel processing environment. The other challenge addresses the need to *Increase Programmer Productivity 100-fold*. Gartner emphasizes that business' and society's demands for software development increased whereas the amount of qualified developers and software experts declined. Consequently, the productivity per programmer had to be improved.

The ADVANCE project tries to deal with the above mentioned trends and challenges by providing a tool chain and platform for multicore software development which shall foster industrial adoption of multicore programming as well as efficient multicore resource management and task scheduling.

2.4 Market Study

The market study arranges and compares the envisaged ADVANCE tool chain with a wide range of other products in the market of multicore software development tools. Data for the market study in this section is based on a survey carried out by the MWare project [3].

The following 18 tools participated in the survey:

- Acumem ThreadSpotter, Acumem
- Akka, Scalable Solutions
- Allinea DDT, Allinea
- DataRush, Pervasive Software
- Enterprise Architect, SparxSystems
- GPI, Fraunhofer ITWM
- IESE Simulink Partitioner, Fraunhofer IESE
- Intel Ct, Intel
- JProfiler, ej-technologies
- Offload, Codeplay Software
- PGI, The Portland Group

- Poly-Platform, PolyCore Software
- Prism, CriticalBlue
- PVS-Studio, Program Verification Systems
- SAMG, Fraunhofer SCAI
- Scalasca, Forschungszentrum Jülich
- Visual Studio 2010, Microsoft
- Zoom, RotateRight

2.4.1 Product Categories

Table 2.1 gives an overview of the various product categories, tries to assign the ADVANCE technologies accordingly and to position ADVANCE in the context of the competitor tools. The product categories are:

Profiling: Measuring runtime characteristics during program execution. Give an overview how much time was spent in various functions. Help identifying long-lasting function calls as basis for further improvements.

Debugging: Support the developer analyzing program execution, understanding the program flow and identifying the source of errors.

Tuning: Measuring runtime characteristics providing feedback on parallel performance. Use metrics like e.g. time for thread synchronization load balancing and scalability.

CASE/Modeling: Design tool to create the structure of a software. May include features for identifying parallelism.

Algorithm: Algorithm tools provide already parallelized code which can be included into the software being developed.

Runtime: Runtime tools provide programming libraries as higher level abstractions or primitives for parallel software.

Framework: Tools providing patterns for parallel algorithms. Developer fills-in missing functionality. Framework handles the execution of the parallel pattern configured.

Compiler: Compilers may provide automatic parallelization of algorithmic structures.

Library: Similar to algorithm tools. Libraries are more general and may offer parallelized versions for common operations (e.g. matrix multiplications).

Language: Programming languages may offer different abstractions for implementing parallel software.

Regarding its core functionality, ADVANCE has most overlaps with Microsoft Visual Studio (though following an entirely different concept). Apart from Visual Studio no other tool examined covers more product categories than ADVANCE.

Table 2.1: Product categories. Competitor data based on [3]. Legend: • = product category applies to given product, ⊙= ADVANCE may target product category, ⊕= ADVANCE may optionally target product category.

Vendor	**Product**	Profiling	Debugging	Tuning	CASE/Modeling	Algorithm	Runtime	Framework	Compiler	Library	Language
ADVANCE			⊕	⊙	⊙	⊕	⊕	⊕	⊙	⊙	⊙
Program Verification Systems Co. Ltd	PVS-Studio		•								
Allinea Software	Allinea DDT	•	•	•							
PolyCore Software	Poly-Platform	•		•	•		•	•			
CriticalBlue	Prism	•		•		•		•			
The Portland Group	PGI Suite	•	•	•			•		•		•
ej-technologies GmbH	JProfiler	•	•								
Pervasive Software	DataRush	•	•			•	•	•		•	•
Codeplay Software Ltd	Offload: CE		•				•		•	•	•
Microsoft	Visual Studio	•	•	•	•	•	•	•	•	•	•
Scalable Solutions AB	Akka						•	•	•	•	

Continued on next page

Continued from previous page

Vendor	Product	Profiling	Debugging	Tuning	CASE/Modeling	Algorithm	Runtime	Framework	Compiler	Library	Language
SparxSystems Central Europe	Enterprise Architect				•						
RotateRight LLC	Zoom	•		•							
German Research School for Simulation	Scalasca	•									
Acumem	ThreadSpotter	•		•							
Intel	Intel Ct						•	•	•	•	•
Fraunhofer SCAI	SAMG					•		•		•	
Fraunhofer IESE	IESE Simulink partitioner			•	•		•				
Fraunhofer ITWM	GPI						•			•	

2.4.2 Application Areas

Table 2.2 gives an overview of application areas and positions ADVANCE among various other competitor tools. Like the majority of other tools also ADVANCE may mainly target the areas on numerical simulations, technical applications and high performance computing. This can mainly be explained by the concept of stream processing used as modeling approach in ADVANCE. Many typical stream processing examples stem from these application areas. With some restrictions ADVANCE may also be suited in various office application scenarios. Overall ADVANCE shows most overlaps with competitor products like PCS-Studio, JProfiler, Visual Studio and Zoom.

Table 2.2: Application Areas. Competitor data based on [3]. Legend: • = product targets application area, ⊙= ADVANCE may target application area, ⊕= ADVANCE may optionally target application area.

Vendor	**Product**	Embedded Software	Networking Software	Numerical Simulations	Office Applications	Technical Applications	High Performance Computing	Computer Games
ADVANCE				⊙	⊕	⊙	⊙	
Program Verification Systems Co. Ltd	PVS-Studio			•	•	•	•	•
Allinea Software	Allinea DDT	•	•	•		•	•	•
PolyCore Software	Poly-Platform	•	•	•		•	•	
CriticalBlue	Prism	•	•	•	•	•		•
The Portland Group	PGI Suite			•		•	•	
ej-technologies GmbH	iProfiler		•	•	•	•	•	•
Pervasive Software	DataRush		•	•		•	•	
Codeplay Software Ltd	Offload: CE						•	•
Microsoft	Visual Studio	•	•	•	•	•	•	•
Scalable Solutions AB	Akka		•				•	•
SparxSystems Central Europe	Enterprise Architect	•			•	•		
RotateRight LLC	Zoom	•	•	•	•	•	•	•
German Research School for Simulation	Scalasca			•			•	
Acumem	ThreadSpotter	•	•	•		•	•	•
Intel Corporation	Intel Ct			•		•	•	•
Fraunhofer SCAI	SAMG			•		•	•	
Fraunhofer IESE	IESE Simulink partitioner	•						
Fraunhofer ITWM	GPI			•		•	•	

2.4.3 Analysis

Table 2.3 compares the analysis capabilities of ADVANCE with the relevant other competitor tools. The comparison is executed along a set of features:

Processor support: The more the analysis tools are aware of the hardware architecture used the better the tool can examine program behavior and the more accurate performance assessments can be.

Language support: Programming languages supported by the analysis tool.

Level: Level on which the analysis is done. Model analysis applies a computation model to identify conceptual errors. Static analysis examines the code with no need for compilation or execution. Communication analysis searches for inefficient communication patterns at runtime. Instruction level analysis examines the execution online and deliver most accurate results.

Restart before analysis: Describes how to set up the analysis and how to attach the tool to the running software. Restart before analysis means that the software is restarted in a special environment, where parameters can be set and specific dynamic libraries can be loaded.

Compilation before analysis: Signifies if a compilation is needed before starting the analysis.

Manual modification before analysis: Indicates if the code can be instrumented before the analysis (e.g. wrapping function calls).

Results are stored: Analysis tools storing results allow to compare information across multiple analysis procedures.

Presentation: Results can be presented in a textual or graphical manner. Additional evaluation features may be offered.

Availability: Indicates when the results are available.

Mapping: Expresses if the tool can map a result to the source code causing it.

Overhead: Overhead involved with the analysis. Should be as low as possible.

Limit: Limits of the analysis, e.g. with respect to the maximum number of threads that can be considered.

Table 2.3 shows that ADVANCE aims at supporting a wide range of processors and also GPUs and FPGAs. This versatility is unrivaled in the market. Generally, the spectrum of programming languages is not restricted with respect to the potential box languages (needed interfaces have to be provided). The specialization towards SAC and S-NET make ADVANCE unique in the market. With its specialization on these languages and on stream processing ADVANCE also targets

a niche where it may be able to attract a specific community. The level of analysis is comparable to many other competitors like Allinea, CriticalBlue, Codeplay, Microsoft and RotateRight. Concerning analysis preparation no functionality restrictions are predicted for the time being. Regarding the treatment of analysis results, the focus may be on research of statistical methods whereas an appealing presentation is out of scope for a research prototype. Textual presentation, which can be interfaced by other visualization techniques shall be sufficient. The range of features in this context would still be competitive in the market.

2.4.4 Tuner

Table 2.5 gives an overview of the tuning features of the various multicore software engineering products in the market and compares them with the envisaged functionality of ADVANCE. Tuning supports the developer and helps assessing and improving the performance of multicore software. The products are compared along the following topics:

Method: Indicates the profiling technique. Statistical sampling examines the program status in pre-defined time intervals. Instrumentation is used to extend the source code by profiling calls. Event-based profiling requires monitoring hardware on the processor to gain the profiling information.

Processor/core load: This metric assesses how the workload is scheduled on the available processors/cores available.

Memory access: Metric measuring the wait time for memory accesses.

Cache access: Metric measuring the wait time for cache accesses.

Thread communication: Metric measuring the time spent for thread communication, e.g. for thread synchronization.

Hardware performance counters: Additional measurements of rather low-level performance indicators.

Library support: Indicates of the tool provides specific profiling support for libraries used by the software.

The overview shows that ADVANCE will put less emphasis on measuring low-level hardware usage characteristics like memory access, cache access and thread communication. The focus is set on controlling and measuring the performance of the software in terms of the metrics latency, throughput and jitter. Statistical analysis and cost-directed control based on these metrics belong to the main research directions of the project and should differentiate ADVANCE as product from its competitors.

Table 2.3: Analysis. Competitor data based on [3]. Legend: •/○ = product does(/not) provide functionality, ⊙= ADVANCE may provide functionality, ⊕= ADVANCE may optionally provide functionality.

Product	**ADVANCE**	**PVS-Studio**	**DDT**	**Poly-Platform**	**Prism**	**PGI Suite**	**JProfiler**	**DataRush**	**Offload: Community Edition**	**Visual Studio 2010**	**Zoom**	**Scalasca**	**ThreadSpotter**	**IESE Simulink partitioner**
Vendor company	ADVANCE project	*Program Verification Systems Co. Ltd*	*Allinea*	*PolyCore Software*	*CriticalBlue*	*The Portland Group*	*ej-Technologies*	*Pervasive Software*	*Codeplay*	*Microsoft*	*RotateRight*	*Forschungszentrum Julich*	*Acumem*	*Fraunhofer IESE*
Processor support	x86, x86_64, Sparc, PowerPC, DEC alpha, NVidia GPUs, Xilinx FPGAs	x86, x86_64	x86, x86_64, ia64, Sparc, PowerPC, CeIIBE, NVidia GPUs	x86, x86_64, ARM, PowerPC, SF561, C55	x86, x86_64, ARM, PowerPC, MIPS32, Reneses SH4, Cavium cn, MIPS, NEC V850	x86, x86_64, NVidia CUDA-enabled devices	x86, x86_64, ia64, Sparc, PowerPC	x86, x86_64, ia64, Sparc, PowerPC, any Java 6 JVM	x86	x86, x86_64, ia64, ARM	x86, x86_64, ARM, PowerPC, Cell BE	x86, x86_64, ia64, Sparc, PowerPC, MIPS, NEC	x86, x86_64	x86, x86_64, ARM, FPGA-based processors
Language support	C, SAC, S-NET	C, C++	C, C++, Fortran, NET, CUDA	C	C, C++	C, C++, Fortran	Java	Pervasive Data-Rush Dataflow Language Analysis	C++	C++, C#, .NET, Python, DirectX Compute	C, C++, Fortran	C, C++, Fortran	C, C++, Fortran, any compiled language	Simulink / generic data-flow based languages
Analysis														
Level	Static analysis, Instruction	Static analysis	Static analysis, Communication, Instruction	Instruction	Static analysis, Communication, Instruction	Instruction	Instruction	Model, Static analysis, Communication	Static analysis, Instruction	Model, Static analysis, Communication, Instruction	Static analysis, Instruction	Communication, Function/routine, Load balance	Instruction, Memory access	Model, Static analysis
Preparation														
Restart before analysis	⊙	N/A	○	N/A	○	○	•	•	•	○	○	•	○	○
Compilation before analysis	⊙	N/A	•	•	○	○	○	○	○	○	○	•	○	○
Manual modification before analysis	⊙	N/A	•	○	○	○	○	○	○	○	○	•	○	○
Results														
Results are stored	⊙	•	•	•	•	•	•	○	N/A	•	•	•	•	○
Presentation	Textual (graphical optional)		Graphical	Graphical	Graphical	Textual, Graphical	Graphical	Graphical	Textual	Textual, Graphical	Textual, Graphical	Textual, Graphical	Textual, Graphical, HTML Output	Textual
Availability	During, After processing		During, After processing	After processing	After stopping	After processing	During, After stopping	During	After stopping	After stopping	After processing	After processing	After processing	After processing
Mapping	⊕	N/A	N/A	N/A	•	N/A	•	•	N/A	•	•	•	•	N/A
Overhead	Depends on analysis method	N/A	2-3%	N/A	Depends on method used	2%	N/A	Low	N/A	Depends on desired analysis details	2%	Depends on configuration / application	Depends on configuration, 10% for long running applications	None, analysis runs at compile time
Limit	No	N/A	No	N/A	No	256 Processes, 64 Threads per process	No	No	○	N/A	○	• Depends on available processors and memory	○	N/A

Table 2.5: Tuner. Competitor data based on [3]. Legend: •/○ = product does(/not) provide functionality, ⊙= ADVANCE may provide functionality, ⊕= ADVANCE may optionally provide functionality.

Product	**AD-VANCE**	**DDT**	**Poly-Plat-form**	**Prism**	**PGI Suite**	**Visual Stu-dio 2010**	**Zoom**	**Thread-Spotter**	**IESE Simu-link Parti-tioner**
Vendor company	AD-VANCE project	Allinea	Poly-Core Soft-ware	Criti-cal-Blue	The Port-land Group	Mi-cro-soft	Rotate-Right LLC	Acumem	Fraun-hofer IESE
Tuning									
Method	Instru-men-tation, statis-tical analy-sis	Sta-tis-tical sam-pling, In-stru-men-ta-tion,	In-stru-men-ta-tion	In-stru-men-ta-tion	Sta-tis-tical sam-pling, In-stru-men-ta-tion, Event-based	Sta-tis-tical sam-pling, In-stru-men-ta-tion, Event-based	Sta-tis-tical sam-pling, Event-based	Statis-tical sam-pling, Instru-menta-tion	Model check-ing, se-mantic anal-ysis, flow analy-sis
Characteristics									
Processor / Core load	⊙	•	○	•	•	•	•	○	•
Memory access	⊕	○	○	•	•	•	•	•	•
Cache ac-cess	⊕	•	○	•	•	○	•	•	•
Thread commu-nication	⊕	○	○	•	•	•	○	•	•
False sharing	⊕	○	○	•	○	○	○	•	○
Misc.									
HW Per-formance Counters	through-put, la-tency and jitter perfor-mance mea-sure-ments	•	○	N/A	•	•	•	○	○

Continued on next page

Continued from previous page

Product	**AD-VANCE**	**DDT**	**Poly-Platform**	**Prism**	**PGI Suite**	**Visual Studio 2010**	**Zoom**	**Thread-Spotter**	**IESE Simulink Partitioner**
Vendor company	AD-VANCE project	Allinea	Poly-Core Software	Critical-Blue	The Portland Group	Microsoft	Rotate-Right LLC	Acumem	Fraunhofer IESE
Library support	⊕(MPI, CUDA, SVP)	MPI	○	○	MPI, OpenMP, PGI Accelerator model	MPI, OpenMP, Intel Threading Building Blocks	○	○	○

2.4.5 Implementation

Tables 2.6 and 2.7 give an overview of the implementation aspects of ADVANCE and competing product in the market.

Concept: Indicates the programming concept supported by the tools. These concepts may include threads, message passing, dataflow concepts, functional programming concepts, asynchronous computation with callbacks, lock-free data structures and parallel patterns.

Memory model: Describes the memory model supported like shared memory, distributed memory or NUMA.

Programmer knowledge: The knowledge required from the programmer results from the implementation concepts applied.

Programming abstraction: Describes the level and type of abstraction for modeling parallelism. Typical abstractions include e.g. library calls, threads or parallel patterns.

Validation of correctness with standard tools:

Error prevention: Indicates the ability of the implementation tools to prevent parallelization errors. This can be achieved either by prevent the construction of such errors or by detecting them at compile or run time. Prevention mechanisms include: deadlock detection, race condition detection, transactional memory, lock-free data structures, functional semantics and dataflow semantics.

Help and documentation: Documents and services assisting the programmer to get acquainted with the tools or serving as reference guides.

Support: This category lists the means of support offered to the programmer including e.g. phone support, e-mail support, active communities, conferences or seminars and workshops.

With stream-processing ADVANCE follows a specific modeling concept which differs from all other tools examined. The focus is on functional programming which is also used by products like Poly-Platform, Prism, DataRush, Offload:CE, Visual Studio 2010, Akka and Simulink partitioner. With respect to asynchronous modeling there exist certain overlaps with Poly-Platform, DataRush, VisualStudio 2010, Akka, Simulink partitioner and GPI. Considering programming languages, concurrency engineering in ADVANCE is limited to S-NET, whereas algorithm engineering inside boxes provides support for C and SAC even though any other box language can be used if the required interfaces are implemented accordingly. Most other tools offer language support for C and C++ whereas FORTRAN and Java appear less frequent. Error prevention is mainly achieved at compile time but also by avoiding errors through stream-processing as the modeling concept applied. Help and support offered by ADVANCE is very extensive like the majority of other products. In addition, ADVANCE provides regular conferences, workshops and tutorials (S-Hack). The only other products providing support to a similar extend are PGI Suite and Intel Ct.

Table 2.6: Implementation part 1. Competitor data based on [3]. Legend: •/○ = product does(/not) provide functionality, ⊙= ADVANCE may provide functionality, ⊕= ADVANCE may optionally provide functionality.

Product	**ADVANCE**	**Poly-Platform**	**Prism**	**PGI Suite**	**DataRush**	**Offload: CE**	**Visual Studio 2010**
Vendor Company	ADVANCE project	PolyCore Software	CriticalBlue	The Portland Group	Pervasive Software	Codeplay Software Ltd	Microsoft
Concept	Functional programming, stream-processing, asynchronous	Threads, Message passing, Dataflow, Functional, Asynchronous	Threads, Dataflow, Functional, Lock-free data structures, Parallel patterns	Threads, Message passing, Parallel patterns	Dataflow, Functional, Asynchronous, Lock-free data structures, Parallel patterns	Dataflow, Functional	Threads, Functional, Asynchronous, Lock-free data structures, Parallel patterns
Memory model	Shared memory, Distributed memory	Shared memory, Distributed memory, NUMA	Shared memory	Shared memory, Distributed memory, NUMA	Shared memory, Distributed memory, NUMA	Shared memory, Distributed memory, NUMA	Shared memory
Programmer knowledge	Functional programming, concurrency modeling, stream processing	Functional programming		Thread programming	Dataflow programming, Parallel patterns	Functional programming	Thread programming
Programming abstraction	Stream processing networks, library calls, automatic parallelization	Library calls, Actors	Threads	Threads, Actors	Parallel patterns	Library calls	Threads, Tasks
Validation of correctness with standard tools	⊙	•	○	•	•	•	•

Continued on next page

Continued from previous page

Product	**ADVANCE**	**Poly-Platform**	**Prism**	**PGI Suite**	**DataRush**	**Offload: CE**	**Visual Studio 2010**
Vendor Company	ADVANCE project	PolyCore Software	CriticalBlue	The Portland Group	Pervasive Software	Codeplay Software Ltd	Microsoft
Special tool support	Performance modelling, SAC Compiler, graphical modeling, hardware virtualization and abstraction through SVP	Poly-Mapper validates the communications Topology	Prism has a set of verification features to help the programmer ensure that the parallel code is safe.	PGDBG parallel debugger	VisualVM plugin		Parallel Debugger
Programming language	SAC, C, S-NET (any other programming language with interfaces to S-NET)	C, C++	C, C++	C, C++, Fortran	Java	C, C++	C, C++, C#, VB.NET, F#, Python, Ruby
Error prevention							
Deadlock detection	N/A	○	○	○	•	○	○
Race condition detection	N/A	○	•	○	•	○	○
Transactional memory	○	○	○	○	○	○	•
Lock-free data structures	⊕	○	○	○	•	○	○
Functional semantics	⊕	○	○	○	○	○	○
Dataflow Semantics	⊕	○	○	○	•	○	○
Further		Thread safe message passing		Language compliance, Runtime error checking			
Help							
Getting started	Support of familiar programming languages, ease of modeling concurrency, APL-like structures in SAC	Familiar programming language, Familiar abstractions, Plugins for IDEs, Specialized IDEs	Familiar programming language, Familiar abstractions, Plugins for IDEs, Specialized IDEs	Familiar programming language, Familiar abstractions, Parallel Infrastructure, Plugins for IDEs	Familiar programming language, Familiar abstractions, Plugins for IDEs	Familiar programming language, Parallel infrascstructure, Plugins for IDEs	Familiar programming language, Familiar abstractions, Specialized IDEs
Tool support							
Documentation							
Meaningful API documentation	⊙		•	•	•	•	•
Beginner tutorials	⊙	•	•	•	•	•	•
Usage examples	⊙	•	•	•	•	•	•
Further	Reference manuals, language reference	Reference Manuals		Language reference, Compiler and Tools manuals			
Support							
On-site support	○	•	•	•	•	•	•
Phone support	⊕	•	•	•	•	○	•
Email support	⊙	•	•	•	•	•	•

Continued on next page

Continued from previous page

Product	**ADVANCE**	**Poly-Platform**	**Prism**	**PGI Suite**	**DataRush**	**Offload: CE**	**Visual Studio 2010**
Vendor Company	ADVANCE project	PolyCore Software	CriticalBlue	The Portland Group	Pervasive Software	Codeplay Software Ltd	Microsoft
Active community (website, wiki, mailing list, newsgroup, etc.)	⊙	•	•	•	•	•	•
Conferences	⊙	•	•	•	•	○	•
Seminars and workshops		•	•	•	•	○	•
Other/Further	Tutorials (e.g. S-Hack)			Webinars, Tutorials			Startup programs for partners

Table 2.7: Implementation part 2. Competitor data based on [3]. Legend: •/○ = product does(/not) provide functionality, ⊙= ADVANCE may provide functionality, ⊕= ADVANCE may optionally provide functionality.

Product	**Akka**	**Enterprise Architect**	**Intel Ct**	**SAMG**	**Simulink partitioner**	**GPI**
Vendor Company	Scalable Solutions AB	SparxSystems Central Europe	Intel Corporation	Fraunhofer SCAI	Fraunhofer IESE	Fraunhofer ITWM
Concept	Threads, Message passing, Dataflow, Functional, Asynchronous, Lock-free data structures, Parallel patterns			Threads	Threads, Message passing, Dataflow, Functional, Asynchronous, Lock-free data structures, Parallel patterns	Lock-free data structures, Parallel patterns, One sided communication, Asynchronous communication
Memory model	Shared memory, Distributed memory		Shared memory, Distributed memory, NUMA	Shared memory, Distributed memory, NUMA	Shared memory, Distributed memory, NUMA	Shared memory, Distributed memory, NUMA
Programmer knowledge			Functional, Parallel patterns, Metaprogramming concepts	Numerical modeling	Dataflow programming, Platform knowledge	Thread programming, Parallel patterns, Understanding of how to achieve maximal locality
Programming abstraction	Actors		Computation kernels, Parallel patterns	Library calls, Actors	Parallelism is derived from dataflow specification	Threads, Communication queues (channels)
Validation of correctness with standard tools	•		•	•	•	•
Special tool support			Data visualization tools are also provided			
Programming language	Java Scala		Is programming language, C, C++ through virtual machine API	C, C++, Fortran	Simulink, ASCET	C, C++, Fortran
Error prevention						
Deadlock detection	○	○	○	○	○	○
Race condition detection	○	○	○	○	•	○
Transactional memory	•	○	○	○	○	○
Lock-free data structures	○	○	○	○	•	○
Functional semantics	•	○	•	○	○	○
Dataflow Semantics	•	○	○	○	•	○
Further			Since race conditions cannot be written, they do not have to be debugged. Deterministic execution is guaranteed.			
Help						

Continued on next page

Continued from previous page

Product	**Akka**	**Enterprise Architect**	**Intel Ct**	**SAMG**	**Simulink partitioner**	**GPI**
Vendor Company	Scalable Solutions AB	SparxSystems Central Europe	Intel Corporation	Fraunhofer SCAI	Fraunhofer IESE	Fraunhofer ITWM
Getting started	Familiar programming language, Familiar abstractions		Familiar programming language, Familiar abstractions, Parallel infrastructure, Plugins for IDEs	Simple library interface, Demo programs	Specialized IDE	Familiar programming language, Familiar abstractions, Parallel infrastructure
Tool support	Standard Java profilers such		Intel(R) Thread Profiler	MPI tracing tools (e.g.	Simulink Real-time Workshop	gdb, vampir, ddt, valgrind
Documentation						
Meaningful API documentation	•		•			•
Beginner tutorials	•	○	•	•	○	•
Usage examples	•	○	•	•	○	•
Further	Complete reference				Internal reference documentation/reports	
Support						
On-site support	•	○	•	•	○	○
Phone support	•	○	•	•	○	○
Email support	•	○	•	•	○	•
Active community (website, wiki, mailing list, newsgroup, etc.)	•	○	•	○	○	○
Conferences	•	○	•	○	○	○
Seminars, workshops	•	○	•	•	○	○
Other/Further			Webinars, Tutorials		Product is intended for in-house use.	

Chapter 3

First Dissemination Plan

3.1 Dissemination Strategy

The objective of this chapter is to provide an overview of the communication activities within the ADVANCE project. In order to develop a communication plan that meets the needs of the ADVANCE project's internal stakeholders and target audiences, it is important to be aware of the project's aims and objectives. For an overview of the aims and objectives cf. to Section 2.1.1.

The following plan of communication activities within the ADVANCE project aims at establishing a detailed outline of dissemination activities; it defines the communication objectives, describes a communication strategy, identifies target audiences and the key messages to be delivered to these audiences, and provides an overview of the communication channels deemed the most credible and popular by the target audiences.

3.1.1 Core Communication Messages

The objective of future communication activities is to convey the following core communication messages to the identified audiences:

- ADVANCE focuses on creating tools facilitating software engineering for heterogeneous multicore systems supportd by statistical performance analysis.
- The goal of ADVANCE is to develop a new and advanced cost-directed hardware virtualisation technology to map programs onto emerging hardware architectures in a way that is both flexible and transparent to the programmer, but which, nevertheless, respects the programmers' expectations and requirements on extra-functional properties, such as resource utilisation or power consumption.
- Throughout the duration of ADVANCE, demonstrators and prototypes will be continuously developed and the results published

- ADVANCE is initiated in the 7th Framework Program of the EU
- partners from academia, research institutions and industry are engaged in ADVANCE

3.1.2 Communication Strategy

To implement the ADVANCE communication objectives, the following strategic focuses have been identified:

- Satisfy communication needs of all target audiences
- Ensure the visibility of ADVANCE results both internally and externally in the context of the framework program

The communication measures to realize the above mentioned strategies are:

- Creation of unique ADVANCE messages contributing to the ADVANCE image
- Identification of external and internal stakeholders and the definition of respective communication tools, channels and deliverables
- Exchange of communication knowledge and the agreement of the communication approach between ADVANCE participants
- Creation of appropriate marketing material and its subsequent distribution
- Organization of and participation in local and global events (general & scientific) to raise the awareness of the project
- Creation and maintenance of a project web site

ADVANCE partner communication activities are based on the top-down principle according to which all ADVANCE activities are to be aligned with the overall ADVANCE communication strategy.

3.1.3 Target Audiences

The key target audience of the communication strategy is divided into two categories.

3.1.3.1 External Target Groups

The following groups of stakeholders for communication activities have been identified:

General public: the general community of people concerned with project results

Special interest groups: future customers, users, decision-makers interested in the development of the project and its outcomes

Scientific community: universities and other research institutions from academia and industry concerned with the project's outcomes

Media representatives: journalists preparing press releases on the project

3.1.3.2 Internal Target Groups

The group of internal stakeholders consists of representatives from different levels:

ADVANCE partners from industry and academia

Company-internal stakeholders involved in the project's advancement or interested in its results.

Communication within these groups is a multi-level process: project results should be communicated within the individual internal stakeholders through the ADVANCE communication contacts as well as within partner organizations on a company-internal level, e.g., to company executive boards, all employees, or other interested groups.

3.1.4 Communication Channels

Communication channels refer to the medium used to convey communication messages to ex- ternal and internal target audiences. The identified communication channels are:

- Internet (ADVANCE website)
- Information material (demos, posters etc.)
- Scientific events
- Industry events, conferences and trade fairs
- Partner-internal events
- Face-to-face contact (presentation, briefings, workshops, etc.)
- Personal contact (including e-mail / telephone)
- Media

3.1.5 Mapping of Target Groups to Communication Channels

The communication plan for utilizing the identified channels and instruments of communication is presented in this section. The list may change during the development of the project and new channels and instruments may be added to the ones mentioned.

Table 3.1 gives a detailed picture of the possibilities of how project dissemination can approach target audiences through the mentioned communication channels. We may consider all opportunities described in the table, taking on an opportunistic approach: whenever there is an opportunity to address a certain target audience through one of the channels in an economically feasible way and given the human capacity of the project, we may use it. However, in order not to get lost in too many opportunities, we will prioritize the measures marked with yellow color. We have set our communication focus areas in a way that each target audience is addressed by at least one focus measure.

3.2 Concrete Envisioned Dissemination Tools

3.2.1 Project Website

The ADVANCE website is reachable at following address:
http://www.project-advance.eu. It provides relevant information about the project, its goals, ongoing work and of course all public results achieved. The web site will be continuously updated in order to provide news regarding issues of importance for the project and its partners.

3.2.2 Scientific Journals

Publications in journals in computer science are a prestigious way of publishing research results. As they often require in-depth discussion of a subject, publication in journals is rather relevant at the end of the project, when there is a enough result content that can be summarized in a journal publication. We intend to publish results in relevant journals. Below is a list of journals we may consider for paper submission.

- IEEE Software
- SPIE Journal of Electronic Imaging
- Elsevier Procedia Computer Science
- Springer International Journal of Parallel Programming
- Springer Real-Time Systems Journal
- Parallel Processing Letters (World Scientific Publishing)

Table 3.1: Mapping of communication stakeholders to communication channels

Target audience / Channel	General Public	Special interest group	Scientific community	Media representatives	Intra-organizational
Web Site including mailing list	Project overview text, related projects, videos	Details on SAC, S-NET, SVP, CAL, white papers	List of publications, related projects	Press contact and press releases, Press kit with images	n/a, internal information is confidential, cannot be added to web site
Information material	same information as web site, but in printed form (brochures, flyers, ...), also referencing web site				Internal white papers showing alignment with portfolio and strategy
Scientific events and publication portals	Posters, Journal papers	Tutorial session, summer schools, discussion panels	Conference / workshop papers, demos, tutorials	Interviews, demo / poster walk throughs	Present project concepts to managment staff
Industry events, conferences and trade fares	Demos, Quizzes, Talks	Tech talks, expert sessions, panel discussions	same as scientific events		
Partner internal events	these stakeholders do not have access to internal events				Internal presentations, talks, demos showing alignment with portfolio and strategy to colleagues, lines of business, and management
Face-to-face contact (presentation, briefings, workshops, etc.)	Evangelize project topics, goals, and approach, hand-out information material	Explain details, offer collaboration, invite to events, hand out information material, offer tutorials, workshops	Explain details, offer collaboration, invite to events, hand out information material, offer tutorials, workshops, include topics in university curricula	Interviews, demo / poster walk throughs	Evangelize project topics, goals, and approach, hand out information material, justify investment in project, show alignment to portfolio and strategy
Media (TV broadcasts, print & electronic)	Videos	Videos, implementation documentation	Journal articles	n/a	Internal videos, blog articles

- Journal of Embedded Systems
- ACM Transactions on Programming Languages and Systems, Software, Practice & Experience
- Journal of Functional Programming
- IEEE Transactions on Audio, Speech and Language Processing
- IEEE Transactions on Pattern Analysis and Machine Intelligence
- IEEE Transactions on Image Processing
- IEEE Transaction on Software Engineering
- Wiley & Sons Concurrency and Computation: Practice and Experience
- Journal of Machine Learning Research
- Machine Learning Journal

3.2.3 In-Person Events

In-person events comprise conferences, workshops, tutorials, third-party industry events, trade fairs, public co-organized events and internal events.

3.2.3.1 Conferences and Workshops

We will publish results at pertinent scientific venues. In an early phase of the project we target workshops for the publication of concepts and architectures. As soon as evaluation results become available, we will focus on scientific conferences. Below is a list of venues we intend to consider for paper submission.

- Conferences:
  - ACM/IEEE International Conference on Software Engineering
  - International Euro-Par Conference on Parallel Processing
  - IEEE International Parallel & Distributed Processing Symposium (IPDPS)
  - IS&T/SPIE Electronic Imaging
  - SAC Developers' Conference
  - Quality Control and Artificial Vision
  - SPICE Conference International Conference on Software Process Improvement and Capability Determination
  - International Forum for Mechatronic
  - Evocop European Conference on Evolutionary Computation in Combinatorial Optimisation

- ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)
- SIREN/NL (national meeting for curiosity driven research on ICT in the Netherlands)
- International Conference on Parallel Processing
- IEEE/ACM International Conference on Cluster, Cloud and Grid Computing (CCGrid)
- ACM International Conference on Computing Frontiers (CF)
- International Symposium on Implementation and Application of Functional Languages (IFL)
- IEEE International Parallel and Distributed Processing Symposium (IPDPS)
- IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)
- IEEE/ACM/IFIP International Conference on Hardware-Software Codesign and System Synthesis (CODES+ISSS)
- International Conference on Dependable Systems and Networks (DSN)
- International Conference on Software Engineering Research, Management and Applications (SERA)
- ACM Conference on Languages, Compilers and Tools for Embedded Systems (LCTES)
- International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES)
- ACM International Conference on Embedded Software (EMSOFT)
- ACM International Conference on Functional Programming Languages (ICFP)
- Symposium on Principles of Programming Languages (POPL)
- International Symposium on Trends in Functional Programming (TFP)
- International Symposium on Static Analysis (SAS)
- Symposium on Logic in Computer Science (LICS)
- European Symposium on Programming (ESOP)
- ACM Conference on Embedded Softare Systems (EMSOFT)
- IEEE Conference on Industrial Informatics (INDIN)
- IEEE Symposium on Field and Custom Computing Machines (FCCM)
- IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- International Conference on Machine Learning (ICML)
- Neural Information Processing Systems

- Workshops:
  - Workshop on Compilers for Parallel Computing (CPC)
  - International Workshop on Multicore Software Engineering
  - Workshop on Statistical and Machine learning approaches to ARchitecture and compilaTion (SMART)
  - Workshop on Streaming Systems
  - ACM SIGPLAN Workshop on Advances in Message Passing(AMP)
  - Declarative Aspects of Multicore Programming (DAMP)

3.2.3.2 Tutorials

To gain adoption of the ADVANCE technology it will be fundamental to present the results of the project didactically, by illustrating the concepts, as well as the use, of the ADVANCE tool chain and platform. To this end, we plan to regularly organize S-Hack events, a comprehensive tutorial, which will also include a hands-on session. In the tutorial, we will teach the attendees on the different constituents of the ADVANCE tool chain. The attendees will gain the knowledge and practical experience required to use ADVANCE in their own projects. The tutorial will target students and researchers working in the domain of multicore and high performance computing. So far it has been organised three times: April 2008 and March 2010 in Hatfield UK, September 2010 in Novosibirsk (RUS).

Furthermore it is planned to submit a tutorial proposal to the International Supercomputing Conference (ISC 2011). Working title: SAC and its Autoparallelising Compiler Technology.

3.2.3.3 Third-party Industry Events and Trade Fairs

The commercial and academic project parters will exploit the project to demonstrate their thought leadership in the area of business-relevant third-party industry events and trade fairs. Table 3.3 gives an overview of potential events, the contributing partner and the envisaged time frame.

Table 3.3: Third-party Industry Events and Trade Fairs

Partner	Event	Planned contribution	When
BioID	Die Lange Nacht der Wissenschaften - demonstrates academic and industry research to the general public	presentation or poster	2011
SCCH	Lange Nacht der Forschung - demonstrates academic and industry research to the general public	presentation or poster	2011

Continued on next page

Continued from previous page

Partner	Event	Planned contribution	When
Twente	STARS Kick-off-meeting	presentation of ADVANCE objectives	2010
Twente	Dutch Functional Programming Day	presentation of ADVANCE objectives	2010
UvA	Central European Functional Programming Summer School (CEFP)	lecture on SAC technology	2011
UvA	Intel European Research and Innovation Conference (ERIC)	poster presentations	2011

3.2.3.4 Public Events Co-organized by Project Partners

The project partners organize many events where they invite the general public or invited guests to inform them about new developments. Table 3.4 gives an overview of potential events, the (co-)organizing partner, and sets a time frame.

Table 3.4: Public events co-organized by project partners

Partner	Event	Planned contribution	When
HERTS, UvA	S-Hack - Practical training on SAC and S-NET	tutorial	2010, 2011, 2012, 2013
HERTS, UvA	SAC Developers' Conference	Workshop and talks	2010, 2011, 2012, 2013
HERTS, UvA	Stream processing and programming workshop (Novosibirsk)	Talks and demos	2010
SAP	SAPPHIRE NOW, SAP's customer and user conference	expert session or demonstrator	2012/2013
SAP	SAP TechEd, annual conference on SAP system management and custom development	expert session or demonstrator	2013
SCCH	Presentation on high performance computing at SCCH board (gremium consisting of about 30 companies and 5 acadamic partners)	Presentation	2012
SCCH	Technokontakte - Effizienzsteigerung durch intelligentes Datenmanagement	expert session and talk	2011

Continued on next page

Continued from previous page

Partner	Event	Planned contribution	When
UvA	Advanced School on Concurrency and Imaging (ASCI) – a Dutch PhD School	tutorials on SAC and SVP	2010, 2012, 2014
UvA	Advanced School on Concurrency and Imaging (ASCI) – a Dutch PhD School	lecture on SAC	2011

3.2.3.5 Important Internal Events of Project Partners

In order to disseminate the ideas of the project and to ensure uptake within the partner organizations, internal events play a vital role. Table 3.5 gives an overview of the potential internal events of the partners and what the contribution of the project to these events could be.

Table 3.5: Important internal events of project partners

Partner	Event	Planned contribution	When
BioID	Internal meetings of the development team to work out concept and implementation strategies of BioID's use case	Biometric optimization framework as first use case	2010
BioID	Workshops to internally discuss further development of the biometric test and optimization framework, implemented with the ADVANCE tools	Maximization of both exploitation opportunities as well as relevance of the contribution to reach the ADVANCE goals	2011, 2012
HERTS, UvA	Development week - refine hardware virtualization and software interfaces	workshop	2010, 2011
HERTS	Training of graduate students and PhD canditates	technical presentations of the ADVANCE technology and interactive discussions	2010, 2011, 2012
HERTS	Presentation of research progress	technical presentation of the students' current research challenges and solutions	2011, 2012, 2013
Philips	Internal meetings	project presentations	2010, 2011, 2012, 2013

Continued on next page

Continued from previous page

Partner	Event	Planned contribution	When
SAP	DKOM - internal annual developer and architecture conference	expert session or demonstrator	2012/2013
SAP	Research Summit - annual summit of the whole research department	poster, session or demonstrator	2012/2013
SCCH	Advances in Knowledge-Based Technologies - PhD-Seminar providing overview of current topics	expert session	2011
Technion	Internal Technion meetings	Project presentation	2011, 2012, 2013
Twente	Internal chair meetings of CAES (Computer Architectures for Embedded Systems)	project presentation	2010
Twente	Strategic Research Orientation - cross-group event for knowledge exchange on research topics	project presentation	2011
UvA	Computer Science - course at UvA	lectures, students use ADVANCE technology	2010
UvA	Institute Seminar - staff meeting	talk on ADVANCE	2010
UvA	Technical training	training introducing new graduate students to ADVANCE technology	2010, 2011
UvA	ASCI graduate school	lecture and practicals about ADVANCE-related technology	2011

3.2.4 Alliances and Networks of Excellence

SAP is part of the European Institute of Technology (EIT). This offers the opportunity to disseminate the results of ADVANCE in the context of this large-scale initiative, which is supposed to set the landscape of research and innovation in ICT for the next decade.

University of Hertfordshire is a member of the ARTEMIS Industry Association (ARTEMISIA), whose mission is to continue the work on the ARTEMIS European Technology Platform. This platform has many members from the academic and industrial sector, and its networking events offer opportunities to disseminate the mission and results of the ADVANCE project.

Several partners in the Consortium (e.g. SAP, UvA) are members of the Eu-

ropean Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC). In this context, partners are engaged in fostering the cooperation in the area of high-performance and embedded computing systems which shall strengthen joint research between academia and industry and between computer architects and tool builders.

USTAN is a major contributor to IFIP Working Group 2.11 (Generative Programming), of which Kevin Hammond is a founder member. USTAN is also a member of the DECOS Interest Group on Dependable Embedded Components and Systems, of the UK's Semi-Autonomous Systems Defence Technology Consortium, of the Autonomous Systems Network and of the HPC Advisory Council Network of Expertise.

Chapter 4

First Exploitation Plan

4.1 Exploitation Strategy

The ADVANCE consortium includes both academic and industrial partners which will exploit the results in different ways. The project results will strengthen European competitiveness both from a commercial and academic perspective.

4.1.1 General Approach

Exploitation is recognized as the key enabler for the success of the ADVANCE project. Hence all partners of this project are aware of and committed to the exploitation of the project results. It is the principle of all exploitation activities to use research results to create value within all participating organizations and thus to improve their competitive advantages. Only by scaling up the results into commercial offerings, all European Commission intended impact/goals can be reached while ensuring profitability through economies of scale. Wherever possible, research results will be exploited for the internal development and support of new products and services. These products and services will lead to a competitive advantage of the participating organizations and will substantially contribute to the benefit of the targeted users. In order for the exploitation to be effective, an integrated approach will be necessary, combining experience and expertise from the development department and solution management, and the involvement of a user base represented by the consortium partners. An integral part of this exploitation approach is the identification of use cases/scenarios which will serve as the validation point throughout the project. To safeguard and incubate the value of the project results, the following non-exhaustive list of actions will be taken:

- Organization of transfer activities of research results into development, product, or service organizations of the industrial partners.
- Continuous analysis of transfer opportunities, adjusting the project when necessary in order to ensure the best possible outcome.

- Investigation into the possible economic benefits and impact of the expected research results.

All consortium partners will develop their own exploitation plan throughout the project; the final version could serve as input to their product & solution management organization and enable them to develop a first business case for their development decisions.

4.1.2 Overview of Exploitable Results

Table 4.1 summarizes the exploitable knowledge respectively the exploitable results that are generated in the ADVANCE project. With all partners working from a common knowledge and result base, we ensure that all ADVANCE exploitation activities follow a common vision.

4.2 Partner Exploitation Strategies

4.2.1 Individual Exploitation Strategy for Academic Partner HERTS

The University of Hertfordshire was awarded "Entrepreneurial University of the Year 2010" by the Times Higher Education. HERTS is the UK's leading business-facing University, committed to improving the performance and competitiveness of organisations in a wide range of sectors. It is innovative and enterprising and challenges individuals and organisations to excel. The University of Hertfordshire is one of the region's largest employers with over 2,600 staff and a turnover of almost GBP 235 million. With a student community of over 27,500 including more than 2,000 international students from over eighty five different countries, the university has a global network of over 200,000 alumni. HERTS offers business services such as R&D and technology transfer. For more information, please visit www.herts.ac.uk.

4.2.1.1 General Exploitation Strategy

As an academic institution, the primary exploitation activities of University of Hertfordshire are centered around the dissemination of research results, continuation of research activities, and technological transfer of research results.

4.2.1.2 Transfer Opportunities

HERTS has a wealth of experience on holding user community workshops on S-Net and SAC, the key technologies provided by HERTS within ADVANCE. Within the course of the ADVANCE the first such user community workshop has been held in March 2010. Since then, several individual meetings with industrial partners have been held as well, in order to help in exploiting these technologies for their

Table 4.1: Key Results of ADVANCE and their Exploitation Potential

Exploitable Knowledge	**Exploitable for which Products and Activities**	**Sectors of Application**	**Partners Involved**
Modeling the concurrency of stream-based applications	Opportunity to use modeling concept for own SAP products (e.g. SAP SCM or TREX). Application to stream-based sub-systems in BioID's biometric applications, e.g. in the biometric test and optimization framework	Applicable to any industry sector with stream-oriented applications and strong performance requirements	BioID, SAP
Aggregation of (extra-)functional properties	Usage of the CAL language and property aggregation framework in other research projects and commercial application	Applicable to any industry sector were resource-aware composition of components is needed	HERTS, USTAN
Integrating legacy software modules into an overall application that is written in a language supported by ADVANCE (e.g., S-NET or SaC) in order to efficiently run on heterogeneous, multi-core platforms	Biometric test and optimization framework, under development by BioID	Any industry sector that uses legacy software which is expensive or difficult to port completely to an ADVANCE-supported language. Instead, a trade-off between development effort and computational efficiency is achieved by incorporating existing software modules	BioID
Getting familiar with state-of-the-art hardware virtualisation and parallelisation techniques, and modeling of concurrency of the whole processing pipeline in the context of quality inspection.	Concepts are basis for long-term strategy facing the requirements in terms of performance, heterogeneity of hardware platforms as well as productivity of the whole software engineering.	Machine learning based computational models for quality monitoring and inspection systems, in particular high-performance vision inspection for industrial quality control based on optical sensors	SCCH
Statistically-based performance models	Any software system with quality-of-service or performance prediction requirements	Real-time applications requiring QoS guarantees; streaming applications, e.g. video-processing; general high-performance computing applications including multicore or other parallel architectures; computer networks; systems with strong user interaction requirements, including medical imaging, computer games; embedded systems and devices	USTAN, HERTS, Twente
Statistical Analysis of primary data using linear models	Deep understanding of the statistical properties of the data, Base line for the ability to predict behavior of some threads given information on other threads.	Applications with a need to automatically allocate resources based on their properties	Technion and Twente
investigate performance characteristics, and performance predictability for applications based on combinations of stream-based and data-dependent image processing algorithms and corresponding middleware solutions	The results of the project will be used by Philips in the next generation high-end X-Ray imaging systems to provide added clinical value. After that the same will be exploited in the value market to achieve cost-reduction.	Applicable to industrial sector with stream and data-dependent imaging applications such as X-ray imaging, video surveillance and consumer multimedia	Philips, Twente

applications. There are further user community workshops planned, giving the chance of demonstrating the latest achievements within the ADVANCE project.

The tools are available for free download on the internet from the individual web sites of S-Net [1] and Sac [2].

4.2.1.3 Planned Exploitation Activities

University of Hertfordshire will develop new technologies around the S-Net approach. Further, in similar way as the S-Net technology builds upon previous research project (most prominently, the FP6 project Æther), we anticipate to exploit the results of the ADVANCE project to inspire future grant applications and research projects on both European and national level.

4.2.1.4 Contribution to Standards

The *Constraint Aggregation Language* (CAL) is under development within the project ADVANCE [2]. The aim of CAL is to provide a generic declaration language of functional and extra-functional properties of software components. First of all, the language has to prove its effectiveness in practical use, to be done with the project ADVANCE. Given that the evaluations will be quite satisfactory, we believe that CAL has a good position to be adapted by the real-time community to serve as the foundation for defining a common annotation language. The benefits of such a common annotation language have been already raised, especially in the timing analysis domain [1].

4.2.2 Individual Exploitation Strategy for Academic Partner USTAN

4.2.2.1 General Exploitation Strategy

USTAN maintains an active Knowledge Transfer Centre with connections to industrial, commercial and governmental sources `http://www.st-andrews.ac.uk/ktc/` and an active *Technology Alert* system that alerts investors, industrialists and other potential exploiters to newly developed technology. One of the university's aims is to promote the discovery, development and application of knowledge through high quality research. During this process, we will extend current knowledge and produce innovative and novel concepts, which may give rise to commercially viable inventions and solutions.

The university will then decide how best to exploit technology, either through collaborative development with industry, by licensing or spin-out company formation. This gives inventors the opportunity to participate in developing commercial products or services and take them to the market. The creation of new spin-out companies benefits local economic development and can create many new, high-tech jobs in the region.

[1] S-Net: `http://www.snet-home.org`
[2] SaC: `http://www.sac-home.org`

4.2.2.2 Transfer Opportunities

We foresee transfer opportunities arising primarily through the development of new statistically-valid performance models and associated analyses that we will develop in WP4; through the feedback mechanisms that we will develop in WP5; and through the annotation and constraint-exposure mechanisms that we will develop in WP2. The models and analyses will assist the development of high-performance parallel computer systems, partcularly the new and increasingly important class of heterogeneous multicore/manycore architectures. We will work closely with the Advance industrial partners in determining how our results may be deployed in their areas of expertise.

4.2.2.3 Planned Exploitation Activities

At this early stage of the project, we have not yet developed directly exploitable intellectual property. We will monitor our production of research results, systems and deliverables, and highlight amy clearly exploitable results to our commercial partners. We will also make our Knowledge Transfer Centre aware of any potentially exploitable results. Exploitation may also occur through reuse of or extension of research that is developed in this project in other EU or national research projects. We will actively seek to use the results of the Advance project in this way. We will also actively work with the industrial partners to extend knowledge and appreciation of our research results in the wider community through joint research papers, posters/presentations/demonstrations at industrial conferences and trade events etc. as appropriate.

4.2.3 Individual Exploitation Strategy for Academic Partner Twente

4.2.3.1 General Exploitation Strategy

The University of Twente is a young, entrepreneurial university. It sets standards in the field of new technology and seeks to stimulate change, renewal and progress in society. Our strength lies in our capacity to combine. We work with the technologies of the future - information technology, biotechnology and nanotechnology - in which behavioral and social science research play a vital role. After all, the most interesting and relevant innovation takes place at the interface between technology and its implications for mankind and society. We are active in areas such as health, water, sustainability, security and education.

Together, 3,300 scientists and professionals carry out ground-breaking research, bring about socially relevant innovation, and provide inspiring teaching for more than 9,000 students. To us, entrepreneurship comes as second nature. The campus is home to around 100 businesses, including student-run businesses. The University of Twente has also generated more than 700 successful spin-off companies. The university's business park, Kennispark Twente, encourages and assists entrepreneurs to start new companies.

Our research is highly regarded at both the national and international levels, and is accommodated within six very active research institutes. These focus on nanotechnology, information technology, biomedical technology and technical medicine, sustainable energy and smart devices, governance, behavioral sciences and geo-information science/earth observation. The research institutes combine scientific excellence with a sharp eye for knowledge valorization (the process of translating newly developed knowledge into economic activity) and social applications. They are highly successful in generating spin-off businesses.

Science depends on developing talent. The large number of research environments at the University of Twente make it an attractive destination for researchers. There are opportunities for junior and senior resRecore Systems, Smart Sign Solutions, Bibix and Homa SoftwareRecore Systems, Smart Sign Solutions, Bibix and Homa Softwareearchers alike. Twente's research output and citation rates are well above the national average.

4.2.3.2 Transfer Opportunities

The transfer opportunities for the university of Twente are found in the many contacts w ehave with industry. First of all there are the intimate relationships with the many spin-off companies such as Recore Systems, Smart Sign Solutions, Bibix and Homa Software. These companies, each in their own specility, all work on the topic of distributed systems, multi-core archtiectures, and concurrent computation. They are in big need of solutions to the fundamental probelsm in this area and they may benefit from the results that Advance produces. Secondly, there is a close cooperation with companies such as Thales Netherlands, NXP, Océ, where the same topics are deeply relevant and information on the work in Advance will naturally flow to the contact persons within these companies during the many meetings that are organised on various topics.

4.2.3.3 Planned Exploitation Activities

So far the planning of concrete exploitation acitivities is not relevant yet, as there are no concrete results yet. But in general, as history showed, the University of Twente as a solid starting point for spin-off companies and when the Advance project is successful, UT will certainly consider to create possibilities for those who see possibilities in bringing the results of Advance to market. Furthermore, results will be incorporated into courses that are taken care of by the chair CAES, such as a course on hardware specification.

4.2.4 Individual Exploitation Strategy for Academic Partner UvA

4.2.4.1 General Exploitation Strategy

UvA is a research organisation and the Computer Systems Architecture group has a breadth of experience in defining and prototyping architecture and system level

support for concurrent systems based on heterogeneous multi-core architecture. Our expertise is in identifying appropriate models to separate the concerns between the specification and exploitation of concurrency in applications and to provide flexible and dynamic mapping and scheduling of that concurrency. This has the positive advantage of making concurrent programming generic providing that the concurrency can be adapted by the run-time system, be it in hardware or software. We are aware that this approach is not so easy as programming concurrency for a specific architecture and hence our aim is to demonstrate the benefits of this to various partners and other contacts. Our overall strategy is to demonstrate this and to have these stake holders adopt the technology. To this end we will disseminate, where possible, results developed in the project under open-source licences. This will achieve the widest possible impact in the community and, as this is long term research, will help us in building a community, which can multiply the investment made in this project by the EU.

4.2.4.2 Transfer Opportunities

Our transfer opportunities, in addition to adoption by the partners has been to build relationships with various multi-core vendors, to obtain their platforms and demonstrate our developments on their platforms. To this end we have a relationship with both Intel and Oracle and have been donated equipment from these sources to investigate this opportunity.

4.2.4.3 Planned Exploitation Activities

As we have no exploitable results to date we have not considered any exploitation activities. Our goal has been in the definition of the interfaces that will maximise the opportunities for optimisation across a range of platforms from conventional multi-cores, for example from the vendors above, through to GPGPUs and other accelerator-based architectures.

4.2.4.4 Contribution to Standards

If this work is successful there will be the potential for contributing to new standards for the capture of concurrency from application-level compilers, which can then be mapped to a range of different architectures. To date there are no such compilers but it is clear from our work to date that such standards are critically lacking, making the development of compilers for a range of platforms expensive and the code generated platform specific.

4.2.5 Individual Exploitation Strategy for Industrial Partner Philips

4.2.5.1 Business Model

Philips Healthcare is a leading supplier of healthcare products and services that deliver value throughout the complete cycle of care - from disease prevention to screening & diagnosis, to treatment, health management & monitoring - in key areas including cardiology, oncology, critical care and women's health. Philips Healthcare is active in the following businesses:

- Home healthcare Respiratory care, Medical alert services, Remote patient monitoring
- Imaging systems X-Ray imaging, Computed Tomography, Magnetic Resonance Imaging, Nuclear imaging
- Clinical care systems Ultrasound, Defibrillators, Electrocardiography Information RIS (radiology information system), PACS (picture archiving and communication system), Clinical imaging applications, Patient monitoring.

One of the Business Units within Philips Healthcare develops and produces high end cardio vascular X-Ray systems for interventional users. During an intervention, a physician acquires X-Ray images to visualize a catheter inside the arteries or heart of the patient, and relies on these X-Ray images for navigating the catheter, and for treatment of pathologies.

4.2.5.2 General Exploitation Strategy

Philips Healthcare plans to investigate performance characteristics, and performance predictability for the image processing algorithms and middleware. Philips Healthcare contributes to project Advance on use cases, requirements and evaluation and will consult the whole consortium from the perspective of a major industrial player. In particular, it will bring in their CV X-Ray system as a use case for Advance.

4.2.5.3 Transfer Opportunities

Philips brings in an application on X-Ray image processing. A cardio vascular X-Ray system is a system that is used for performing interventions: a physician is maneuvering a catheter or a needle inside a patient, while the X-Ray images (acquired at frame rates from 15 till 60 fps) give him the necessary high-quality visual feedback. For the required hand-eye coordination it is important that the entire image processing chain has a maximum latency of 110 msec. This is the time between the moments an X-Ray pulse has been generated, until the X-Ray image is displayed on the monitor. And the variation in latency (jitter) is to be within strict bounds as well. When designing such a low-latency and jitter-free

image processing system on a PC platform, it is important to be able to predict the performance, and to be able to influence those characteristics. This is especially true in cases where the algorithms themselves have non-deterministic performance behavior (i.e. the number of operations per pixel depends on the pixel contents), or when multiple image pipelines run concurrently.

Philips is considering heterogeneous multi-core hardware as platform for all future image processing. Multi-core technologies in theory provide Philips with the computation power to introduce the next generation image processing algorithms in its X-Ray products without cost price increase. The next generation image processing algorithms enable new minimal invasive interventions for diseases that can now only be treated by traditional surgery. In particular, minimal invasive treatment will reduce costs for the society through:

- Faster patient recovery reduces the length and costs of the stay in hospital
- Better well-being of the patients, enabling fast return to normal (professional) activity
- Early detection of (possible) diseases leading to lower or no subsequent care costs
- Increasing number of diseases can be treated with a higher likelihood of success

Minimal invasive intervention has a drawback for the patient and the medical professional, because they will be exposed to X-ray for long times. In order to improve the situation it is necessary to make images with lower X-ray dose, which reduces the signal/noise ratio. New advanced denoising techniques implemented on multicore hardware allow however an X-Ray dose reduction by a factor 2 - 4, while maintaining the same image quality that is in Philips systems today.

4.2.5.4 Planned Exploitation Activities

The results of the project will be used by Philips in the next generation high-end X-Ray imaging systems to provide added clinical value. After that the same will be exploited in the value market to achieve cost-reduction. Philips will show demonstrators at medical equipment fairs such as

- RSNA - Yearly conference of the Radiological Society of North America in Chicago.
- ESC - Conference of the European Society of Cardiologica

Exploitation will be undertaken by the generation of patents when critical and innovative results are obtained in the fields of technology or in case of new openings for applications. The increased level of knowledge, technology, and/or product portfolio will enable new customer projects and/or R&D projects, in the fields

of medical diagnosis. Philips will use the project results to leverage the development of commercial products in the field of healthcare with fully integrated image processing solutions.

4.2.5.5 Contribution to Standards

Recently, there has been an increased awareness and interest about standardisation of image processing components.

Modeling. Standards and model homogenization used through all fields of research must be discussed and adopted in order to facilitate data and algorithm distribution for various levels of image processing. When healthcare information is exchanged between different medical information systems, it is essential that the clinical meaning expressed by the original author within a set of record entries is faithfully preserved and presented by the receiving system, even though this may have internal system architecture different from that of the sending system. Parallel computing and imaging, open standards & APIs for parallel computing, graphics and dynamic media could be targeted, focusing for instance on standard programming environment for image processing chains.

Interoperability and image exchange. Interoperability and instruments are to be provided as an open platform for asset development by third parties. Recently, there has been an increased awareness and interest about standardisation of image processing components, e.g. in the Khronos and ONVIF groups. Also IEEE standards may be addressed, and especially the new 1722 & 1733 drafts related to 802.1 Audio/Video Bridging (AVB) that intend to provide an environment for the highest quality streaming A/V experience.

4.2.6 Individual Exploitation Strategy for Academic Partner Technion

4.2.6.1 General Exploitation Strategy

For more than eight decades the Technion - Israel Institute of Technology has been Israels primary technological university and the largest centre of applied research. It is ranked among the leading technological universities in the world. Many innovations in all fields of science, technology, engineering and life sciences have their origins in research conducted at the Technion. The Technion offers access to state-of-the-art in scientific and technological innovation at Technion via Technion Technology Transfer (TTT) office. Our agenda includes: licensing technologies developed at the Technion, Incorporation of spin-off companies based on Technion IP, and negotiation and approval of the IP and business aspects of agreements with industry.

4.2.6.2 Transfer Opportunities

We except primary transfer opportunities in the development of the resource management layer that will be developed as part of WP6. That includes method to be developed to choose the (near) optimal hardware from the available platforms, based on a characteristic of the application at hand. Criteria for the best choice have to be derived from the expectations and requirements given in the application description, from the virtual hardware description, and from the results of the aforementioned observations.

4.2.6.3 Planned Exploitation Activities

Currently, as the project is in early stage, we have not developed yet exploitable components. We will make our TTT office, as well as our business partners, aware of any results that may be exploitable. Such technology may also be communicated to a wider circle via scientific publications in leading conferences or leading journals.

4.2.7 Individual Exploitation Strategy for Industrial Partner SAP

4.2.7.1 SAP's Business Model

SAP's mission is to make every customer a best-run business. The goal is to enable them by delivering the world's best business applications: on premise, on demand, on device with overarching orchestration integrating all these channels. The application portfolio covers applications for large, medium-sized and small enterprises. In addition to business applications, SAP also offers a stable technology platform on which customers and partners can build their own solutions or extend the solutions from SAP.

In addition to software, the company also delivers services. Those include software related services like custom development, support, and on-demand software services, but also non-software related services like consulting and education. In order to deliver highest value to its customers, SAP builds on co-innovation and a strong partner ecosystem.

4.2.7.2 SAP Research's General Exploitation Strategy

Figure 4.1 illustrates SAP Research's invention process. Project execution mainly overlaps with phase 3 ("Co-innovative Research"). Any concepts, or artefacts built during the course of the project will be used in the transfer phase. They can lead to completely new technologies, new product ideas, improvement of existing products. Alternatively, customer pilots and spin-offs are also considered as possible exploitation paths.

As ADVANCE is a research project, it is obvious, that transferability can not be guaranteed. Although the topic of the project was aligned with the corporate

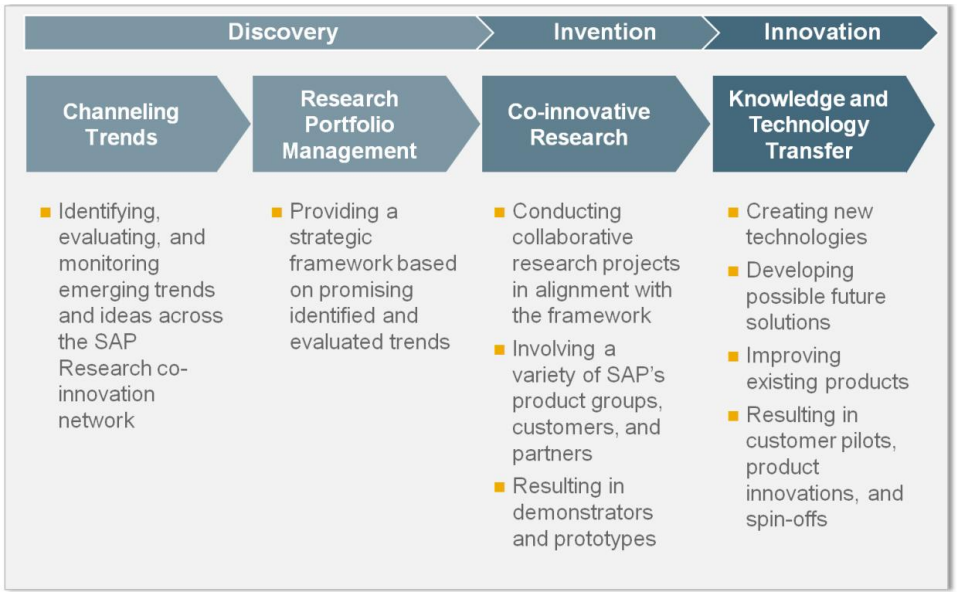

Figure 4.1: SAP Research Exploitation Strategy

strategy during the "Discovery" phases, it might happen that the artefacts produced in ADVANCE do not fit into the evolving strategy of the departments to uptake the results or are simply not requested by the market.

4.2.7.3 Transfer Opportunities

SAP is the world-leading provider of e-business software covering solutions for a wide range of industries including manufacturing (e.g. aerospace, automotive, high tech, life sciences etc.), financial and public services (e.g. banking, defence & security etc.) as well as services (e.g. telecommunications, travel & logistics etc.). Applications in these areas evolve with a growing demand on computational speed and data throughput. At the same time, cost for system acquisition, maintenance and energy consumption have to be reduced. Multi-core technologies further advance to tackle the above-mentioned demands. SAP considers (heterogeneous) multi-core platforms as a key technology for the acceleration of a wide spectrum of applications in its portfolio. ADVANCE has the potential to enable SAP applications to deal with increasing masses of data in expedited time.

One goal of the participation of SAP in the ADVANCE project is to foster the industrialisation of parallelism. Current multicore systems and tools still suffer from a lack of user-friendly parallelization, suitable modelling of parallelism and high-level programming language support. The increasing degree of heterogeneity further widens the gap between device innovations and model/tool support. The ADVANCE project has the potential to cope with these problems as it will deliver

an end-to-end tool chain and programming language support which may provide for a high productivity of SAP developers, facilitate the adoption by SAP customers and shorten time-to-market.

The wide range of SAP applications with different requirements (in terms of, e.g., concurrency, real-time constraints, adaptivity at runtime) demand a highly flexible platform which can be used for multiple purposes. The ADVANCE platform and tools may suitably approach this demand: application requirements are analyzed to find a best fit for the target hardware, and tools are aware of the state of the running system and may be able to efficiently execute and monitor the scheduled applications. SAP applications may also benefit from the adaptive placement capabilities, e.g., in the case of dynamic changes of input parameters.

The following list of application areas have been identified as prospective transfer areas for the ADVANCE project: SAP Supply Chain Management and Enterprise Search and Analytics. This selection is preliminary, code analysis and architectural decisions during the execution phase will influence the best-fitting choice of applications for this project.

Architectural decisions, problems faced and lessons learned are expected to flow towards this direction. We will pursue a market reality check and closely collaborate with business units of SAP in order to make sure that the efforts of ADVANCE can be realized in the real-world.

4.2.7.4 Planned Exploitation Activities

Due to the early stage of the project, no concrete transfer activities are planned to date. However, ADVANCE has already been introduced internally through several potential organizations that could act as transfer targets. Negotiations on concrete internal transfers can be started as soon as first tangible results or concepts can be demonstrated. Potential transfer targets include the product groups of SAP Supply Chain Management and SAP TREX (which is a solution for enterprise search and analytics). SAP expects to generate significant IP from ADVANCE in the form of patents, trade secrets, or competitive advantage both technologically and w.r.t. business insight.

4.2.7.5 Contribution to Standards

Work carried out in ADVANCE has the potential to contribute to standardization. SAP is member in several organizations including OMG, Open Group, OASIS, W3C, IPSO alliance etc. Results that extend existing standards wrt. software engineering, parallelization and resource managment can be evaluated and possibly given as feedback to the respective groups.

4.2.8 Individual Exploitation Strategy for Industrial Partner BioID

4.2.8.1 Business Model

BioID's strategy is to develop and continuously improve a multi modal, biometric framework of pattern recognition algorithms for face finding and recognition, speaker recognition, iris finding and recognition and finger print recognition. On the basis of these software modules, which can be seen as components of a "biometric engine", end-user applications and solutions are built. Due to the modular approach of the biometric framework and the highly portable source code of the core algorithms, various target markets can be addressed, including PC applications (possibly bundling BioID authentication solutions with devices), mobile applications, cloud computing / Web 2.0, and authentication in connection with e-passport documents. All these potential markets for BioID's biometric solutions are visualized in figure 4.2.

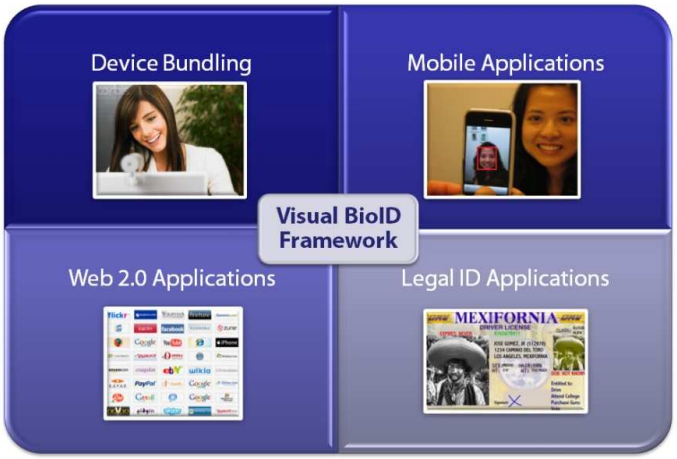

Figure 4.2: BioID's target markets

4.2.8.2 General Exploitation Strategy

BioID plans to exploit the implementations within the ADVANCE project for improving the existing biometric test and optimization framework.

BioID's biometric solutions are based on complex pattern recognition algorithms in the domain of image, video and speech processing. In general, pattern recognition problems cannot be solved with recognition rates of 100 %. In order to achieve reliable and robust software as well as to adapt modules or sub systems to end user applications with particular requirements, testing and optimizing recognition accuracy is crucial to the success of the application. For this testing and optimization processes, BioID has developed a suite of tools for the distinct

biometric traits like speaker recognition, face finding and recognition, iris finding and recognition and others. Typically, these tools rely on executing suitable test protocols on appropriate, representative databases of images or speech recordings, and compute recognition or error rates for certain recognition problems. In order to avoid over-optimization, these databases should be sufficiently large, which usually results in a very computation-time intensive process, and it is important to make most efficient use of current hardware.

BioID's exploitation strategy therefore consists in porting or re-writing certain parts of the biometric test and optimization framework to make use of the tools and techniques that are investigated and developed within the ADVANCE project. This should make it possible to use current hardware, with respect to multi core CPUs and GPUs, in a most efficient way. It is expected that the use of ADVANCE technologies will reduce development effort when hardware architectures change, e.g., if new CPUs with an increasing number of cores come into market, or if GPUs are increasingly suitable for scientific calculations. Furthermore, making better use of existing, heterogeneous computation power can also lead to increased recognition accuracy of biometric algorithms, either by better adaptation of certain models that are in use e.g. for face finding, or by better parameter optimizations of biometric algorithms in general.

4.2.8.3 Transfer Opportunities

In principle, it is conceivable that some of the software modules that are being adapted or ported to use ADVANCE technologies are transferred to BioID's end user solutions. This could make it possible to make most efficient use of the hardware platform that the respective end user has. However, as the most computation-intensive part connected to our biometric solutions is the test and optimization framework, as described above, the goal is to transfer the ADVANCE results to this framework rather than to the end user applications, especially in the short- and medium-term view. The end user applications that we are currently targeting are to some extent mobile devices, most of which are currently still less suitable for running concurrent software for optimum performance, so the application of the tools of the ADVANCE suite to the biometric test and optimization framework is more straightforward.

4.2.8.4 Planned Exploitation Activities

After the first 12 months of the project, a test protocol on the basis of a corpus of recorded speech utterances has been implemented in the SaC language, which is one of the languages supported within the ADVANCE framework by the compiler "sac2c". This SaC application is a first example and makes it possible to compute the recognition rate of a speaker identification problem in a very efficient way on current multi core hardware, as the number of concurrent threads are easily adapted to the number of available cores. First experimental results show that a

very efficient use of the multi core CPUs is actually achieved.

The planned exploitation is based on further extensions of this example application such that it can be used for optimizations of the speaker recognition algorithms. Additional biometric modules, like feature extraction components, should be incorporated in the new, ADVANCE-enabled biometric optimization framework to achieve this. More ADVANCE techniques, especially S-NET, are expected to be suitable for these extensions and subsequent exploitation, which eventually have the potential to increase the reliability, robustness and recognition accuracy of the speaker recognition.

Taking the experiences that will result from this into account, in a second step, exploitation might also be extended to other biometric traits, like face and iris recognition.

4.2.8.5 Contribution to Standards

Being an SME, BioID's primary focus in its role as ADVANCE consortium partner has not yet been on contributions to standards. However, if the opportunity emerges within the ADVANCE project to make a contribution to standards in collaboration with other ADVANCE partners, BioID will be happy to support these efforts.

4.2.9 Individual Exploitation Strategy for Industrial Partner SCCH

SCCH aims at integrating concepts from ADVANCE into its industrial research activities and system development in order to reduce processing time. By this more sophisticated methods can be employed which increases reliability and quality of the overall inspection system for customers with high quality standards.

4.2.9.1 Business Model

The business model of SCCH is based on contract research. This applies also to the topics of ADVANCE, be it as supplementing competence which increases software engineering productivity and the resulting performance or, be it as research project of its own as for example in the field of scientific computing.

4.2.9.2 General Exploitation Strategy

By its role as applied RTO SCCH pursues the strategy to closely collaborate with world-leading academic partners to establish and further develop key technologies together with its industrial partners. This is organized by conducting common applied research with its industrial partners, and strategic research aiming at scientific publications, PhD theses and patents. Following this general strategy SCCH exploits its benefit from the ADVANCE project by building up a new core competence which supplements the well-established research domains at SCCH like software engineering and knowledge-based technologies.

4.2.9.3 Transfer Opportunities

SCCH, which is organized as a non-profit GmbH (Ltd), is an Austrian research and technology organization with primary focus on applied research in the field of software engineering and computational intelligence. SCCH's mission is to increase and maintain the competitiveness of its partner companies through nationally and internationally networked state-of-the-art research.

The scope of applications span a) fault diagnosis, causal reasoning, decision support for process and production engineers by inductive rule extraction, regression models, Bayesian networks and automated reasoning techniques, and b) automation of defect analysis, motion analysis, quality reasoning, early fault detection and diagnosis, predictive maintenance and predictive control by supervised and semi-supervised learning, and advanced model-fitting and model calibration techniques. This scope of applications provides opportunities to transfer the competitive edge of the cross-sectional technology of ADVANCE to research domains of SCCH beyond image processing applications.

4.2.9.4 Planned Exploitation Activities

SCCH plans to exploit the results and findings of ADVANCE by

- taking its concepts into account in the software design process of various applications as outlined above,
- conducting common benchmark research with the academic partners of ADVANCE in order to achieve visibility and establish high-performance computing competence in the application domain of SCCH,
- acquiring customers and research partners specially interesting in high-performance computing.

4.2.9.5 Contribution to Standards

The concepts of ADVANCE have the flavor of a cross-sectional technology for many of SCCH's applied research projects. By this the concepts and technologies of ADVANCE influence the conception and design of software systems of the above mentioned field of applications.

Chapter 5

Conclusion

In deliverable D9 we described opportunities and plans on how effective market analysis, dissemination and exploitation will be conducted in the course of the project ADVANCE. This document is the first of a series. It will be edited as a living document and extended by more concrete exploitation and dissemination plans and reports of activities that have already been engaged or finished. The next version will be delivered after month 32 of the project.

Bibliography

[1] Raimund Kirner, Jens Knoop, Adrian Prantl, Markus Schordan, and Ingomar Wenzel. WCET analysis: The annotation language challenge. In *Proc. 7th International Workshop on Worst-Case Execution Time Analysis*, pages 83–99, Pisa, Italy, July 2007.

[2] A. Shafarenko and R. Kirner. CAL: A Language for Aggregating Functional and Extrafunctional Constraints in Streaming Networks. *ArXiv e-prints*, January 2011.

[3] D. Spath, A. Weisbecker, and E. Hebisch, editors. *Market Overview of Tools for Multicore Software Development*. Fraunhofer Verlag, 2010.