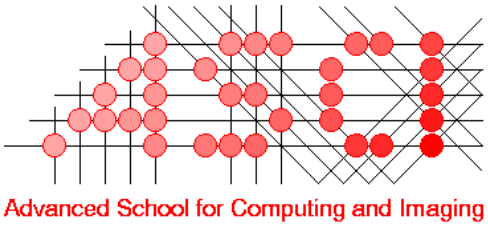

Bi-Annual Report 2007-2008

Advanced School for Computing and Imaging

p/a DUT / EWI
P.O. Box 5031
2600 GA DELFT
The Netherlands

Phone: + 31 15 27 88032
Fax: +31 15 27 86632
asci@ewi.tudelft.nl
http://www.asci.tudelft.nl

ASCI is a Dutch graduate school established in 1993 and accredited by the Royal Netherlands Academy of Arts and Sciences. Research groups of Delft University of Technology, Vrije Universiteit, University of Amsterdam, Leiden University, University Utrecht, University of Twente, University of Groningen, Eindhoven University of Technology, Erasmus University Rotterdam and Radboud University Nijmegen participate in ASCI.

Contents

Preface

The world and many of its processes are progressing ever faster and faster. Digital media and the Internet provide instant access to documents of many kinds. The digitization has had a profound effect on science and surely on the consultation of research manuscripts, papers and conference proceedings. In daily practice of science, papers are no longer of paper, but rather made up from bits. Papers circulate ever faster and the reaction speeds are ever shorter. By building fast networks, parallel machines, search engines, and massive visualizations, ASCI plays its part in the speeding up of the process.

As the world's information is delivered ever faster and faster, calmness is of increasing value as well. As a consequence, we have decreased the frequency of our summary of the research result from annually to bi-annually. In a world where individual scientific results are easily accessible through the web, the overview of all results no longer serves the purpose to inform about the content of the research, but rather to stimulate the generation of new ideas (to which search engines are not the answers) by reading about the existence of research. This overview aims to evoke just that: new ideas.

That is why we have collected for each research group in ASCI a view on current and past research, future plans and key publications. And, a few people have given a shot at making the text better readable outside the research topic. We hope to have succeeded.

In 2007 and 2008 the 13th and 14th ASCI-conferences have been held at Heijen, Limburg. The audience consisted of some 100 ASCI PhD's and 20 staff members. Keynote speakers have been Markus Gross from ETH Zurich, Robbert van Renesse from Cornell University and Johan Suykens from the University of Leuven in 2007, and in 2008 Alan Smeaton from Dublin City University, Dirk Harryvan from IBM and Jon Weissman from the University of Minnesota. In the years 2007 and 2008 two ASCI GNARP workshops have been held on Parallel Systems, and one ASCI Springschool on Embedded Systems.

Enough to enjoy reading this report, if anything, indicative for the vitality of ASCI.

Arnold Smeulders
Scientific Director

1 ASCI and its Research

1.1 About ASCI

ASCI is a national research school on advanced computer and imaging systems. The school was founded in December 1993, and it was approved by the KNAW (Dutch Royal Academy of Sciences) in May 1995. In 2005 the school got its new accreditation for the coming six years.
Participants in ASCI are groups from Delft University of Technology, the University of Amsterdam, the Vrije Universiteit, Leiden University and the University of Utrecht; the University of Twente, the University of Groningen, Eindhoven University of Technology and Radboud University Nijmegen have joined ASCI by association agreements.

ASCI performs research in two main fields: computing and imaging. In the course of time 'imaging' broadened to 'multimedia data processing'. The activities within these fields are further classified based on their main target, either Methods and Algorithms (development of models and tools for scientific and industrial applications) or Systems and Architecture (large-scale integration in areas like telematics, embedded systems, communication and networks). In both categories fundamental and applied research is done within ASCI. Much of the ASCI research is interdisciplinary, involving multiple groups and areas from computer science, electrical engineering, physics, and other departments.

The school organizes a graduate program and a research program covering all major subjects concerning parallel, distributed, embedded, and real-time systems, performance analysis, image processing, image analysis, image synthesis, sensor interpretation, pattern recognition and computer vision. Every year ASCI organises the Annual ASCI Conference, the scientific meeting place for all participants in ASCI. Another annual activity is the GNARP workshop (GNARP Graduate Network of Applied Research in Parallel systems) which is organized by PhD students and which is a platform for presenting work in progress. Every two years ASCI organizes a winter/springschool, in 2008 on Embedded Systems.

1.2 Participating Groups

The following reseach groups participate in ASCI. They are represented together with their abbreviations. For each group the members are listed (situation January 2009).

VU-EW-CS	Vrije Universiteit, Faculty of Sciences, Division of Mathematics and Computer Science, Dept. of Computer Science http://www.cs.vu.nl *Prof.dr. A.S. Tanenbaum, Prof.dr.ir. H.E. Bal, Prof.dr.ir. M.R. van Steen, Dr.ing. T. Kielmann, Dr. G.E.O. Pierre, Dr. R. van Nieuwpoort, Dr.ir. H.J. Bos, Dr.B. Crispo, Dr.ir. C. van Reeuwijk, Dr. F.J. Seinstra*
UvA-FNWI-SNE	University of Amsterdam, Department of Computer Science,System and Network Engineering http://www.science.uva.nl/research/sne *Dr. C.T.A.M. de Laat*
UvA-FNWI-IAS	University of Amsterdam, Faculty of Science, Informatics Institute, Intelligent Autonomous Systems http://isla.science.uva.nl *Prof.dr.ir. F.C.A. Groen, Prof.dr. D.M. Gavrila, Dr.ir. B.J.A. Kröse, Dr.ir. L. Dorst, Dr. G. Pavlin*
UvA- FNWI-ISIS	University of Amsterdam, Faculty of Sciences, Informatics Institute, Intelligent Sensory Information Systems Group http://isla.science.uva.nl *Prof.dr.ir. A.W.M. Smeulders, Dr.ir. R. van den Boomgaard, Dr. M. Worring, Dr.Ing. J.M. Geusebroek, Dr. Th. Gevers, Dr. C.J. Veenman, R.F. Aldershoff, Drs. C.G.M. Snoek, Dr. N. Sebe*
UvA-FNWI-CSA	University of Amsterdam, Faculty of Sicence, Informatics Institute, Computer Systems Architecture Group http://www.science.uva.nl/research/csa *Prof.dr. C. Jesshope, Dr. A.D. Pimentel*

TUD-EWI-ST-PGS	Delft University of Technology, Faculty of Electrical Engineering, Mathematics & Computer Science, Parallel and Distributed Systems Group http://www.pds.ewi.tudelft.nl *Prof.dr.ir. H.J. Sips, Dr.ir. D.H.J. Epema, Prof.dr.ir. A.J.C. van Gemund, Prof.dr. C. Witteveen, Dr. K.G. Langendoen*
TUD-EWI-MM-CGCC	Delft University of Technology, Faculty of Electrical Engineering, Mathematics & Computer Science, Computer Graphics and CAD/CAM Group http://graphics.tudelft.nl *Prof.dr.ir. F.W. Jansen, Dr. W.F. Bronsvoort, Ir. F.H. Post, Dr. C.P. Botha*
TUD-EWI-MM-ICT	Delft University of Technology, Faculty of Electrical Engineering, Mathematics & Computer Science, Department of Mediamatics, Information and Communication Theory Group http://msp.ewi.tudelft.nl and http://prb.tudelft.nl *Prof.dr.ir. J. Biemond, Prof.dr.ir. R.L Lagendijk, Dr. E.A Hendriks, Dr.ir. J.C.A. van der Lubbe, Prof.dr.ir. M.J.T Reinders, Dr. A Hanjalic, Dr.ir. D. de Ridder, Dr. D. Bellomo, Dr. M.J.L. de Groot, Dr.ir.J.R. Taal*
TUD-TNW-QI	Delft University of Technology, Faculty of Applied Physics, Imaging Science & Technology, Quantitative Imaging Group http://www.ist.tudelft.nl/qi *Prof.dr. L.J. van Vliet, Dr. B. Rieger*
UL-LIACS	Leiden University, Faculty of Mathematics and Natural Sciences, Leiden Institute of Advanced Computer Science (LIACS) http://www.liacs.nl/research *Prof.dr. H.A.G. Wijshoff, Dr. M.S. Lew, Dr. A.A. Wolters, Dr. D.P. Huijsmans, Dr. E.M. Bakker, Dr.ir. T.P. Stefanov, Dr.ir. B. Kienhuis, Prof.dr.ir. E.F. Deprettere*
UU-ICS-GMT	Utrecht University, Faculty of Science, Department of Information and Computing Sciences, Center for Geometry, Imaging and Virtual Environments http://www.cs.uu.nl/centers/give/give-center.html and http://www.cs.uu.nl/groups/MG *Prof.dr. M.H. Overmars, Dr. M. van Kreveld, Dr.ir. F. van der Stappen, Dr. R. Veltkamp, Prof.dr P.J. Werkhoven*
TUD-L&R-FRS	Delft University of Technology, Faculty of Aerospace Engineering Department of Earth Observation and Space systems (DEOS), Optical and Laser Remote Sensing Group http://www.lr.tudelft.nl/olrs *Dr.ir. B.G.H. Gorte*
RUG-CS-SVCG	University of Groningen, Faculty of Mathematics and Physical Sciences, Institute of Mathematics and Computing Science, Scientific Visualization and Computer Graphics http://www.cs.rug.nl/svcg *Prof.dr. J.B.T.M. Roerdink, Dr. H. Bekker*
RUG-CS-IS	University of Groningen, Faculty of Mathematics and Natural Sciences, Institute of Mathematics and Computing Science, Intelligent Systems http://www.cs.rug.nl/is *Prof.dr. N. Petkov, Dr. M.H.F. Wilkinson*
TUE-ET	Technische Universiteit Eindhoven, Faculty of Electrical Engineering, Design Methodology for Electronic Systems http://www.es.ele.tue.nl *Prof.dr.ir. R.H.J.M.Otten, Dr.ir. T. Basten, Prof.dr. H. Corporaal, Dr.ir. M.C.W. Geilen, Prof.dr.ir. G. de Haan, Dr.ir. J.P.M. Voeten*

TUE-WI	Technische Universiteit Eindhoven, Department of Mathematics and Computer Science, Visualization Group http://www.win.tue.nl/vis *Prof.dr.ir. J.J. van Wijk, Prof.dr.ir. R. van Liere, Dr.ir. H.M.M. van de Wetering, Dr. M.A. Westenberg, Dr. A.C. Jalba*
TUE-BMT	Technische Universiteit Eindhoven, Department of Biomedical Engineering Biomedical Image Analysis http://bmia.bmt.tue.nl *Prof.dr.ir. B.M. ter Haar Romeny, Prof.dr. F.A. Gerritsen, Prof.dr. L.M.J. Florack, Dr. A. Vilanova Bartroli, Dr.ir. H.C. van Assen, Dr.ir. R. Duits, Dr.ir. B. Platel*
UL-LUMC	Leiden University Medical Center, division of Image Processing, laboratorium voor klinische en Experimentele Beeldverwerking http://www.lumc.nl/con/1010/83058/87360/87377 *Prof.dr.ir. Johan H.C. Reiber, Dr.ir. B.P.F. Lelieveldt, Prof.dr. R. Nelissen*
UT-EWI-DACS	University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science, Design and Analysis of Communication Systems http://dacs.ewi.utwente.nl P*rof.dr. B. Haverkort, Dr.ir. G.J. Heijenk, Dr.ir. A. Pras, Dr.ir. P.T. de Boer, Dr.ir. G. Karagiannis*
UT-EWI-CAES	University of Twente, Faculty of Electrical Engineering, Mathematics an Computer Science, Computer Architecture, Design & Test for Embedded Systems http://caes.ewi.utwente.nl *Prof.dr.ir. G.J.M. Smit, Dr.ir. A.B.J. Kokkeler, Ir. E. Molenkamp*
EUR-UMCR	Erasmus MC, Departments of Radiology and Medical Informatics, Biomedical Imaging Group Rotterdam http://www.bigr.nl *Prof.dr. W.J. Niessen, Dr. J. Veenland, Dr.Ir. H. Vrooman*
RUN-UMCR	Radboud University Nijmegen Medical Centre, Radiology Department, Diagnostic Image Analysis Group (DIAG) Nijmegen http://www.umcn.nl Dr.ir. N. Karssemeijer, Dr.ir. H.J. Huisman, Dr. M. Velikova

1.3 ASCI Research Themes

ASCI research comprises historically two main interlocked themes. The 'C' in ASCI stands for Computing and the 'I' for Imaging. One half of ASCI is still best represented by the 'C' for Computation, but the 'I' is gradually developing into Sensory Information, which still justifies the 'I' if the emphasis is shifted from the word Imaging to Information.

Within the two ASCI themes, the scientific activities can be divided in Methods & Algorithms on the one hand and Systems & Architectures on the other. Methods & Algorithms deal with the development of models and tools as such. They are eventually directed towards particular applications in science or industry and non-profit organizations. Systems & Architectures deal with the large-scale design and integration of tools, and the evaluation thereof at the system level. They are eventually directed to exemplary systems such as embedded systems, communication networks, information analysis systems, search engines, and visualization systems. Both types of activity are targets for fundamental and applied research.

These themes and target areas are combined in the following matrix, in which the different computer science research disciplines covered by ASCI can be placed.

	Methods & Algorithms	*Systems & Architectures*
Computing	A 1. High Performance Computing 2. Computational Science	B 1. Large-Scale Information Systems 2. Distributed Systems 3. Embedded Systems 4. Sensor networks
(Sensory) Information	C 1. Image and Multimedia Sensing 2. Processing 3. Interpretation and Visualization	D 1. MM Analysis & Search Systems 2. Sensing and Learning Systems 3. Acting and Visualization Systems

Within the area of Computing we have seen several important developments. The field of high performance computing (HPC) and computational science (theme A) has shifted focus from exclusively computation-intensive computing to also include data-intensive computing, which is required by many e-Science applications. Apart from traditional HPC, ASCI now also studies data-centric aspects, including networking, high-performance communication, I/O, and security. The result is a more balanced treatment of computation-intensive and data-intensive applications.

The field large-scale distributed information systems and embedded systems (theme B) has been extended considerably. The work on operating systems is now focusing on the highly successful Minix-3. The distributed systems subtheme has grown with many new activities on grids, peer-to-peer systems, and sensor networks, especially resource management, network management, programming, and workflow systems. Much of this work is related to the national BSIK projects VL-e and Freeband, and to the NWO program I-Science. ASCI´s research on embedded systems has expanded significantly. Many ASCI groups are studying Systems-on-a-chip (SoC) designs, for example multiprocessor SoCs. They collaborate extensively in ASCI and in large STW Progress projects. The field of sensor networks is gaining a lot of momentum within the ASCI community. Many groups deal with aspects of sensor networks, including distributed communication protocol such as gossiping, low-power sensor nodes, distributed information processing, and security aspects.

In the field of Sensory Information and Systems (theme D), large developments have taken place. Where sensory information has been an academic topic of study for twenty years, it recently became part of mainstream information and communication technology for two reasons. Firstly, massive digitization of all sensory data is taking place, for robots, science or popular use alike. Secondly, very large-scale archives are disclosed through digital media networks, again in science and society. Multimedia systems are no longer academic playgrounds but real platforms with many applications in science and the arts, cultural heritage, safety services, medical imaging, industry and the population at large.

In the field of Sensory Data Methods and Algorithms (theme C), steady developments take place to unravel the structure of multimedia data of many different sources. Examples are the understanding of the (deep) structure of images in for instance medicine, the structure of space observed through sensor networks or mobile robots, the learning of facts from multimedia information, and the understanding and exploitation of multimedia exchange, consumption, alteration and annotations in on-line social communities.

At the onset, ASCI anticipated the massive computation needs generated by the digitization of massive sensory data. This expectation has come true as can be seen from the processing of popular resources such as Hyves and Flickr, or professional archives like broadcast archives, or scientific resources with large archives in ecology, astronomy or geo-sciences. The research program and the educated PhD students of ASCI have and will remain to contribute to this confluence.

Furthermore, for large repositories of data and knowledge, also the structuring and computation of heterogeneous and multimedia sources of knowledge in ontologies and databases is increasingly important. This has lead to more co-operation with members of the Research School SIKS.

2 Cooperation within ASCI

ASCI researchers collaborate through the DAS infrastructure and throught numerous externally funded collaborative projects. Below, we describe DAS and some of the largest collaborative projects.

The DAS projects

DAS, Distributed ASCI Supercomputer, is the experimental infrastructure shared among all ASCI researchers. The first DAS system was set up in 1997, while DAS-3 was operational in 2006. The successive systems were funded largely by three equipment grants from NWO. Each generation consists of four to five clusters at different locations, connected by a wide-area network and integrated into a single, shared, distributed system. DAS-3 pioneers a novel flexible dedicated optical interconnect Starplane provided by SURFnet.

DAS is unique in that it is designed specifically for experimental computer science and (unlike grids) it is designed and managed by a single organization (ASCI). The DAS systems have been highly successful and clearly demonstrated that computer scientists need such a dedicated distributed infrastructure. Over a 100 researchers currently use DAS, including dozens of PhD students. In the period 2005-2008, 36 PhD-students used DAS for their research project. Only few other countries have such a facility; the most prominent similar system is the French Grid'5000, with which we collaborate intensively.

Large-scale collaborative projects

Ever since the start of ASCI, its researchers collaborated in numerous joint projects. ASCI researchers also participated in (or initiated) several very large collaborative programs, each involving many dozens of scientists, often from different areas. ASCI played a leading role in most of these programs, and we feel that the efforts we invested during the preceding decade in building a coherent research community paid off in these programs. The programs are described below.

MultimediaN is a 30 MEuro BSIK program that runs from 2004 to 2009 exploring (large-scale) digital sensory data, their storage and their interaction. Sensory mostly focused on image and video data, with small excursions to audio, speech and text processing. The program is well connected to industry and non-profit organizations. The operational model is based on the work-table model where intensive co-operation at the lowest level of the organization takes place and immediate transfer of knowledge and know-how is achieved.

VL-e (Virtual Laboratory for e-Science) is a 40 MEuro BSIK program running from 2004 to 2009 that studies virtual laboratories for e-Science. About one third of this program consists of Computer Scientists (mostly from ASCI) who design generic methods and tools for scientific applications from high-energy physics, medicine, bioinformatics, biodiversity, and other areas. This program boosts our research on grid programming environments, workflow systems, problem solving environments, resource management, and networking.

STW Progress program. ASCI has set up a research program together with Philips Corporation on future consumer electronics (CE) devices, resulting in three related projects that were granted from STW (the Dutch national technology foundation) in the Progress program on embedded systems. They study different aspects of System-on-a-Chip (SoC) architectures, including software engineering methods (SCALP), mapping domain specific (video) applications onto a domain specific Network-on-a-Chip platform (Artemisia) and the design of NoC-based real-time systems (PreMaDoNa). The projects collaborate intensively in regular tri-partite meetings.

NWO i-Science program. NWO has set up a cluster of three programs (GLANCE, VIEW, STARE) to advance the research in e-Science in the Netherlands. Numerous collaborations within ASCI and between ASCI and other partners were funded by these programs, including large-scale distributed systems projects in GLANCE (GUARD-MM, StarPlane, GUARD-G, MicroGrids), visualization projects in VIEW (EIO, IMOVIS, MFMV, Multi-Vis), and research on astronomy applications in STARE (Astrostream, SCARI, ASTROVIS). Almost two third of the i-Science projects contain ASCI researchers.

3 Scientific Output along Research Themes

This chapter contains the contributions of the various research groups in ASCI. The following table lists the enrollment of the groups with respect to the themes.

	Methods & Algorithms	Systems & Architectures
Computing	A 1. High Performance Computing & 2. Computational Science Wijshoff (UL)	B 1. Large-Scale Information Systems Tanenbaum (VU) 2. Distributed Systems Bal (VU), De Laat (UvA), Sips (TUD), Pras (UT) 3. Embedded Systems Corporaal (TUE), Smit (UT), Jesshope (UvA), Haverkort (UT) 4. Sensor networks Langendoen (TUD), Van Steen (VU)
(Sensory) Information	C 1. Image and Multimedia Sensing Van Vliet (TUD), Ter Haar Romeny (TUE), Groen (UvA) 2. Processing Petkov (RUG) 3. Interpretation and Visualization Roerdink (RUG), Jansen (TUD), Van Wijk (TUE), Overmars (UU)	D 1. MM Analysis & Search Systems Veltkamp (UU), Smeulders (UvA), Lagendijk (TUD) 2. Sensing and Learning Systems Niessen (EUR), Lelieveldt (LUMC), Karssemeijer (RUN), Reinders (TUD) 3. Acting and Visualization Systems Gorte (TUD)

3.1 Contribution of VU-EW-CS

Vrije Universiteit, Faculty of Sciences, Division of Mathematics and Computer Science, Dept. of Computer Science
Prof.dr. A.S. Tanenbaum, Prof.dr.ir. H.E. Bal, Prof.dr.ir. M.R. van Steen, Dr.ing. T. Kielmann, Dr. G.E.O. Pierre, Dr. R.van Nieuwpoort, Dr.ir. H.J. Bos, Dr.B. Crispo, Dr.ir. C. van Reeuwijk, Dr. F.J. Seinstra

Research 2007-2008

The Section Computer Systems of the VU comprises three programs: Secure and Reliable Computer Systems (Tanenbaum), High Performance Distributed Computing (Bal), and Large-Scale Distributed Computer Systems (Van Steen).

Secure and Reliable Computer Systems

It is our belief that if the operating system is unreliable and insecure, there is little hope for building reliable and secure applications. Our focus has been on research concerning operating systems with a different architecture than current ones. Current systems have millions of lines of code linked together in a single huge binary program running in kernel mode. A bug in any line of code can bring down or corrupt the system. We consider this situation highly undesirable. In our model, the kernel should be reduced to a minimum, handling only interrupts, MMU management, scheduling, and message passing. The rest of the operating system, especially the device drivers, should run as a collection of user-mode processes. Much of the research is concerned with finding ways to put up barriers between the components so that problems in one component cannot spread to other ones. Also, we looked at ways to be able to restart operating system components, should they fail, without having the system crash and without affecting running processes.

While improving operating system reliability also indirectly improves security, we also looked at security directly, in the context of networked and distributed systems. Often there are security policies that extend over machine boundaries, especially for mobile computers. One well known issue here is digital rights management, in which the owner of some intellectual property, such as music, wants to control how the IP is used by its users. For example, a music vendor may

authorize a customer to play the music on any of the customer's computers or other devices, but not in those belonging to the customer's neighbors. However, we went beyond DRM to look at more general issues of remotely enforcing security policies. For example, most e-commerce Websites require the customer to agree to the site's terms and conditions. Suppose the customer also had terms and conditions to which the e-commerce vendor had to agree in order to get the sale. For example, suppose that the customer had a requirement that the e-commerce vendor would not distribute the customer's e-mail address to third parties. As a different example, imagine that future e-mail systems support e-mails tagged with policies, such as "Do not forward" or "For distribution only within the company." The research issue is how the policies can be enforced at a technical (as opposed to a legal) level. Using advances in secure hardware technology, encryption, secure boot, and remote attestation of running processes, we made strides in remote policy enforcement.

Closely related to this is how to specify and enforce more general security contracts. This involves what kind of infrastructure is present, how security semantics can be described, and how users can be convinced that the stated security policies are being enforced. S3MS addresses these concerns and is creating a framework and technological solution for trusted deployment and execution of communicating mobile applications and services in a mobile, heterogeneous environment. S3MS R&D activities may enable the opening of the software market of nomadic devices (from smart phones to PDA) to trusted third party applications beyond the sandbox model and without the burden of complex roaming trust infrastructure and without compromising security and privacy requirements of all involved stakeholders.

Besides reliability and security, we are also interested in privacy. At the low end, we are interested in RFID chips, especially their security and privacy aspects. We did work on RFID viruses. We also looked at privacy aspects. RFID chips are going to be everywhere in the future, in supermarket products, clothes, animals, banknotes, just about everything. In particular, we have been looking at how to safeguard privacy in a world full of RFID chips. Our approach is to design a battery-powered device that we call the RFID Guardian. This portable device monitors RFID chips in the vicinity of this owner as well as RFID scans of the owner's chips. It can be programmed to recognize friendly from hostile chips and scans and act accordingly, including selectively blocking hostile scans.

High Performance Distributed Computing

The landscape of distributed computing systems has changed many times over the previous decades. Modern real-world distributed systems consist of clusters, grids, clouds, desktop grids, and mobile devices. Writing applications for such systems has become increasingly difficult. We feel there is an urgent need to drastically simplify the programming of such applications. It is important to study the underlying fundamental problems of distributed computing (e.g., performance, heterogeneity, malleability, fault-tolerance, and connectivity) hand-in-hand with major applications.

Most of our efforts here are integrated into a single large project called Ibis, funded by VL-e, VU-CvB, and NWO. An important insight from our work is that these problems require functionalities from both programming systems (traditionally associated with languages) and from deployment systems (traditionally associated with operating systems). Ibis thus consists of two subsystems:

1. The Ibis High-Performance Application Programming System. It contains a Java-centric communication library (IPL) that is designed for dynamically changing distributed environments, such as grids or mobile devices. Many programming models have been implemented on top of the IPL, including adaptive divide-and-conquer and a high-level imaging model. The lowest layer is a library (SmartSockets) that automatically solves connectivity problems due to firewalls, etc.
2. The Ibis Distributed Deployment System supports application deployment and management, file I/O, etc. An intermediate programming interface (JavaGAT) makes the operations available in a middleware-independent manner, a high-level GUI (IbisDeploy) is used for running and managing grid applications, and the Zorilla peer-to-peer layer performs resource management.

The resulting system is used by many researchers elsewhere and has resulted in several awards. We use the system ourselves for several e-Science applications and for applications like model checking, game tree search, multimedia content analysis.

We also use Ibis for distributed applications that involve mobile devices, such as smart phones and PDAs. Such systems bring new fundamental problems like how to handle context-awareness (e.g., location) and how to deal with resource-limited devices. In our vision, as much processing as possible should be pushed from such devices to less constrained environments (for instance, a cloud). In one of our demo applications, we use a phone's camera to capture an image and use the cloud to run image recognition software. The phone runs the same Ibis software as described above for Grids. In another application domain, we conduct research in providing comprehensive security for mobile devices. As phones and PDAs are too resource constrained to run heavy-weight security checks themselves, our solution is to capture a minimal execution trace, which is subsequently replayed at a remote security server.

We also work on networking and security for distributed systems. High-speed networking environments pose new challenges to existing operating systems. The Streamline/FFPF project reconfigures the network stack at runtime to be able to improve I/O by optimizing copying, virtual memory operations, context switching, and use of the various caches. It also facilitates the integration of embedded hardware in the network stack, which we demonstrated with FPGA boards and Intel IXP network processors. From a network-specific solution, Streamline has evolved to a generic I/O architecture that encompasses pipes, disk, local IPC, etc.

In the NWO project AstroStream (with TU Delft and Astron), we investigate overlay networks for improving the throughput of data streaming across the Internet, driven by the requirements of large-scale radio telescope projects, such as LOFAR and SKA. In another NWO project we investigated novel multicast algorithms for data-intensive grid applications. The new approach is receiver-initiated, and automatically optimizes throughput based on the (fluctuating) bandwidth.

As future plans, there are many new application areas, especially multimedia, model checking, and knowledge-based reasoning, in which we have just started new initiatives. The forthcoming DAS-4 system (expected in 2010) will allow us to do unique experiments with modern architectures (GPUs, Cell, manycores, etc) and in a much more controlled way than other systems. It will also be an excellent platform to investigate energy consumption issues. Also, we will study the opportunities and challenges that cloud computing will bring. We will strengthen our role in system security research in a new European Network of Excellence (SysSec). We intend to use Streamline in a new European project (I-Code) to detect new attacks in the network.

Large-Scale Distributed Computer Systems

CONTENT DELIVERY NETWORKS

Globule. We have developed a collaborative CDN called Globule. The key issue in Globule is that it fully supports automated replication of (static) Web pages based on real-time analysis of access traces. To support sites that required high performance, we developed the concept of distributed server: a collection of relatively unreliable hosts that jointly appear to its clients as a high-end server. Using technology from mobile IPv6 (MIPv6), we have devised an efficient solution by which a large collection of such hosts can share a single IPv6 address. We have been able to establish full-fledged connection handoffs in less than 10 ms, implying that clients will hardly even notice handoffs in the middle of a streaming video. Our distributed server technology is now being further developed to support wide-area grid computing.

Globule initially concentrated on automatically replicating static Web pages. More challenging is the case of automated replication of dynamically generated Web content, which essentially involves replicating the application and database layer of multi-tiered Web servers. As with static Web pages, differentiating replication strategies showed to be crucial. In collaboration with Alonso's group at ETH Zurich, we have devised edge-server solutions that combine partial replication, content-aware caching, and content-blind caching. An important part of this research consisted of developing feedback control mechanisms that allow us to automatically (de)allocate resources to a multi-tiered Web site. The results of this part of the Globule research have been incorporated into Amazon's infrastructure.

Wikipedia. To further our research in collaborative CDNs, we have initiated a project to develop a scalable solution for Wikipedia. We have collected a huge access trace in collaboration with the Wikimedia foundation, containind more than 25 billion requests and still growing. We used this trace to analyze potential replication strategies and to design a collaborative, decentralized and scalable Wikipedia system.

EPIDEMIC-BASED DISTRIBUTED SYSTEMS

Wireline systems. W have extensively investigated the role of epidemics in developing very large-scale decentralized distributed systems. We have been the first to investigate the impact that specific protocols have on the functional and nonfunctional behavior of an epidemic-based system (as reported in ACM TOCS, 2007). In particular, we concentrated on developing an epidemic system for randomly sampling live peers from a large network. Results of this have found their way into our research on Globule (for disseminating information), Tribler (decentralized recommendations), and our Wikipedia system (large-scale dissemination of consistency information).

Wireless systems. Since approximately 2004, we started to investigate epidemics for wireless systems, notably large-scale wireless sensor networks. This research is conducted in collaboration with Chess, a Dutch-based high-tech SME. They have implemented our epidemic protocols into their sensor nodes, and are now exploiting our protocols in numerous commercial R&D projects. Initial research was moderate and concentrated on exploring epidemic protocols through simulations. This has led to a PhD thesis (Gavidia, 2009). Since 2007, we have been expanding our research in this area. In 2008, we set up a first version of a 60-node TelosB-based sensor network, which has proven to be of immense value for validating results. At the moment, four PhD students are working in this area of wireless distributed systems, in collaboration with TU-Delft and Chess, First results at visible venues are being published (best student paper IPSN'09 and best paper EWSN'09).

Key Publications 2007-2008

Herder, J.N., Bos, H., Gras, B., Homburg, P., Tanenbaum, A.S.: Failure Resilience for Device Drivers, Proc. Dependable Systems and Networks 2007, pp. 41-50, 2007 (Best Paper award).

Mancina, A., Herder, J.N., Gras, B., Tanenbaum, A.S., and Lipari, G.: Enhancing a Dependable Multiserver Operating System with Temporal Protection via Resource Reservation, Proc. 16th Int'l Conf. on Real-Time and Network Systems, pp. 41-50, 2008 (Best Paper Award).

Nair, S.K., Gerrits, R., Crispo, B., and Tanenbaum, A.S.: Turning Teenagers into Stores, IEEE Computer, vol. 41, pp. 58-62, Feb. 2008.

Georgios Portokalidis and Herbert Bos: Eudaemon: Involuntary and On-Demand Emulation Against Zero-Day Exploits, Proceedings of ACM SIGOPS EUROSYS'08, pp. 287–299, April 2008, Glasgow, Scotland.

Rob V. van Nieuwpoort, Thilo Kielmann and Henri E. Bal: User-Friendly and Reliable Grid Computing Based on Imperfect Middleware, Proc. of the ACM/IEEE Conference on Supercomputing (SC'07), Nov. 2007, Reno, NV, USA.

Gosia Wrzesinska, Jason Maassen and Henri Bal: Self-adaptive Applications on the Grid, ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'07), March 2007, San Jose, CA.

F.J. Seinstra, J.M. Geusebroek, D. Koelma, C.G.M. Snoek, M. Worring, and A.W.M. Smeulders: High-Performance Distributed Video Content Analysis with Parallel-Horus. IEEE Multimedia,
14(4):64-75, October-December 2007.

M. Szymaniak, D. Presotto, G. Pierre, M. van Steen. Practical Large-Scale Latency Estimation.
Computer Networks, vol. 52(7):1343-1364, May 2008.

M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, M. van Steen. Gossip-based Peer Sampling.
ACM Transactions on Computer Systems, vol. 25(3) 2007.

Z. Wei, J. Dejun, G. Pierre, C.-H. Chi, M. van Steen. Service-Oriented Data Denormalization for Scalable Web Applications.
Proc. 17th Int'l World Wide Web Conference (WWW 2008), Beijing, China, April 2008.

3.2 Contribution of UvA-FNWI-SNE

University of Amsterdam, Department of Computer Science,System and Network Engineering
Dr. C.T.A.M. de Laat

Research 2007-2008

The System and Network Engineering science group (SNE) focuses its research on emerging new local and wide area optical networks and the associated models, systems and protocols. The group is building tools and proof of concept applications that promote optimal use of these high speed networks. The group develops grid middleware to empower applications to optimally allocate and use these infrastructures. Security of the required mechanisms, infrastructure, middleware, applications and the privacy of data in distributed processing environments is an essential aspect of the research. The research group is closely working together with the SNE Master and the Computer (Grid) Science Master programs to disseminate knowledge through education.

Advanced Networking

In 2007 and 2008 within our participation in the GigaPort project and with support from TNO we have further developed the semantic framework for describing networked resources named the Network Description Language (NDL). NDL is now a multi-schema language that can be used by network control planes to provision (configure) dedicated optical connections in a hybrid network. To facilitate path finding in such environments we have developed technology-independent schemas that wrap the multi-layer components of the network in standard interfaces. We have also further developed NDL to make it suitable for network emulations and simulations, so that new protocols can be tested in realistic network topologies. The group has also put significant effort into enlarging the community of NDL users and to cooperate internationally with other researchers toward a standardized language for network descriptions. This effort has been carried out within the Open Grid Forum NML-WG (Network Markup Language Working Group).
The StarPlane project has moved forward with the final delivery of all the photonic hardware necessary to provision dynamically lightpaths. We developed the integration of the StarPlane Management Plane (SPMP) with the services

provided by the Dynamic Resource Allocation Controller (DRAC) in charge of the lightpaths setup in the SURFnet6 network. We have also built active and passive monitoring services to maintain real-time information on the availability of paths.
The SCARIe project is focused on providing a software correlator for Very Long Base Interferometry, VLBI. In our implementation the correlation task follws a hierarchical master-worker model. In StarPlane we implemented this model on top of MPI and of the SATIN programming model, and we we tested the high performances aspects of StarPlane. In our experiments Starplane is delivering 400MB/s of throughput between the cluster sites and thus permits us to distribute medium correlation jobs to the DAS-3 grid.
CineGrid, of which the SNE group is a founding member, is a new collaboration in the field of high-quality media content delivery. CineGrid recognizes the need of new collaborative tools to enable the transport of this content. Data need to move seamlessly from the location where it is stored to the end user requesting it in a fast and automatic way: we need IT infrastructures that offer enhanced Quality-of- Service suitable for real-time data distribution. In 2007 we have started to build in the first European CineGrid distribution center. We have in particular focused on the creation of a suitable storage environment where the large amount of CineGrid data can be saved and easily retrieved.
Proof of concept demonstrations of the research and development described above were presented at: the GLIF meeting in Prague in Sep.2007, the e-Challenges conference in The Hague in Oct. 2007, Terena conferences in may 2007 and 2008, and the SuperComputing conference SC07 in Reno NV in Nov.2007 and SC08 in Austin TX in nov 2008.

Authorization Concepts and Architectures

In the AAA field our research focussed on the concept of using tokens to enable access control, workflow management and resource management for lightpath sessions within optical networks. A conceptual demonstration using the US DRAGON projects code during Supercomputing 2006 proved that the token concept could be applied within the GMPLS signaling plane of a lightpath. Driven by the interest shown by Internet2, we decided in 2007 to integrate the token concept into the Supercomputing 2007 demonstration of the Internet2 Dynamic Circuit Network (DCN). We hereto developed a Token Validation Service (TVS) which modules were integrated into the InterDomain Controller (IDC) of the DCN. The demo showed a network using multiple domains, each with their own IDC communicating with each other. The tokens were transported using the inter-IDC protocol based on Web Services mechanisms, where each domain interacts in a chain: The first domain in the chain generated a Global Resource Identifier (GRI) which was signed by the last domain in the chain, indicating that all intermediate domains in the chain are able to make reservations honor the request. The signed GRI acted as a commit of the entire service chain in the form of a token. The token would be passed back to the first domain in the chain and replied to the user. A copy of the token would be handed to the signaling part of the network, enabling subsequent enforcement at the moment the user inserts the token into the signaling to proof that the application has the right to use the lightpath. Our research and development in the area of applying tokens in the dataplane using in-band enforcement continued. Based ib a Intel Network Processor Unit (NPU) development platform, we implemented a Token Based Switch allowing packets, containing valid tokens, to be switched to a specific port, that could be connected to a lightpath. The IETF FORCES architecture, which describes ways to build a router from functional element, was used to implement the token handling functions.

Complex Resource Provisioning

The Generic AAA Authorization Framework architecture and its toolkit functionality aims to support Complex Resource Provisioning (CRP) in Grid based applications, including Grid enabled Network Resource Provisioning. The management of the security context and authorization sessions in multi-domain applications have specific attention. Further research has been conducted to use the Trusted Computing Platform Architecture (TCPA) for inter-domain trust management and ensuring virtual workspace and environment trustworthiness. This initiative uses the proposed authorization session management functionality that allows binding application related security context to those of virtual environment and trusted platform. Proposed solutions are being implemented in the framework of the Phosphorus project, by extending the AAA Toolkit AuthZ framework. Contributions were also made to enable AuthZ session management and policy obligations handling by extending EGEE gLite Java Authorization Framework (gJAF). The functionality has been implemented as pluggable AAA Toolkit components to support CRP. The newly developed components provide extensive support for XACML based access control policies expression and evaluation. The implementation also provides support for authorization session management wit the XML based AuthZ ticket that uses both a proprietary and a SAML based format. Actively contributed to Grid AuthZ interoperability initiative between EGEE, OSG and Globus consortia and projects.

Security and privacy

As part of our cooperation within VL-e, we initiated the development of a framework for managing privacy-sensitive data storage, access control and data distribution in the Grid. A paper on this framework has been accepted for the workshop on Security, Trust and Privacy in Grid Environments at IEEE CCGRID 2008.

Sensor Grids

It is already a mainstream vision that orders of magnitude more sensors and actuators than humans will interact through the internet with themselves or with applications. With this vision in mind the SNE group started research into the area of sensor (and actuator) networking. Complementary to the popular studies of battery saving, self organizing sensor-to-sensor and sensor-to sink communications SNE embarked on the field of sensor telecommunications. In this field one studies the

development of internet technologies that allow an application specific behavior of the internet. Here one strives to program a part of the Internet in such a way that it optimizes, in an application specific manner, the way how groups of sensors communicate to groups of traditional internet access points (that connect to applications). Typically the programmable part of the internet is most likely a set of autonomous systems with mobile links, forming a wireless networks that span huge geographic distances. These networks will not be over dimensioned as costs will inhibited that. In this situation the behaviors and topologies of sensors, applications and networks must be optimized as a system. SNE addresses the system optimization issue through Grid concepts. In this case network elements, sensors and applications are regarded as resources and are represented by objects in computer programs that interact with them. In 2006 SNE started this research, concentrating on the creation of a concept that represents ("virtualizes") network elements as objects in computer programs. In 2006 a collaboration on the theme of "Intelligent Sensor Networks" with the Dutch Institute for applied research TNO lead to a joint knowledge centre on this topic. In the Centre for Intelligent Observation Systems (CIOS) the research of programmable networks will be continue where a PhD study starts on the subject of ICT system programming. This R&D activity develops methods how specifications of the time behavior of sensors, applications, computers and network elements can be translated in actions of system elements to contribute to that behavior. Closely related to the work of sensor networks, intelligent observation systems is the subject of decision making on basic of sensor and human inputs. Results of SNE's work on sensor networks will disseminate into projects as the IJkdijk, that develops for the Dutch system of 17000 km of dike a sensor infrastructure, including an internet based, adaptive and programmable telecommunication infrastructure (www.ijkdijk.nl).

Future research plans

The year 2008 marks the end of the GigaPort project. The SNE group has introduced hybrid networking as a paradigm to give Internet communications for large sciences one more means of scaling while using 10's of gigabits of network links. The technology will approach Terabit/s capacities soon. The coupling of that kind of networks to computational, data and visualisation resources will require innovation in the ways the applications can express their needs to the infrastructure. Current protocols and API's can not handle that. Protocol stacks and latency issues hinder the applications to optimally use the connectivity capacity and functionality. Important themes in the next years will be: semantic descriptions of networked resources and their relationships, data security and privacy, and green-it.

Key Publications 2007-2008

Cees de Laat, Paola Grosso, "Lambda Grid developments, History - Present - Future", Lighting the Blue Touchpaper for UK e-Science - Closing Conference of ESLEA Project, Proceeding of Science, june 2007.
http://pos.sissa.it//archive/conferences/041/027/ESLEA_027.pdf

Jeroen van der Ham, Freek Dijkstra, Paola Grosso, Ronald van der Pol, Andree Toonk, Cees de Laat, "A distributed topology information system for optical networks based on the semantic web", Elsevier Journal on Optical Switching and Networking,Volume 5, Issues 2-3, June 2008, pp 85-93.

Freek Dijkstra, Bert Andree, Karst Koymans, Jeroen van der Ham, Paola Grosso, Cees de Laat, "Multi-Layer Network Model Based on ITU-T G.805", Computer Networks, Volume 52, Issue 10, 16 July 2008, Pages 1927-1937.

Paola Grosso , Li Xu, Jan-Philip Velders, Cees de Laat, "StarPlane - A National Dynamic Photonic Network controlled by Grid Applications", Emerald Journal on Internet Research, 2008, Vol 17, Issue 5, pp 546-553.

Yuri Demchenko, Olle Mulmo, Leon Gommans, Cees de Laat, Fred Wan, "Dynamic Security Context Management in Grid-based Applications", Future Generation Computer Systems, vol 24 (2008) 434-441.

Guido J. van 't Noordende, Silvia D. Olabarriaga, Matthijs R. Koot, Cees Th.A.M. de Laat, "A Trusted Data Storage Infrastructure for Grid-based Medical Applications", Proc. Eighth IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2008), Lyon, France, pp. 627-632.

Zhiming Zhao, A. Belloum, C. de Laat, P. Adriaans, B. Hertzberger, "Distributed execution of aggregated multi domain workflows using an agent framework", 2007 IEEE Congress on Services, Salt Lake City, UT, 9-13 July 2007, ISBN: 978-0-7695-2926-4, page 183 – 190.

3.3 Contribution of UvA-FNWI-IAS

University of Amsterdam, Faculty of Science, Informatics Institute, Intelligent Autonomous Systems
Prof.dr.ir. F.C.A. Groen, Prof.dr. D.M. Gavrila, Dr.ir. B.J.A. Kröse, Dr.ir. L. Dorst, Dr. G. Pavlin

Research 2007-2008

We focus on intelligent autonomous systems operating in a world inhabited by humans, explored within two themes: 'Perception and Modeling' of the world, and `Decision Making' to use these models for goal-directed behaviour.

Perception and Modeling
This theme develops cognitive sensors that enable machines to interact intelligently and naturally with a human-inhabited environment. We want to localize humans in their environment, track them over time, and recognize their activities. Also, we are interested in modelling the environment the humans are living in, typically from distributed or moving sensor systems. We have a number of projects addressing these issues.

In the project CARE (Context Awareness in Residences for Elderly), we inferred 'Activities of Daily Life' (ADL's) from a network of simple sensors in an elderly home. This project is a collaboration between the Centre for Intelligent Observation Systems (CIOS), the Hogeschool van Amsterdam and an elderly care institute. In the context of the 'Cogniron' project (a European FP-6 Integrated Project) we developed novel methods for spatial cognition and space learning for a 'cognitive' robot.

In the MultimediaN – Professional Dashboard project we perform 3-D human pose estimation from multi-view images. We developed a system that combines single-frame pose recovery, temporal integration and (texture-based) model adaptation. A second project in MultimediaN uses Dynamic Bayesian Networks to outperform existing methods on speaker diarization using audio-visual information.

In an STW Valorisation Grant, we are making a pilot system to test the commercial viability of our system to track humans using distributed non-overlapping cameras based on an efficient way to perform EM optimization.

For ego-motion estimation of moving robots in collaboration with TNO, we developed a novel fast and robust outlier-filtering method called EM-SE(3) that utilizes Expectation Maximization (EM) on a local linearization of the rigid body motions group i.e. SE(3).

For the 3D reconstruction of environments and (articulated) structures, we are extending the data processing aspects of the new field of geometric algebra. in the NWO project 'Discovery of Articulated Structures in Image Sequences'. To disseminate the foundation of these new techniques, we completed a tutorial book 'Geometric Algebra for Computer Science'.

Decision Making research theme
This theme focuses on the action side of the complete perception-action loop of an intelligent autonomous system, which requires robust handling of uncertainty due to incomplete information. We are interested in real-world applications with direct relevance to society. Examples include crisis management, wherein intelligent systems must analyze crises and help human managers make decisions in real time; rescue robot teams, which help people in disaster areas; and traffic management, which can be optimized via intelligent agents. Funding is obtained from the BSIK ICIS project and the FP7 DIADEM project. These projects run in the context of the D-CIS lab, collaboration with TU-Delft, Thales and TNO.

We developed methods to allow traffic control systems to better handle missing perceptual information or partial observability, using belief state estimation techniques and combining these with the value functions estimated using reinforcement learning.

Our theoretical research in distributed perception has focused on sound Bayesian information fusion which is particularly relevant for situation assessment problems in crises. We derived design and assembly rules for modeling components which guarantee correct belief propagation using factor graphs. The new theory has been implemented in the Distributed Perception Networks (DPN) architecture. Through an information theoretic approach we can test conditional independence relations between certain random variables in different Bayesian modules to discover crucial Bayesian module interdependencies which can significantly improve classification. Applications are in the FP7 DIADEM project on gas-detection in collaboration with DCMR.

Our intelligent Decision Making methods are also being applied in the RoboCup international robot competitions, notably the Virtual Rescue League aims to develop robot control systems in crisis environments. We re-implemented our previously developed simultaneous localization and mapping (SLAM) to allow real-time creation of an occupancy grid-map and steer

multiple robots through an unknown environment. Our effective methods have had us consistently end up in the top three places in the past few years.
In the problem of decentralized planning under uncertainty for teams of collaborating agents, we have focussed on the case where the agents are fully cooperative. Finding a solution for the corresponding stochastic game model (known as DEC-POMDP model) is NEXP-complete. We have provided theory concerning the existence of an optimal value function for DEC-POMDPs, as well as tighter upper bounds to this optimal value function. We also proposed algorithms that use these tighter bounds to either obtain more accurate approximate results or to find exact solutions faster. We collaborated with international partners to develop a multi-agent decision library, which provides functionality for planning in multi-agent domains, as well as a set of benchmark problems.

Future research plans

Perception and modelling
We will continue the collaboration with TNO in the expertise Centre on Intelligent Observation Systems (CIOS). The field lab 'caring and comfortable neighbourhood' is a collaboration with the Institute for Information Engineering of the Hogeschool van Amsterdam in Almere, where a 'Digital Life Centre' was started jointly with UvA. In the "Looking at People" area, we will develop 3D human pose estimation algorithms that are robust to occlusion and work with several people in the scene. We aim to integrate the richer 3D pose human pose features (MultimediaN – Professional Dashboard) into our aggression detection system (CASSANDRA). In the field of image interpretation, we will develop the geometric data processing capabilities of geometric algebra and integrate this with the effort in 'Looking at People' and 3D reconstruction of crime scenes together with NFI (National Forensic Institute).

Decision Making
We will continue our collaboration on this theme within the D-CIS lab. We plan to continue our efforts to develop novel methods for efficiently finding solutions to DEC-POMDPs. In particular, we intend to extend our work on factored DEC-POMDPs to create a more general solution method that uses collaborative graphical Bayesian games to exploit locality of interaction. Furthermore, we plan to continue our work developing robust neuro-evolutionary controllers for helicopter hovering. Our existing method is robust to changes in wind levels in the helicopter's environment, but we aim to extend this method to be robust to a wider range of fluctuations in the environment dynamics.

Key Publications 2007-2008

Dorst, L., Fontijne, D.H.F. & Mann, S. (2007). Geometric Algebra for Computer Science: An Object-Oriented Approach to Geometry. San Francisco, USA: Morgan-Kaufmann Publishers.

Dimitrakakis, C. (2008). Exploration in POMDPs. OGAI-Journal, 27(1), 24-31.

Gavrila, D.M. & Munder, S. (2007). Multi-Cue Pedestrian Detection and Tracking from a Moving Vehicle. International Journal of Computer Vision, 73(1), 41-59.

Oliehoek, F.A., Spaan, M.T.J. & Vlassis, N. (2008). Optimal and approximate Q-value functions for decentralized POMDPs. Journal of Artificial Intelligence Research, 32, 289-353.

Withagen, P.J., Groen, F.C.A. & Schutte, K. (2007). Ccd color camera characterization for image measurements. IEEE Transactions on Instrumentation and Measurement, 56(1), 199-203.

Zivkovic, Z., Booij, O., Kröse, B.J.A., Topp, E.A. & Christensen, H.I. (2007). From sensors to human spatial concepts: an annotated data set. Journal of Robotics and Autonomous Systems, 55(5) , 411-418.

3.4 Contribution of UvA FNWI-ISIS

University of Amsterdam, Faculty of Sciences, Informatics Institute, Intelligent Sensory Information Systems Group
Prof.dr.ir. A.W.M. Smeulders, Dr.ir. R. van den Boomgaard, Dr. M. Worring, Dr.Ing. J.M. Geusebroek, Dr. Th. Gevers, Dr. C.J. Veenman, R.F. Aldershoff, Drs. C.G.M. Snoek, Dr. N. Sebe

Research 2007-2008

We consider the research areas of content-based access of visual data, computer vision, and systems for retrieval of video, more specifically:

Visual search engines

For the retrieval of videos the aim is to make multimedia archives as accessible as their textual counterpart. To that end, our research efforts concentrate on automatic semantic indexing and interactive retrieval of multimedia sources. We have developed the MediaMill semantic video search engine, which uses a lexicon of detectable concepts in combination with several advanced user interfaces. To value the merit of our efforts on the highest international standards, all research is evaluated within the international TRECVID benchmark for multimedia retrieval, organized by NIST. In 2008, we obtained top rank performance in both the concept detection task and the supervised search task, and we secured the second rank in the automatic search task. Furthermore, we participated in the international PASCAL-VOC Visual Object Categorization challenge for image categorization, organized by the EU funded PASCAL network. We obtained top rank performance in the VOC challenge. These results, in such strong international benchmarks, confirm our research efforts over the past few years have given us a leading position in the field of image and video retrieval.

Computer vision and visual cognition

An import asset in our endeavor for semantic access is indexing of semantic concepts. One approach is scene categorization by modeling ambiguity in the popular codebook approach. There are two drawbacks to the traditional codebook mode, which stem from the hard assignment of visual features to a single codeword. We have studied the learning of soft relations between visual words, and the effect on categorization of many popular datasets.

Another line of research is on emotion recognition and categorization. We considered the question: can a machine learn to perceive emotions as evoked by an artwork? Therefore, we proposed an emotion categorization system, trained by ground truth from psychology studies. The training data contains emotional valences scored by human subjects on the International Affective Picture System (IAPS), a standard emotion evoking image set in psychology. Our approach is based on the assessment of local image statistics which are learned per emotional category using support vector machines.

Colour in computer vision

Our research on colour has led to an evaluation of edge-based colour constancy algorithms. The goal of color constancy is to measure image colors despite differences in the color of the light source. Traditionally, the computational method of obtaining this ability is by using pixel values only. Recently, methods using edges instead of pixel values have been proposed. However, different edge types exist, such as material, shadow and specular edges. Therefore, we analyzed the influence of different edge types on the performance of edge-based color constancy. We have shown that, on generated data without color clipping, specular edges deliver near-perfect color constancy and that shadow edges are more valuable than material edges. However, with color clipping, the performance using the specular edges decreases significantly, while the performance using the material or shadow edges is less affected.

Computer vision for humans

The main direction of research here addresses the problem of sensing and understanding users' interactive actions and intentions for achieving multimodal human-computer interaction in natural settings. To that end, we developed and evaluated and algorithm for accurate eye center location and tracking. The ubiquitous application of eye tracking is precluded by the requirement of dedicated and expensive hardware, such as infrared high definition cameras. Our aim is to perform very accurate eye center location and tracking using a simple web cam. The proposed method makes use of isophote properties to gain invariance to linear lighting changes (contrast and brightness), to achieve rotational invariance and to keep low computational costs. We have tested our approach for accuracy and robustness using the BioID and the Yale Face B databases. Our system can achieve a considerable improvement in accuracy over state of the art techniques.

Spatial and extensible databases

The activities are realized in close co-operation with the CWI-database group. Amongst the results obtained the Monet database kernel and its modules for image and geo-spatial reasoning stands out. In the area of database kernels, an innovative experimental analysis uncovered the lack of performance improvement in database technology over the last decade. The underlying reason is the relative progress in CPU- and RAM-technology, which shows an increasing performance bottleneck. This observation has led to novel techniques to measure the resource waste and new database algorithms to avoid resource stales. The open-source version of the MonetDB system has been downloaded more than 30.000 times.

Future Research Plans

Colour in computer vision will be structurally positioned within the Faculty by the recent VICI award for Dr. Gevers.

For video search engines, we are starting a close collaboration with the Faculty of Humanities, aiming at combined research in a Centre for Content Culture and Technology. Furthermore, we intent to collaborate with the Cognition Centre of the university on various aspects of vision. For computer vision and learning, we intent to concentrate on recognition and similarity-learning of scenes, objects, and actions.

We will continue the line of research as initiated for computer vision for humans and for visual data mining for safety. Recent EU funding has strengthened this line of research.

Key Publications 2007-2008

A. Diplaros, N. Vlassis and Th. Gevers. A Spatially Constrained GenerativeModel and an EM Algorithm for Image Segmentation. IEEE Trans. on Neural Networks, 18(3), 798-808, 2007.

J.C. van Gemert, J.-M. Geusebroek, C.J. Veenman, A.W.M. Smeulders, Kernel Codebooks for Scene Categorization, ECCV, Volume 3, page 696-709, 2008.

H.T. Nguyen, Q. Ji, and , A.W.M. Smeulders, Spatio-Temporal Context for Robust Multitarget Tracking. IEEE Transactions on Pattern Analysis andMachine Intelligence, 29, 52-64, 2007.

K. van de Sande, Th. Gevers, C. G. M. Snoek, Evaluation of Color Descriptors for Object and Scene Recognition, IEEE Conference on Computer Vision and Pattern Recognition, 2008.

Cees G. M. Snoek, M. Worring, O. de Rooij, K. E. A. van de Sande, R. Yan, A. G. Hauptmann, VideOlympics: Real-Time Evaluation of Multimedia Retrieval Systems, IEEE Multimedia, Volume 15 (1), page 86-91, 2008.

H.M.G. Stokman and Th. Gevers, Selection and Fusion of Colour Models for Image Feature Detection. IEEE Trans. on Pattern Analysis and Machine Intelligence, 29, 371-381, 2007.

R. Valenti and Th. Gevers, Accurate Eye Center Location and Tracking Using Isophote Curvature, IEEE Conference on Computer Vision and Pattern Recognition, 2008.

A. Gijsenij, Th. Gevers, M. Lucassen, A Perceptual Comparison of Distance Measures for Color Constancy Algorithms, European Conference on Computer Vision, Volume 1, page 208-221, 2008.

J. van de Weijer, Th. Gevers and A. Gijsenij, Edge-Based Colour Constancy. IEEE Trans. on Image Processing, 16, 2207-2214, 2007.

M. Worring and G. Schreiber, Semantic, image and video indexing in broad domains, IEEE Transactions on Multimedia, 9, 909-911, 2007.

3.5 Contribution of UvA-FNWI-CSA

University of Amsterdam, Faculty of Sicence, Informatics Institute, Computer Systems Architecture Group
Prof.dr. C. Jesshope, Dr. A.D. Pimentel

Research 2007-2008

The CSA research group is tackling significant conceptual problems in computer architecture while maintaining a firm grounding in the engineering constraints that future integrated circuits impose. The main issue is concurrency at all levels of systems from the instruction-level upwards. The work spans from compilers to computer architectures including their design and implementations as systems-on-a-chip. Our strategy with respect to concurrent computer architecture therefore, is to promote the conceptual design of systems that exploit concurrency without intervention from the end user providing deterministic and deadlock-free programming solutions. Our policy is to achieve concurrent instruction issue such that: the silicon area used is proportional to issue width; the power dissipated is proportional to the issue width; and the performance is proportional to the power dissipated for a given clock frequency. Moreover, our policy is that such systems should be programmable from existing sequential code and should be backwards compatible, so that unmodified code runs on the new processors but that once translated using binary-to-binary translation or by recompilation of the source code may run on any number of processors to obtain either performance or power targets. This policy creates difficult targets to achieve but we have already made significant progress in defining abstractions that support this in silicon, with tool chains from sequential, data-parallel and functional languages to support these underlying architectural abstractions.

As a second major theme, our group investigates system-level design methodologies for the design of (embedded) multiprocessor system-on-chip (MP-SoC) architectures. In terms of strategy, this work focuses on system-level performance/power analysis and architectural design space exploration (DSE) during the very early stages of design, where design decisions have great impact on (the success of) the final product. To this end, we study analytical and simulative

techniques for system-level (performance and power) analysis, as well as search mechanisms for the actual exploration and pruning of the design space.

Future Research Plans

Future work on the first theme will focus on making our model completely generic, with autonomous systems for managing resources across diverse targets. This will promote a write-once deploy anywhere approach to parallel programming.
Future work on the second theme will focus on technology-aware system-level DSE, and DSE for adaptive MP-SoCs. More specifically, we intend to address topics such as scenario-based workload modeling, scenario-aware DSE, system-level power modeling, modeling and exploration of dynamic task mapping, and clustering techniques for design space pruning.

Key Publications 2007-2008

CR Jesshope (2008) A model for the design and programming of multi-cores, in Advances in Parallel Computing, 16, High performance Computing and Grids in Action (Ed. L. Grandinetti), IOS Press, ISBN 978-1-58603-839-7, pp37-55.

C.R. Jesshope, “Operating Systems in Silicon and the Dynamic Management of Resources in Many-core Chips”, in Parallel Processing Letters (PPL), pp. 257-274, Vol. 18 (No. 2), 2008

K. Bousias, L. Guang, C.R. Jesshope and M. Lankamp, “Implementation and evaluation of a microthread architecture”, Journal of Systems Architecture, Volume 55, Issue 3, March 2009, pp 149-161

M.W. van Tol, C. R. Jesshope, M. Lankamp and S. Polstra, “An implementation of the SANE Virtual Processor using POSIX threads”, Journal of Systems Architecture, Volume 55, Issue 3, March 2009, pp 162-169.

A.D. Pimentel, M. Thompson, S. Polstra and C. Erbas, "Calibration of Abstract Performance Models for System-level Design Space Exploration", in Journal of Signal Processing Systems for Signal, Image, and Video Technology, pp. 99-114, Vol. 50 (No. 2), Feb. 2008, Springer.

A. D. Pimentel, "The Artemis Workbench for System-level Performance Evaluation of Embedded Systems", in Int. Journal of Embedded Systems, pp. 181-196, Vol. 3 (No. 3), 2008.

C. Erbas, A. D. Pimentel, M. Thompson, and S. Polstra, "A Framework for System-level Modeling and Simulation of Embedded Systems Architectures", in EURASIP Journal on Embedded Systems, 2007, available online: DOI 10.1155/2007/82123.

H. Nikolov, M. Thompson, T. Stefanov, A. D. Pimentel, S. Polstra, R. Bose, C. Zissulescu, and E. F. Deprettere, "Daedalus: Toward Composable Multimedia MP-SoC Design", invited paper, in the Proc. of the ACM/IEEE Int. Design Automation Conference (DAC '08), pp. 574-579, Anaheim, USA, June, 2008.

M. Thompson, T. Stefanov, H. Nikolov, A. D. Pimentel, C. Erbas, S. Polstra, and E. F. Deprettere, "A Framework for Rapid System-level Exploration, Synthesis, and Programming of Multimedia MP-SoCs", in the Proc. of the ACM/IEEE/IFIP Int. Conference on Hardware-Software Codesign and System Synthesis (CODES+ISSS '07), pp. 9-14, Salzburg, Austria, Oct., 2007.

3.6 Contribution of TUD-EWI-ST-PGS

Delft University of Technology, Faculty of Electrical Engineering, Mathematics & Computer Science, Parallel and Distributed Systems Group
Prof.dr.ir. H.J. Sips, Dr.ir. D.H.J. Epema, Prof.dr.ir. A.J.C. van Gemund, Prof.dr. C. Witteveen, Dr. K.G. Langendoen

Research 2007-2008

In High Performance Computing, we concentrate on *grid computing* and *multi-core programming*. The research in *grid computing* focuses on the problem of scheduling and resource management in multicluster systems and grids, and is centred around the design, the implementation, and the analysis of the KOALA grid scheduler. Among other things, KOALA, which has been deployed in the Dutch National Research Grid system (the DAS), supports processor co-allocation, load balancing, and cycle scavenging. Other research activities include the design of a grid workloads archive, and the investigation of grid interoperability and prediction methods.

In *parallel programming* we are doing research in parallel languages and programming environments, more specifically in languages and compilation techniques for distributed memory architectures such as multicore systems. We focus on HPC-extensions to Java (SPAR) and compilers that semi-automatically generate code for distributed memory systems. More recently, we focus is stream processing, which stems from both consumer electronics applications and scientific applications where data streams are generated by a large amount of sensors (e.g., radio telescopes).

In Distributed Systems, we target *peer-to-peer systems* and *wireless sensor networks*. The research in *peer-to-peer networks* focuses on adding social features such as friends and taste buddies, and adding support for Internet TV and video distribution (both live and video on-demand) to BitTorrent, which is one of the most popular p2p systems. Among the most important features we concentrate on are efficient gossip-based algorithms for doing recommendations for content, support for improved download performance and streaming videos across p2p systems, IP support for the notion of friends, and decentralization of a number of p2p management and content&resource discovery mechanisms.

In *wireless sensor networks*, we focus on the development of new protocols and algorithms for the efficient management of the resource- and energy-limited sensor nodes making up such networks. In particular, we are interested in developing medium access control (MAC) protocols with duty cycles below 1%, opportunistic routing protocols to leverage "long" links, and localization algorithms for highly mobile usage scenarios. To assess the performance, as well as the impact of the unpredictable wireless medium on such protocols, we have created a test bed for experimental use that has the unique capability of measuring the power consumption of all nodes in real time.

Future Research Plans

In grid computing, we will extend our research to cloud computing and to new application areas, such as massive multiplayer online games. In addition, we will study performance guarantees in grids and clouds, and the occurrence of failures in such systems.

In parallel programming, future focus will be on programming and optimization strategies for multicores and system-wide dimensioning techniques for complete applications that include multicore components.

In peer-to-peer systems, future focus will be on measurements, understanding, and improvements of BitTorrent, distributed trust management, freeriding prevention, spam elimination, and the removal of central component structures. Furthermore, the integration and applicability of our BitTorrent based systems substrate with existing web services will be investigated.

In wireless sensor networks, we are participating in the build-up of (federated) test beds for public use, such that algorithms can be tried out under different conditions before engaging in a real deployment. This infrastructure will be used to develop localization/tracking algorithms for indoor situations, in-network data processing algorithms, and MAC protocols handling burst traffic. Also the connection with the fixed, IP-based Internet will be investigated.

Key Publications 2007-2008

A.I.D. Bucur and D.H.J. Epema, Scheduling Policies for Processor Co-Allocation in Multicluster Systems, *IEEE Trans. on Parallel and Distributed Systems*, Vol. 18, pp. 958-972, 2007.

H.H. Mohamed and D.H.J. Epema, KOALA: A Co-Allocating Grid Scheduler, *Concurrency and Computation*, Vol. 20, pp. 1851-1876, 2008.

Fabrizio Petrini, Gordon Fossum, Juan Fernández, Ana Lucia Varbanescu, Michael Kistler, Michael Perrone, Multicore Surprises: Lessons Learned from Optimizing Sweep3D on the Cell Broadband Engine, *IPDPS* 2007.

Ana Lucia Varbanescu, Alexander S. van Amesfoort, Tim Cornwell, Andrew Mattingly, Bruce G. Elmegreen, Rob van Nieuwpoort, Ger van Diepen, and Henk J. Sips, Radioastronomy Image Synthesis on the Cell/B.E.. *Euro-Par 2008*, August 2008.

P. Garbacki, D.H.J. Epema, and M.R. van Steen, Optimizing Peer Relationships in a Super-Peer Network, *Int'l Conference on Distributed Computing Systems* (ICDCS), June 2007.

J.A. Pouwelse, P. Garbacki, J. Wang, A. Bakker, J. Yang, A. Iosup, D.H.J. Epema, M. Reinders, M. van Steen, and H.J. Sips, Tribler: a social based peer-to-peer system, *Concurrency and Computation*, 20, pp. 127-138, 2008.

K.G. Langendoen, Medium Access Control in Wireless Sensor Networks, in *Medium Access Control in Wireless Networks*, Nova Science Publishers, H. Wu and Y. Pan (editors), pp. 535-560, 2008.

A.G. Baggio and K.G. Langendoen, Monte-Carlo Localization for Mobile Wireless Sensor Networks, *Ad Hoc Networks*, Vol. 6(5), pp. 718-733, 2008.

3.7 Contribution of TUD-EWI-MM-CGCC

Delft University of Technology, Faculty of Electrical Engineering, Mathematics & Computer Science, Computer Graphics and CAD/CAM Group
Prof.dr.ir. F.W. Jansen, Dr. W.F. Bronsvoort, Ir. F.H. Post, Dr. C.P. Botha

Research 2007-2008

Research in this programme is concerned with modelling of 3D objects for industrial products (CAD/CAM) as well as for virtual worlds (gaming), and with visualisation for scientific and medical applications.

Geometric modeling is concerned with the mathematical modeling of 3D shapes. We focus on feature modeling to enhance 3D geometry with functional (semantic) information to support analysis through all stages of 3D product development. Within our modeling approach, constraint specification and constraint solving plays a major role. Several new constraint specification, solving and checking techniques have been developed, including a new, more powerful method for solving geometric constraints and for solving topological constraints.

Visualisation is concerned with analysis and display of large data sets. The purpose is to provide insight by extracting important information from data sets, and supporting interactive exploration of the data. The focus of our research is on improved 3D interaction tools, on data reduction and data exploration techniques. The developed techniques are integrated within a virtual envrionment. A typical application is Large Eddy Simulations (LES) of cumulus clouds.

Medical visualisation is concerned with model generation and simulation from various medical imaging modalities, such as CT and MRI. Our focus is on the visualisation of multi-field data and higher order data such as diffusion tensor imaging. In addition, we develop surgical simulation tools for the pre-operative planning of shoulder replacement surgery, as well as new approaches for polyp detection in virtual colonoscopy. In 2008, we released DeVIDE, an extensive medical visualization software laboratory, as open source.

Game technology is concerned with advanced modelling for the next generation of entertainment and serious games, in particular the dynamic generation and consistency maintenance of virtual worlds. Ultimate goal is to effectively assist game level designers in expressing and consistently maintaining in a model of the virtual world all intent specified throughout the various iterations of the design process.

Future Research Plans

For the coming years, we will continue to work on the above topics. In modeling we will address the integration of design and analysis through incremental meshing techniques and idealised analysis models. For visualisation several projects has been initiated to visualize environmental data (point clouds) and climate data (e.g. for extreme weather predictions). Real-time simulation and integrated visualisation and computational steering is now feasible on desktop computers with GPU support. For video surveillance tasks, we are exploring the integration of video streams with synthetic and reconstructed 3D models. For medical visualization we started a project on molecular imaging in order to visualize biological processes at cellular and molecular level over multiple timepoints and across multiple subjects.

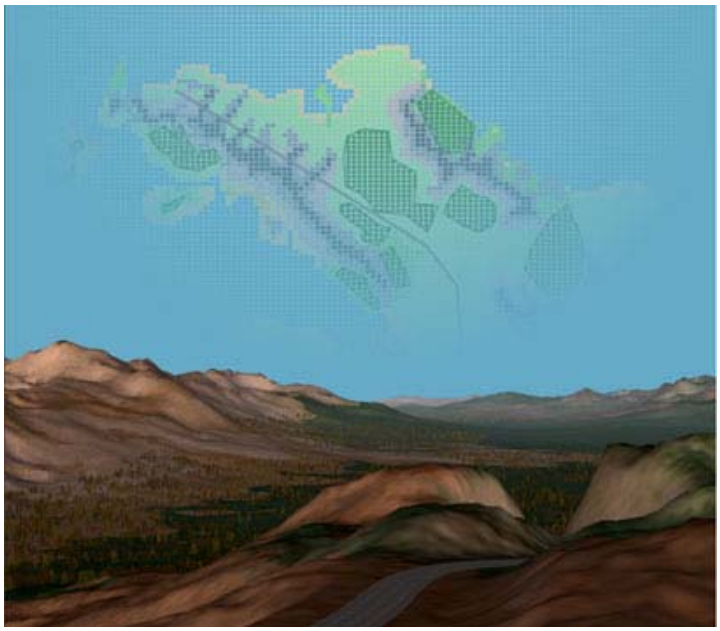

3D terrain generated on the basis of a rough sketch containing earth, road and vegetation information

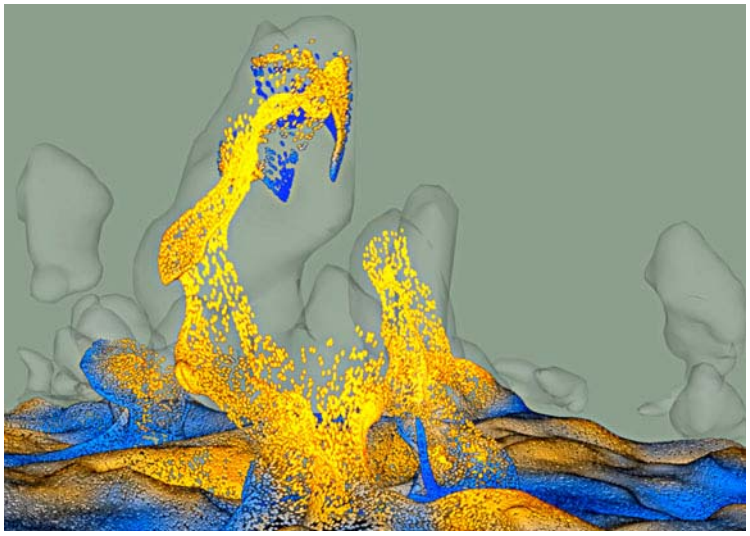

Visualisation of atmospheric simulation data using real-time particle advection. Millions of moving particles can be released in large time-varying data sets.

Key Publications 2007-2008

van der Meiden HA, and Bronsvoort WF (2007) Solving topological constraints for declarative families of objects. Computer-Aided Design 39(8): 652-662.

Gerwin de Haan, René Molenaar, Michal Koutek, Frits H. Post, Consistent Viewing and Interaction for Multiple Users in Projection-Based VR Systems, Computer Graphics Forum 26(3), September 2007, pp. 605-704 (ISSN 0167-7055).

Jorik Blaas, Charl P. Botha, Frits H. Post . Extensions of parallel coordinates for interactive exploration of large multi-timepoint data sets, IEEE Transactions on Visualization and Computer Graphics, 14(6), November/December 2008, pp. 1436-1443 (ISSN 1077 2626).

Tutenel T, Bidarra R, Smelik RM and de Kraker KJ (2008) The role of semantics in games and simulations. ACM Computers in Entertainment 6(4):a57.

Blaas, J, CP Botha, FH Post. Interactive visualization of multi-field medical data using linked physical and feature-space views. In: K Museth, T Msller, A Ynnerman (eds.), The joint eurographics/IEEE VGTC symposium on visualization. Norrkoping, Zweden, May 23 - 25, 2007, Eurographics Association, Zwitserland, 2007, p. 123 - 130.

Nyirenda, PJ & Bronsvoort, WF (2008). Numeric and curve parameters for freedom surface feature models. Computer-aided design, 40(8), 839-851.

Sypkens Smit, M & Bronsvoort, WF (2008). Variational tetrahedral meshing of mechanical models for finite element analysis. Computer-aided design and applications, 5(1-4), 228-240.

Zhao, L, Botha, CP, Truyen, R., Vos, FM & Post, FH (2008). Efficient seeding and defragmentation of curvature streamlines for colonic polyp detection. Proceedings SPIE, 6916.

3.8 Contribution of TUD-EWI-MM-ICT

Delft University of Technology, Faculty of Electrical Engineering, Mathematics & Computer Science, Department of Mediamatics, Information and Communication Theory Group
Prof.dr.ir. J. Biemond, Prof.dr.ir. R.L Lagendijk, Dr. E.A Hendriks, Dr.ir. J.C.A. van der Lubbe, Prof.dr.ir. M.J.T Reinders, Dr. A Hanjalic, Dr.ir. D. de Ridder, Dr. D. Bellomo, Dr. M.J.L. de Groot, Dr.ir.J.R. Taal

Research 2007-2008

The mission of the ICT Group is "*to research and educate theory and algorithms for signal processing and pattern recognition of complex and/or large volumes of data".* The research that we carry out is fundamental in nature, but the choice of specific research topics is often guided by applications in two broad domains, namely multimedia and bioinformatics. The figure below illustrates how the research subjects of the group map onto our fundamental and application domains.

		Research Themes	
		Signal Processing	Pattern Recognition
Application Domain	Multimedia	Parametric compression Signal enhancement (3-D) Video analysis Secure signal processing Watermarking and fingerprinting	Content-based information retrieval Recommendation engines Hyperspectral imaging Similarity-based classification
	Bioinformatics	Medical image analysis In-vivo biomolecular imaging	Molecular classification Inferring genetic networks

We shifted attention in video compression to fundamental approaches in new domains such as multiple description coding (MDC) and distributed source coding of video. We contributed to channel-codes based MDC of video to the (inter-)nationally highly visible peer-to-peer (P2P) research using the open source Tribler software. Our research in distributed source coding of video has contributed to the understanding of the potential and limitations of low-complexity video encoding.

In content-based information retrieval we focused on affective content analysis/retrieval of video in which user emotions are factored into the retrieval queries in addition to purely content-driven queries. Our work in content-based (information) retrieval was increasingly combined with social networking aspects. Our triple synergy approach includes collaborative filtering and tagging in P2P and social networks. We introduced distributed collaborative filtering techniques in which user-item ratings are stored and exchanged only locally (a resulting Buddy Cast algorithm is the recommendation engine within the Tribler P2P software).

In the area of video analysis we focused on extracting depth, motion, poses and gestures from multi-viewpoint video recordings. Our unique sign language recognition research led to an electronic learning environment for young deaf children, which operates under ill-controlled environments. In the field of secure signal processing a prominent result is our secure face recognition solution, in which face recognition is carried out while keeping both the query face and face database private.

In the area of medical image analysis we looked into the fusion and analysis of mainly 3D heart data (ultrasound, MRI, CT, functional data). For applications like surveillance or traffic control we investigated detection, recognition and tracking of (human) objects (faces, bodies, cars) with a focus on unsupervised learning of object classifiers. In image recognition we concentrated on the integration of multiple representations, e.g. resulting from hyper-spectral sensors, filter banks or a multi-scale analysis. We successfully showed that combining spatial information with multiple image representations improves segmentation and recognition results.

In pattern recognition, we introduced geometric-based one-class classifiers and showed that this concept can be elegantly integrated in the support vector machine paradigm. Other major contributions dealt with severely imbalanced data problems, ROC analyses and cost optimizations of multi-class problems including reject strategies. A new topic emerged in which we study proximity-based representations, leading to discriminant functions that are built on the proximity matrix to a set of prototypical objects.

In bioinformatics we advanced molecular classification by including gene annotation knowledge and additional dataset from different labs when building classifiers. To integrate different high-throughput data sources we proposed different integration strategies and several selection and evaluation strategies to establish which data sources are important. Another major contribution was our systematic analysis of the impact of different environmental conditions on the regulation of genes. Finally a prominent aspect of our contributions is that we introduced the concept of kernel convolution to analyze genomic data arranged along sequence location.

Future Research Plans

In the coming years, research will be structured along 6 main directions:

- Multimedia and information retrieval, concentrating on the triple-synergy of multimedia content analysis, social-P2P networks, and users collaborative tagging.
- Signal and information processing, concentrating on theory and algorithms for distributed processing in sensor and actuator networks.
- Information security and privacy, concentrating on cryptography and privacy of RFID technology and secure signal processing in, for instance, social networks.
- Vision research focuses on segmentation and analysis of multidimensional data with applications to 3D imaging, biomedical imaging, social or human signal processing and surveillance.
- Pattern recognition studies the classical trinity of representation, generalization, and evaluation and focuses on developing tools and theories and gaining knowledge and understanding
- Bioinformatics focuses on data-driven analysis of high-throughput biomolecular measurements and integrating that with biological insights and prior knowledge

Key Publications 2007-2008

Landgrebe, TCW, RPW Duin, “Efficient multiclass ROC approximation by decomposition via confusion matrix perturbation analysis.” IEEE transactions on pattern analysis and machine intelligence, 30, 5, 2008, p. 810 – 822.

Patras, I, EA Hendriks, RL Lagendijk, “Probabilistic confidence measures for block matching motion estimation.” IEEE transactions on circuits and systems for video technology*, 17, 8, 2007, p. 988 – 995.*

Wang, J, AP de Vries, MJT Reinders, “Unified relevance models for rating prediction in collaborative filtering.” ACM transactions on information systems*, 26, 3, 2008, p. 16:1 - 16:41.*

Hanjalic, A, R Lienhart, W.-Y. Ma, J.R. Smith, “The Holy Grail in Multimedia Information Retrieval: So Close or Yet So Far Away?.” *Proceedings of the IEEE*, 96, 4, 2008, p. 541 - 547.

Prins, J., Z. Erkin, RL Lagendijk, “Anonymous fingerprinting with robust QIM watermarking techniques.” *EURASIP journal on information security*, 2007, 2007, p. 1 - 13.

Pouwelse, JA, PJ Garbacki, J Wang, A Bakker, J Yang, A Iosup, DHJ Epema, MJT Reinders, MR van Steen, HJ Sips, “Tribler: a social-based peer-to-peer system.” *Concurrency and computation: practice & experience*, 20, 2, 2008, p. 127 - 138.

Knijnenburg, TA, JH de Winde, PAS Daran-Lapujade, JT Pronk, MJT Reinders, LFA Wessels, “Exploiting combinatorial cultivation conditions to infer transcriptional regulation*.”* BMC genomics, 8, 2007, p. 25/1 - 25/16.

Klijn, CN, H Holstege, J de Ridder, X Liu, MJT Reinders, J Jonkers, LFA Wessels, “Identification of cancer genes using a statistical framework for multiexperiment analysis of nondiscretized array CGH data.” *Nucleic acids research*, 36, 2, 2008, p. 1 - 16.

3.9 Contribution of TUD-TNW-QI

Delft University of Technology, Faculty of Applied Physics, Imaging Science & Technology, Quantitative Imaging Group
Prof.dr. L.J. van Vliet, Dr. B. Rieger

Research 2007-2008

The Quantitative Imaging Group invents new image-based measurement principles through a combination of imaging physics and digital imaging leading to novel algorithms for image processing, image analysis, image reconstruction, and image recognition. We perform fundamental research with a focus on applications in Life Sciences and Health in a multi-disciplinary environment of technical and (pre-)clinical scientists and in close collaboration with leading industrial parties. The core competence of the group is on quantitative imaging, i.e. extracting quantitative information from multi-dimensional image data sets. Important contributions herein have been made to the field of digital measurement theory, super-resolution reconstruction, virtual colonoscopy, DTI imaging, and multi-object registration in the presence of global constraints.
In digital measurement theory we analysed the influence of sampling on the size distributions obtained with morphological sieves. This work provides sampling and processing recipes to avoid aliasing artefacts in the size distributions and was applied successfully to the characterization of physical properties of pharmaceutical tablets and dried crackers.
Combining breakthroughs in filtering, estimation and inverse problems have led to methods capable of achieving super-resolution reconstruction in undersampled image sequences. We have obtained super-resolution for moving point targets against a highly cluttered background and for medium sized moving objects, i.e. objects containing interior pixels with sufficient contrast to permit object registration.
In medical image analysis we have worked on virtual colonoscopy, analysis of DTI, and multi-object registration subject to global constraints. We realized computer cleansing for Virtual Colonoscopy and Computer Aided Detection of polyps in CT colonography, a diagnostic tool for minimally invasive early detection of colon cancer, the third leading cause of cancer-related death in the Western world. In the coming years we further improve our methodology by permitting limited patient preparations and low-dose CT recordings. Extension to MR colonography poses another image processing challenge to handle severe contrast inhomogeneities and tools for disease grading. In neuro image analysis we try to differentiate in an early stage between physiologically (aging) and pathologically (e.g. Alzheimer's disease) in the brain, we search for spatiotemporal biomarkers of neurological disorders using sophisticated image processing and pattern recognition methods in Diffusion Weighted MRI data.
Computational imaging offers nanometer resolution in 3D electron-tomographic image volumes of frozen hydrated biological specimen opens new avenues for studying the molecular machinery inside the cell. To achieve our goals we are developing quantitative forward models for designing new acquisition and reconstruction strategies, sparsity promoting image reconstruction techniques for solving underdetermined problems, and image processing algorithms capable of handling very noise data.

Future Research Plans

The next challenge in multi-frame super-resolution reconstruction is a super-resolution scheme for small moving objects, in which all object pixels mix the foreground information with a space-time variant background contribution in cluttered surroundings. In medical image analysis for CT colonography we expect to reach several milestones: to extend our surface-based technique for protrusion detection to a volumetric approach; to incorporate both techniques in a CAD system for polyp detection; to extend our CAD system to handle large lesions; and to develop novel algorithms for computer cleansing of various patient preparations. In diffusion weighted MRI we aim to develop advanced algorithms permitting the quantification of diffusivity in both constituents of crossing fibre bundles in the human brain. In computational microscopy we are developing novel algorithms for super-resolution. In optical microscopy we work on exploiting the blinking characteristics of fluorophors to disentangle overlapping point-spread-functions and in electron tomography we are working on better algorithms for modelling the imaging process.

Key Publications 2007-2008

I.W.O. Serlie, F.M. Vos, R. Truyen, F.H. Post, L.J. van Vliet, Classifying CT Image Data Into Material Fractions by a Scale and Rotaion Invariant Edge Model, IEEE Transactions on Image Processing, 16(12), 2007, 2891-2904.

A. W. M. van Eekeren, K. Schutte, O. R. Oudegeest, L. J. van Vliet, Performance evaluation of super-resolution reconstruction methods on real-world data, EURASIP Journal on Advances in Signal Processing, vol. 2007, Article ID 43953, 11 pages, 2007. doi:10.1155/2007/43953.

S.T. Nevo, M. van Stralen, A.M. Vossepoel, J.H.C. Reiber, N. de Jong, A.F.W. van der Steen, J.G. Bosch, Automated tracking of the mitral valve annulus motion in apical echocardiographic images using multidimensional dynamic programming, Ultrasound in Medicine & Biology, Vol. 33(9), 1389–1399.

J. Dijk, A.W.M. van Eekeren, K. Schutte, D.J. de Lange, L.J. van Vliet, Super-resolution reconstruction for moving point target detection, Optical engineering, 47 (09)(096401), 096401-1-096401-12, 2008.

I.W.O. Serlie, A.H. de Vries, L.J. van Vliet, C.Y. Nio, R. Truyen, J. Stoker, F.M. Vos, Lesion Conspicuity and Efficiency of CT Colonography with Electronic Cleansing Based on a Three-Material Transition Model, Am. J. Roentgenol. Vol. 191, no. 11, 1493-1502, 2008.

M.W.A. Caan, L.J. van Vliet, C.B.L.M. Majoie, E.M. Aukemea, C.A. Grimbergen and F.M. Vos, Spatial Consistency in 3D Tract-Based Clustering Statistics, in: D. Metaxas, L.Axel, G. Fichtinger, G. Szekely (eds.), Medical Image Computing and Computer-Assisted Intervention (MICCAI 2008), 2008, Lecture Notes in Computer Science, vol. 5241, Springer-Verlag Berlin Heidelberg, 2008, 535-542.

3.10 Contribution of UL-LIACS

Leiden University, Faculty of Mathematics and Natural Sciences, Leiden Institute of Advanced Computer Science (LIACS)
Prof.dr. H.A.G. Wijshoff, Dr. M.S. Lew, Dr. A.A. Wolters, Dr. D.P. Huijsmans, Dr. E.M. Bakker, Dr.ir. T.P. Stefanov, Dr.ir. B. Kienhuis, Prof.dr.ir. E.F. Deprettere

Research 2007-2008

The mission of LIACS research program 'Computer Systems and Imagery & Media' is to improve the state of the art in systems and media analysis, and to guide new research areas which are of clear importance to society. In our research, we rigorously demonstrate the effectiveness of the resulting novel techniques and contribute to the leading evaluation and benchmarking projects. This programme covers a wide variety of research topics, including embedded systems and software, parallel and distributed computing, and media research and technology, including imaging and its applications to bio informatics. The research effort is clustered around three themes.

Embedded Systems and Software

The research in embedded systems and software deals with modeling of applications and multi-processor architectures, and mapping methods in these domains, at various levels of abstraction, for exploration and design, theoretically and practically, down to real platforms. The working methodologies relies on advocating and applying modern state–of–the–art Software Engineering Practice both in the way the group's projects are integrated, documented, and assessed, and in the way software that implements research results is written, tested and assessed. There are several themes that are researched in several projects and are jointly contributing to the knowledge, expertise, and skills on local, national, and international competitive settings. The research focuses on strategically well-chosen core activities that are briefly described below:

- Research on Hardware dependent Software (HdS) solutions to improve IP integration in the SoC design process (quality / productivity). This includes: separation of the Operating System (OS) and the application software from the underlying hardware and HdS for efficiency, dependability, flexibility and manageability; systematic, highly automated HW/SW-integration of IPs.
- Research on design and application of multi-core processor architectures targeting tera-flops performance. Specific topics include: In-Network Cache Coherence, Protocol on NoC-based architecture, Scalable Tera Bit/s IO design, Application Specific GALS NoC design, optimal Compilation for Multi-Core/Multi-Thread Processor.
- Research on tools to help the parallelization and porting of applications on an MPSoC architecture. These tools will simplify the manual effort that is currently needed by the application developers to transfer sequential software in a concurrent programming model.
- Research on techniques based on a streaming application model for power and performance analysis considering voltage scaling, frequency control, and alternative mapping of application tasks to different parts of a parallel embedded platform.
- Research on novel methods, techniques, and tools for modeling and parallelization of adaptive streaming applications. Theoretical studies are performed to develop a modeling framework for multiple adaptive streaming applications to be executed simultaneously onto a single heterogeneous MPSoC.
- Research on novel methods, techniques, and tools for designing Embedded Multi-processor Systems-on-Chip with adaptivity and reliability support. The static approach to embedded MPSoC design will be extended to address adaptivity requirements by setting up a methodology and the corresponding tool support for adaptive (i.e., run-time) mapping of application tasks to the underlying architecture resources to cope with QoS and/or dependability demands.

Parallel and Distributed Computing Systems

The research in parallel and distributed computing concentrates on: (optimizing) compilers, grid computing, application drivers for large-scale applications, and large-scale database systems. As within the embedded system research activities the approach taken evolves around the mapping problem of (large scale) applications onto (existing) computing platforms. The research focuses on strategically well–chosen core activities that are briefly described below:

- Analysis and Modeling of the Circadian Pacemaker at the Cellular Level
  The project for developing computer models and performing large-scale simulation of the circadian pacemaker (SCN) at the cellular level started in 2004 as collaboration between the Leiden Institute of Advanced Computer Science (LIACS) and the Leiden University Medical Center (LUMC).
- Performance Predictions and Resource Management in a Grid
  To deliver nontrivial quality of service is one of the most challenging problems in the Grid. Specifically, mining performance data on Grid resources to extract useful information helps to improve performance or manage the system itself. In 2007 the project has been finalized and it resulted in a PhD thesis, which has been defended in January 2008.
- HIRLAM on a Grid Environment
  We made a start with an investigation on the performance of the operational numerical weather forecast system HIRLAM on the new DAS-3 grid environment. DAS-3 consists of five clusters located at five different locations. Futhermore, we investigated the possiblities in overlapping computations and communications for HIRLAM to decrease the communication overhead resulting due to explicit message passing.
- Integration, Analysis and Logistics (DIAL) project
  By combining the vast knowledge on a number of disease areas with the latest research technologies including genomics and bioinformatics, the aim is to elucidate the causes of various common diseases and apply these new insights in the development of new methods for diagnosis, new drugs and new means for prevention.
- Data Structure Independent Programming
  Traditional analyses break down when they encounter pointer based codes. Therefore, as the potential performance of processors increases, new means must be explored to extract parallelism. In the mid-nineties, research was conducted at LIACS to let the compiler decide on which data structures to use. We have extended these techniques to broaden the application domain to pointer-based codes.
- Reshaping Memory Access Patterns
  Using a combination of compile-time and run-time techniques, the access patterns of pointer based codes are restructured such that the resulting code is analyzable by existing techniques for regular codes. These steps result in an intermediate code which is amenable to data dependency analysis and as such can be compiled into a highly optimized executable.

Media Systems

Media systems evolves around two approaches. The first approach focuses on the scientific investigation of novel directions and paradigms in the field of multimedia retrieval with emphasis on content-based methods in images, video, audio, and scientific data. As a special application, the integration bio-imaging information and image information with other bio-molecular information resources is studied. The second approach is the embedding of human beings in a computerized world, which is studied by creating such embeddings.
Media systems evolve around the scientific investigation of novel directions and paradigms in the field of multimedia retrieval with emphasis on content-based methods in images, video, audio, and scientific data. As a special application, the integration bio-imaging information and image information with other bio-molecular information resources is studied.

In multimedia retrieval, we completed the first stage of research and experiments on the new paradigm of Artificial Imagination in multimedia information retrieval in the area of texture retrieval. The latest survey of the state of the art and grand challenges in the field was completed and published in the leading ACM Multimedia journal, and we served as advisers on the leading video analysis research evaluation project, TRECVID.

In the area of bioinformatics, we developed new algorithms for genotype/phenotype data mining. The development of an Hidden Markov Model based classifier for biological imagery was completed and tested under a wide variety of different tissue samples. A new method for viewing the development of micro-tubules was completed.

In the next phase, we will be investigating the integration of very large image and video databases towards creating a general artificial imagination model. The artificial imagination will be used for both multimedia retrieval and learning visual concepts.

Future Research Plans

The research activities changed over the last couple of years from a strong focus on high performance computer systems towards more and more data-oriented computations. Specific examples of these shifts can be found in the themes “Content

based image retrieval", in which the main challenge lies in indexing and classifying tremendous amounts of visual images, "HIRLAM on the Grid", in which the main challenge will be to integrate and reformat tremendous amounts of meteorological data distributed over a number of different networks (the Internet, the world-wide meteorological World Weather Watch network, etc.), "Data compilation", in which the direct target is the integration and coupling of large legacy database applications and "Hardware dependant Software (HdS) solutions to improve IP integration in the SoC design process (quality / productivity)", in which the operating system and the application software are separated from the underlying hardware.
As our programme has grown, our research activities have also evolved to address important new research areas and critical challenges in transferring technology from research labs to society. In particular, it is well known that research systems, which work well in a laboratory environment, may not work well in the real world. In some cases, this is because there was a significant lack of communication between the researchers and the practitioners/users. In other cases the problem has to do with developing high quality evaluation and benchmarking procedures, which reflect real world usage.
The changes as described above became firmly incorporated into the research activities in this programme. In fact our research activities became more focused on data representation, processing data representation, and the reduction of data representation by employing (semantic) transformation and interpretation of media data.

Key publications

H. Nikolov, T. Stefanov, E.F. Deprettere: Systematic and Automated Multiprocessor System Design, Programming, and Implementation. IEEE Trans. on CAD of Integrated Circuits and Systems 27(3): 542-555 (2008)
IEEE 2009 Donald O. Pederson best paper award

H.L.A. van der Spek, S. Groot, E.M. Bakker, H.A.G. Wijshoff: A Compile/Run-time Environment for the Automatic Transformation of Linked List Data Structures. International Journal of Parallel Programming 36(6): 592-623 (2008)

M.J. Huiskes, M.S. Lew. The MIR Flickr Retrieval Evaluation. ACM International Conference on Multimedia Information Retrieval (MIR'08), Vancouver, Canada, 39-43, 2008.

B. Thomee, M.J. Huiskes, E. Bakker, M.S. Lew. Large-Scale Image Copy Detection Evaluation. ACM International Conference on Multimedia Information Retrieval (MIR'08), Vancouver, Canada, 59-66, 2008.

Hristo Nikolov, Todor Stefanov, and Ed Deprettere, "Automated Integration of Dedicated Hardwired IP Cores in Heterogeneous MPSoCs Designed with ESPAM", EURASIP Journal on Embedded Systems, vol. 2008, Article ID 726096, 15 pages, 2008. doi:10.1155/2008/726096.

Hui Li, Michael Muskulus, and Lex Wolters, "Modeling Correlated Workloads by Combining Model Based Clustering and A Localized Sampling Algorithm", in Proceedings of 21st ACM International Conference on Supercomputing (ICS07), Seattle, USA, June 16-20, 2007, ACM Press.

3.11 Contribution of UU-ICS-GMT

Utrecht University, Faculty of Science, Department of Information and Computing Sciences, Center for Geometry, Imaging and Virtual Environments
Prof.dr. M.H. Overmars, Dr. M. van Kreveld, Dr.ir. F. van der Stappen, Dr. R. Veltkamp, Prof.dr P.J. Werkhoven

Research 2007-2008

In game technology, research was performed in path planning, crowd simulation, and manipulation planning. We developed new path planning approaches, partially based on the PRM method, for planning in dynamic and changeable environments. In dynamic environments the location of obstacles can change over time (for example cars drive around) and path planning has to take such changes into account to compute feasible paths. In changeable environments, the virtual character for which we plan routes is able to change the environment, for example by pushing an obstacle out of the way. This is a basic form of a combination of navigation and manipulation. We devised efficient algorithms that plan the correct way to move objects out of the way to reach a desired goal.

In crowd simulation the goal is to simulate the motions of large crowds of people. Such people can have a common goal, e.g. when people move in a large group towards a particular location, or people can simply wander around. We devised a new approach, based on the corridor map method, in which a group of characters can move to a particular goal. By using a new force model we can guarantee that the group stays coherent and does not split us (as you often see in other approaches, like flocking). We also developed algorithms to keep the group in certain formations or to let them obey certain tactical constraints. For simulating large crowds it is crucial that the underlying path planning approach is very fast. Our approach

based on the corridor map method has this property and enables us to simulate thousands of characters at interactive frame rates.

Research in manipulation planning has focused on two major tasks in industrial automation, namely part feeding and grasping. Part feeding is concerned with orienting of parts. The research on part feeding has concentrated on the design of smart mechanisms for the vibratory bowl feeder. We have studied the automated design of various classes of filtering mechanisms and removed several of the severe idealizations present in prior solutions. In addition we have proposed a new device, termed 'blade', that combines reorientation with filtering functionality, and developed algorithms for its automated design. The blade enables, amongst others, a drastic increase of the throughput of the vibratory bowl. The research on grasping has concentrated on the computation of immobilizing and caging grasps. For immobilizing grasps, preventing all motions of the grasped object, we have studied grasps that satisfy a century-old conservative condition as well as grasps that satisfy a recent necessary and sufficient condition. Caging grasps are loose grasps as they do not exclude all motions of the grasped object but just prevent it from escaping, which is sufficient for transportation purposes. We have studied algorithms for verifying whether a certain grasp of an object is caging, and for computing all caging grasps of an object.
In media technology, research was performed in the following areas: image retrieval, geometric algorithms, GIS algorithms, 3D modeling and recognition, and music retrieval. We have developed a number of different methods for shape recognition, shape indexing with vantage objects, and layout indexing with Laplacian and Hermitian graph spectra. Apart from provable properties, the behavior in practice has been rigorously evaluated on the basis of ground truth test sets, and compared with many other content-based image retrieval methods. In a collaboration with Yahoo! Research Barcelona, we have developed a method for query diversification, a recent topic to increase the diversity of retrieval result lists.

With respect to triangulations and terrains we have introduced and developed the concept of higher-order Delaunay triangulations as a class of well-defined triangulations. In a sequence of papers we have explored their properties and algorithmic possibilities, and also applied them to realistic terrain modeling. We have analyzed how geometric imprecision in the data can influence the results of a geometric computation, like the diameter, closest pair, area of the convex hull, and many more. A number of different algorithms to deal with imprecision have been developed. Concerning visibility in 3D we have developed extended notions of visibility in 3D scenes and developed algorithms for it. Further, we have defined new realistic input models in order to be able to prove better running times of geometric algorithms of various sorts.
We have developed a 3D model recognition based on the Earth Mover's Distance, and compared it to other methods. Another method we developed was more meant for indexing, not necessarily recognition. It is based on representing the object as a 3D skeleton, on whose matrix representation we can compute the spectrum, which is used for indexing. We have organized the 3D Shape Retrieval Contest (SHREC), which has a large impact on the community.
We have performed research on the mapping of music notation and audio into a symbolic representation. Then the symbolic representation is indexed and matched to a query. One specific application is to find folk songs that are similar to a query, yet not exactly the same, due to oral transmission. Our retrieval methods are integrated with the Nederlandse Liederenbank, the Dutch folk songs database at the Meertens Institute for ethnological research.

Future Research Plans

While we plan to extend results in the above areas, new research directions will be animation, multi-modal interaction, computer vision, and mobile interaction, and spatial data mining.
The animation research will mainly focuses on generating realistic virtual character motion. In the animation community, character motion is generally divided into facial and body animation, which require different approaches. In collaboration with TNO, we will start a new line of research in multi-modal interaction. The goal of this research is to develop new concepts and techniques with which users can steer the behavior of virtual characters. In particular we are interested in techniques with which users can control the navigation and manipulation performed by such characters
We will start working on detecting, tracking, modeling, and recognition of persons, their pose, and gestures, from multiple video cameras. Our approach is to exploit 3D information that we derive from coherence in views. Research on mobile interaction will be directed towards new, better ways for accessing, handling, and manipulating different kind of media on handheld devices. With respect to spatial data mining we will study concepts like flocking, meeting, convergence, etc., for set of trajectories, and developed algorithms to detect these patterns.

Key Publications 2007-2008

R. Geraerts, M.H. Overmars, The corridor map: A general framework for real-time high-quality path planning, Computer Animation and Virtual Worlds 18 (2007), 107–119.

M. Vahedi, A.F. van der Stappen, Caging polygons with two and three fingers, International Journal of Robotics Research 27 (2008), 1308–1324.

J. van den Berg, M.H. Overmars, Planning time-minimal safe path amidst unpredictably moving obstacles, International Journal of Robotics Research 27 (2008), 1274–1294.

D. Nieuwenhuisen, A. Kamphuis, M.H. Overmars, High quality navigation in computer games, Science of Computer Programming 67 (2007), 91–104.

T. de Kok, M. J. van Kreveld, M. Löoffler: Generating realistic terrains with higher-order Delaunay triangulations. Computational Geometry 36(1), 2007, 52-65.

M. Casey, R. Veltkamp, M. Goto, M. Leman, C. Rhodes, M. Slaney. Content-based music information retrieval: current directions and future challenges. Int. journal 'Proceedings of the IEEE', 96(4), 2008, 668-696. DOI 10.1109/JPROC.2008.916370.

Wolfgang Hurst. Video Browsing on Handheld Devices-Interface Designs for the Next Generation of Mobile Video Players. IEEE MultiMedia, 15(3), 2008, 76-83.

Fatih Demirci, Reinier van Leuken, Remco C. Veltkamp. Indexing through Laplacian Spectra. Computer Vision and Image Understanding, vol 110/3, 312-325, 2008. DOI 10.1016/j.cviu.2007.09.012.

3.12 Contribution of TUD-L&R-FRS

Delft University of Technology, Faculty of Aerospace Engineering Department of Earth Observation and Space systems (DEOS), Optical and Laser Remote Sensing Group
Dr.ir. B.G.H. Gorte

Research 2007-2008

The chair of Optical and Laser Remote Sensing (OLRS) of the Faculty of Aerospace Engineering at TU Delft is part of the Remote Sensing department, where the topic is application of space engineering with a focus on observation of the Earth and other planets.

Earth observation from satellites operating in the visible, near infra-red and thermal infra-red ranges of the electromagnetic spectrum is gaining significance with the increasing needs for high-frequency uniform monitoring of processes related to vegetation and the water and carbon cycles, against the background of increasing population pressure and climate change.

In addition to imagery originating from satellite sensors also airborne and 'close range' image data are analysed at OLRS, for example to measure and monitor driver behaviour in traffic congestion at motorways by analyzing image sequences recorded from a helicopter.

Interesting technology is provided by laser range measurement devices that nowadays can be operated 'close range', airborne and even from satellite (ICESat). These three forms are studied with similarly large interest for very diverse applications in medicin, forestry, flood modeling, urban climate modeling and land use/land cover mapping.

During the reporting period (2007/2008) the chair of OLRS was vacant, as were two of the lecturer positions. This resulted in a temporary stagnation of output production and project acquisition. Since December 2008 the chair is re-occupied. Subsequently the lecture positions were filled and the number of PhD projects increased from 6 to 13, with some more to start soon. This will result in an increasing ASCI-participation shortly.

Future Research Plans

The range of subjects studied at the chair is being extended. The computer vision/automatic interpretation based analysis of image and range data, as described above, will remain one of the the fields of focus. A new activity (within the chair) is physics based retrieval of parameters of geo- and biodynamical processes on the basis of multi-temporal, multi-angular and multi-spectral EM-radiation measurements, from both active and passive sensors. Naturally these two fields will not be seen in isolation. Their integration into a comprehensive modeling and monitoring strategy will be the focus of the chair during the coming years.

Key Publications 2007-2008

F. Karimi Nejadasl, B.G.H. Gorte, M. Snellen, S.P. Hoogendoorn. Energy function behanior in optimization based image sequence stabilization in presence of moving objects. XXI ISPRS conference, Beijing, China, 3-11 July 2008. Volume XXXVII, part B3b, commission III, p.p. 669-674.

M.Z.A. Rahman, B. Gorte. Individual tree detection based on densities of high points of high resolution airborne lidar. Geobia 2008 – Pixels, Objects, Intelligence GEOgraphic Object-Based Image Analysis for the 21st Century,University of Calgary, Calgary Alberta, Canada, August 05-08, 2008, pp 1-6.

D.L. Esme, A.Bucksch, Werner H. Beekman, 3D Laser imaging as a valuable tool to specify changes in breast shape after augmentation mammaplasty, Aesthetic Plastic Surgery, Vol 33, Nr. 2, March 2009, pp. 191-195.

Hieu van Duong, Roderik Lindenbergh, Norbert Pfeifer, George Vosselman. Single and two epoch analysis of ICESat full waveform data over forested areas. International Journal of Remote Sensing, Vol. 29, No. 5, 10 March 2008, pp. 1453-1473.

A. Bucksch and R. Lindenbergh. CAMPINO - A skeletonization method for point cloud processing. ISPRS Journal of Photogrammetry and Remote Sensing, Vol.63, No.1, 2008, pp. 115-127.

3.13 Contribution of RUG-CS-SVCG

University of Groningen, Faculty of Mathematics and Physical Sciences, Institute of Mathematics and Computing Science, Scientific Visualization and Computer Graphics
Prof.dr. J.B.T.M. Roerdink, Dr. H. Bekker

Research 2007-2008

The research group Scientific Visualization and Computer Graphics carries out research in the area of scientific visualization, information and software visualization, computer graphics and innovative interfaces using large, touch-sensitive displays. With respect to applications, the research concentrates on fundamental and applied problems from the life sciences, in particular functional brain imaging and bioinformatics, and astronomy.

In interactive data visualization, the speed of the data processing stage should be comparable to that of the visualization step. We address this demand by developing efficient algorithms and/or by mapping the involved computations to programmable Graphics Processing Units (GPUs), which are capable of outperforming CPUs for compute-intensive applications. A rapidly growing area within the visualization research field is perception-based visualization. Here one takes advantage of knowledge about the human visual system to improve current visualization techniques. An important issue is which `visual cues' of a scene (such as shape, size and distance of objects) can be used to encode independent information dimensions.

In Software Visualization and Program Understanding, our goal is to provide methods, techniques, and tools that assist the entire range of activities in the software engineering discipline. Effective use of software visualization in practice requires a tight integration of software visualization techniques and tools within the classical software analysis, forward, and reverse engineering pipelines.

Non-photorealistic rendering (NPR) is a sub-area of computer graphics that is inspired by a long tradition of artistic and illustrative depiction. We apply NPR techniques to illustration and visualization problems in medical, technical, and other domains. We employ novel touch-sensitive, large displays that enable users to make use of a larger screen area, interact with applications using direct-touch or pens, and work in groups to profit from collaborations.

Applications. In functional neuroimaging, we focus on functional MRI and diffusion tensor imaging (DTI), which is an MRI-based technique enabling the visualization of nerve fibers and connectivity of brain regions. In bioinformatics, we work on the visualization of gene expression data from time series experiments in both a gene regulatory network and metabolic pathway context. In astronomy, we use analytic and explorative visualization methods to study the relations between the spatial arrangement of galaxies and the distribution of various attributes in parameter space.

Future Research Plans

There are two new international initiatives in which the group wants to extend its efforts, i.e., Visual Analytics and Neuroinformatics. Visual Analytics concerns the integration of visualization with other analytical methodologies, such as statistics, data-mining, and cognition. The group already collaborates within the European project *Vismaster*, which intends to develop a European roadmap for Visual Analytics. There are also interesting connections to be explored with artificial intelligence and cognitive science. The second area is Neuroinformatics, which concerns the support of discovery and innovation in neuroscience by developing, maintaining and evaluating worldwide programs, databases, standards, guidelines and infrastructures in neuroinformatics to further our understanding of the human brain and its diseases. The group will actively extend its participation in local, national and international initiatives in this area.

Key Publications 2007-2008

Andrei C. Jalba and Jos B.T.M. Roerdink. Efficient surface reconstruction using generalized Coulomb potentials. IEEE Trans. Visualization and Computer Graphics, 13(6):1512-1519, 2007.

Ronald van den Berg, Jos B.T.M. Roerdink and Frans W. Cornelissen. On the Generality of Crowding: Visual Crowding in Size, Saturation, and Hue Compared to Orientation. Journal of Vision, 7(2):1-11, 7 2007.

Michel A. Westenberg, Jos B.T. M. Roerdink and Michael H.F. Wilkinson. Volumetric attribute filtering and interactive visualization using the Max-Tree representation. IEEE Transactions on Image Processing, 16(2):2943–2952, 2007.

Michael ten Caat, Natasha M. Maurits and Jos B.T.M. Roerdink. Design and Evaluation of Tiled Parallel Coordinate Visualization of Multichannel EEG Data. IEEE Trans. Visualization and Computer Graphics, 13(1):70-79, 2007.

Evert-Jan Blom, Dinne W.J. Bosman, Sacha A.F.T. van Hijum, Rainer Breitling, Lars Tijsma , Remko Silvis, Jos B. T. M. Roerdink and Oscar P. Kuipers. FIVA: Functional Information Viewer and Analyzer extracting biological knowledge from transcriptome data of prokaryotes. Bioinformatics, 23(9):1161-1163, 2007.

Wim H. Hesselink and Jos B.T.M. Roerdink. Euclidean skeletons of digital image and volume data in linear time by the integer medial axis transform. IEEE Trans. Pattern Analysis and Machine Intell., 30(12):2204–2217, December 2008.

Heorhiy Byelas and Alexandru Telea. Towards Visual Realism in Drawing Areas of Interest on Software Architecture Diagrams. J. of Visual Languages and Computing, 20(2):110–128, 2008.

Ronald van den Berg, Frans W. Cornelissen, and Jos B.T.M. Roerdink. Perceptual Dependencies in Information Visualization Assessed by Complex Visual Search. ACM Transactions on Applied Perception, 4(4):1–21, 2008.

Michael ten Caat, Natasha M. Maurits, and Jos B.T.M. Roerdink. Data-Driven Visualization and Group Analysis of Multichannel EEG Coherence with Functional Units. IEEE Trans. Visualization and Computer Graphics, 14(4):756–771, 2008.

3.14 Contribution of RUG-CS-IS

University of Groningen, Faculty of Mathematics and Natural Sciences, Institute of Mathematics and Computing Science, Intelligent Systems
Prof.dr. N. Petkov, Dr. M.H.F. Wilkinson

Research 2007-2008

Research area: interrelated topics from image processing and analysis, computer vision, pattern recognition, machine learning and brain-like computing.

The senior group members have complementary areas of expertise and collaborate with each other. The group leader, Petkov, focuses on biologically motivated image processing and pattern recognition (brain-like computing) and, more recently, on arts and health care applications. Biehl works on machine learning, including the theory of learning processes, the development of efficient algorithms and their application to practical problems. Wilkinson focuses on mathematical morphology, in particular theory and algorithmics of connected filters and multi-scale analysis, segmentation and 3-D volume data analysis.

Biologically motivated Image processing and computer vision, brain-like computing
Models of the visual cortex are developed and used in computer algorithms. This research is relevant for the areas of image processing, computer vision, pattern recognition, visual perception, and computational neuroscience. Our goal is to understand how humans see and deploy principles of natural vision in artificial vision systems. Using facts from neuroscience and visual perception, we build models of visual neurons and use them in computer simulations to obtain insights and derive practical computer vision algorithms. We focus on modeling the processing of edges, contours, texture and shape. A new direction in our work is the development of image processing operators that add artistic effects to photographic images.

Connected filters and morphological operators
Research on connected operators in our group entails algorithm development (including parallelization), development of new classes of filters, applications to 2-D and 3-D medical images, and the development of new connectivity measures for these

filters for increased robustness. One line of this research links to visual cortex modeling: developing morphological analogues of texture operators based on models of certain visual cortical cells. It is hoped these morphological counterparts will be an order of magnitude faster, whilst retaining the useful properties of the cortical cell models. Finally, fast visualization based on connected attribute filters is being explored.
Segmentation is a core problem in image analysis, and methods based on both simple thresholding methods and more advanced methods such as watersheds and deformable models are being explored. Application areas are many, but the focus lies on biomedical imaging, both macroscopic (MRI, CT) and microscopic.

Machine learning
Currently we are focusing on prototype based learning schemes. The technique of Relevance Learning plays a most important role in our attempts to design novel, efficient algorithms. Relevance Learning performs LVQ training and feature selection or weighting, simultaneously. Formally, the distance measure employed in defining the classification scheme is parameterized and adapted to the training data in an iterative prescription. We are currently extending this basic method in order to account for correlations of features. The aim of these matrix-based relevance algorithms is two-fold: to achieve better classification performance and to obtain deeper insight into the structure of the data.

Applications in life sciences and health care
On the applications side, we collaborate with researchers from the medical and life sciences on various applications (see below).

Future Research Plans

Various applications in other disciplines, such as life sciences, medicine, and astronomy, will give an additional inspiration for our work. Notably, life-science, biomedical and health-care applications form an important focus for the future research of the group and a number of long-term collaborations have been started. A joint project with the department of Dermatology targets the automatic analysis of dermatoscopic images; it makes use of the image database of that department which comprises 75000 images and grows with 15000 images per year. In another joint project with the University Medical Center Groningen we will contribute to the analysis of the medical data collected from a large group of (165000) people within the LifeLines study. In a collaboration with the Groningen Bioinformatics Centre, machine learning techniques are applied to the analysis of tiling microarray data. In a joint project with researchers from the Medical School of the University of Birmingham, machine learning is applied in the context of tumour classification and a novel practical tool for the diagnosis and monitoring of adrenal tumours is being developed (patent pending). Further collaborations have been started in the areas of neurology (with the SVCG group and the department of neurology) and psychiatry (with the DSSE group and the department of psychiatry). We collaborate with the Ophthalmology department of Columbia University on the orientation analysis of retinal nerve fibres.

Key Publications 2007-2008

M. Biehl, A. Ghosh, and B. Hammer: Dynamics and generalization ability of LVQ algorithms, J. Mach. Learning Res., 8, 2007, 323-360.

G. Papari, P. Campisi, N. Petkov and A. Neri: A biologically motivated multiresolution approach to contour detection. EURASIP Journal on Advances in Signal Processing, 2007, Article ID 71828, 2007, 28 pages.

G. Papari, N. Petkov, P. Campisi: Artistic Edge and Corner Preserving Smoothing. IEEE Trans. Image Proc., 16 (10), 2007, 2449-2462.

N. Petkov and E. Subramanian: Motion detection, noise reduction, texture suppression and contour enhancement by spatiotemporal Gabor filters with surround inhibition,
Biological Cybernetics, 97 (5-6), 2007, 423-439.

G.K. Ouzounis and M.H.F. Wilkinson: Mask-Based Second Generation Connectivity and Attribute Filters, IEEE Trans. Pattern Anal. Mach. Intell., 29, 2007, 990-1004.

E.R. Urbach, J.B.T.M. Roerdink, and M.H.F. Wilkinson: Connected Shape-Size Pattern Spectra For Rotation and Scale-invariant Classification Of Gray-scale Images, IEEE Trans. Pattern Anal. Mach. Intell., 29, 2007, 272-285.
M.A. Westenberg, J.B.T.M. Roerdink, and M.H.F. Wilkinson: Volumetric attribute filtering and interactive visualization using the Max-Tree representation, IEEE Trans. Image Proc., 16, 2007, 2942-2952.

M.H.F. Wilkinson: Attribute-Space Connectivity and Connected Filters, Image Vis. Comput., 25, 2007, 426-435, special ISMM-2005 issue.

G. Papari and N. Petkov. Adaptive Pseudo-Dilation for Gestalt Edge Grouping for Contour Detection. IEEE Trans. Image Proc. 17 (2008), 1950-1962.

N. Petkov, E. Alegre, M. Biehl, L. Sanchez: Automatic classification of the acrosome status of boar spermatozoa using digital image processing and LVQ, Computers in Biology and Medicine, 38, 2008, 461-468.

E.R. Urbach and M.H.F. Wilkinson. Efficient 2-D gray-scale morphological transformations with arbitrary flat structuring elements. IEEE Trans. Image Proc. 17 (2008) 1-8.

M.H.F. Wilkinson, H. Gao, W.H. Hesselink, J.E. Jonker and A. Meijster. Concurrent computation of attribute filters on shared memory parallel machines. IEEE Trans. Pattern Anal. Mach. Intell. 30 (2008) 1800-1813.

A.Witoelar, M. Biehl, A. Ghosh, B. Hammer: Learning Dynamics of Neural Gas and Vector Quantization, Neurocomputing, 71, 2008, 1210-1219.

3.15 Contribution of TUE-ET

Technische Universiteit Eindhoven, Faculty of Electrical Engineering, Design Methodology for Electronic Systems
Prof.dr.ir. R.H.J.M.Otten, Dr.ir. T. Basten, Prof.dr. H. Corporaal, Dr.ir. M.C.W. Geilen, Prof.dr.ir. G. de Haan, Dr.ir. J.P.M. Voeten

Research 2007-2008

Our work has focused on the research themes of embedded multi-media systems, architectures, models, programming and design and on video processing algorithms and architectures. On the former, the focus of our work is on mapping applications onto multiprocessor systems. System level design methods allow the systematic development of abstract executable models. These models are used to verify correctness and performance properties of the system, or to synthesise systems or system components. The results allow system designers to take well-founded design decisions about the architecture of the system, hardware/software partitioning, the choice of processors, etc. Multiprocessor implementation platforms require novel programming techniques, mapping, and synthesis techniques. These techniques must optimize multiple objectives, execution time, memory usage, and energy usage and allow trade-offs. Platform development is necessary to provide hardware and software components with predictable behaviour.
We have completed an automated mapping trajectory for multi-media applications targeting multiprocessor architectures. As a basic model of computation, we take dataflow models such as synchronous dataflow (SDF) graphs. We have developed new analysis techniques for SDF graphs, and have integrated them in the publicly available tool set SDF3 (http://www.es.ele.tue.nl/sdf3). We have also considered extensions of the dataflow models with new features, particularly to capture the concept of application scenarios that capture common execution modes of a system implementation from the resource usage point of view. Analysis techniques for this model, Scenario-Aware Dataflow Graphs, have been developed and implemented. We have also worked on techniques to exploit scenario characterizations at run-time.
Our performance modeling and analysis techniques have been extended with exhaustive techniques to compute performance numbers. In addition techniques to automatically generate performance models based on abstract application and platform specification have been developed. An exhaustive technique to compute timing errors of the property-preserving synthesis approach has been developed. The work on our hybrid VLIW/SIMD processor template has been completed.
The focus of our video processing research theme is targeted on high-end consumer applications. The current activities fall into two categories, format conversion and video enhancement. The format conversion activities have focused on performance enhancement of the motion estimation unit. (Patent) literature research was carried out to determine state of the art motion compensation/estimation knowledge in industry and academia. Selected subcomponents of algorithms found in literature were tested; of particular interest were (I) candidate generation and selection methods, (II) motion estimator (discontinuity preserving) smoothness constraints and related convergence performance, and (III) occlusion detection and correction methods. A new hierarchical motion estimator has been designed where the candidates from the first hierarchical level are generated in the radon transform domain (topic I). We plan to put together the results of our format conversion activities in a software package.
The focus of the video enhancement activities shifts from resolution up-conversion, necessary since modern displays have more pixels than the video signal contains, to the enhancement of compressed video. A new classification method for image processing using Gaussian mixture models has been developed. Furthermore, a new blur estimation algorithm has been designed, patented and published.

Future Research Plans

On the embedded multimedia systems theme, our major objective is to design platforms and systems, in a predictable way, such that we can guarantee non-functional requirements, while attempting to minimize resource and energy usage and being able to dynamically match quality with available resources. We focus on MPSoC platforms, high-tech industrial systems and wireless sensor networks. The research will be centred around three tracks: (1) The definition and modelling of a predictable MPSoC architecture, its implementation in the form of a software-based simulation environment, and its implementation on an FPGA platform. We will focus on TDMA as well as round-robin type of arbitration mechanisms to achieve predictability and composability. (2) Methods and models for predictable design. Techniques to efficiently and accurately analyze the resource usage, latency, throughput, energy usage, etc., of systems, and to synthesize systems with guarantees are the major topics. The challenge is how to cope with variability and dynamism. We are pursuing a scenario based approach to deal with this. We want to extend single-objective optimisation techniques towards multi-objective optimisation and run-time exploitation of trade-offs for QoS management. (3) The development of a design-flow and run-time environment for sensor-network platforms that allows us to program and configure such a platform in such a way that desired quality levels are guaranteed, and can be maintained during operation. The major challenges are the facts that the optimization space has many dimensions (network usage, processing, energy, reliability, timeliness, security, etc.) and is highly dynamic.

On the video processing theme, in the past, the focus in the field was on sharpness, contrast, and colour improvement, and on noise reduction. Typically, noise entered the system during analogue transmission and/or storage leading to an additive Gaussian type of noise with either a flat or a triangular noise spectrum. More recently, we see a transition to digital transmission and storage leading to different types of signal degradation, like blockiness and mosquito noise. Although the older noise reduction techniques can have a positive effect even though they have not been designed to reduce these newer artifacts, a dedicated design shows more effect. Several concepts have been proposed, both in the scientific literature and in patents, while it is often unclear what is being used in consumer electronics products on the market. Our work will focus on techniques for coding artifact reduction that are optimally adapted to these changing circumstances.
The state of the art picture quality enhancement techniques furthermore rely heavily on state-of-the art motion estimation algorithms that critically determine the performance of a growing list of features, like Natural Motion, MC de-interlacing, motion-blur reduction of LCD-panels, noise filtering and video compression. In the current dynamic consumer-electronics market, platform technology is changing rapidly, from hardware platforms to programmable platforms. Furthermore, a large variety of applications is being proposed and developed. Finally, technological advances, as reflected by Moore's law, relax the implementations constraints for motion estimation algorithms. Our work will therefore focus on the development of new motion estimation techniques for modern platforms and applications.

Key Publications 2007-2008

Zhao, M., Bosma, M. and Haan, G. de , Making the best of legacy video on modern displays. Journal of the Society for Information Display, 15(1), 49-60, 2007.

Beric, A., Meerbergen, J. van, Haan, G. de and Sethuraman, R. Memory-centric video processing. IEEE Transactions on Circuits and Systems for Video Technology, 18(4), 439-452, 2008.

Geilen, M.C.W., Basten, T., Theelen, B.D. and Otten, R.H.J.M. An algebra of Pareto points. Fundamenta Informaticae, 78(1), 35-74, 2007.

Gheorghita, Valentin, Basten, T. and Corporaal, H. Application scenarios in streaming-oriented embedded-system design. IEEE Design and Test of Computers, 25(6), 581-589, 2008.

Kumar, Akash, Fernando, S., Ha, Y., Mesman, B. and Corporaal, H. Multiprocessor systems synthesis for multiple use-cases of multiple applications on FPGA. ACM Transactions on Design Automation of Electronic Systems, 13(3), 1-27, 2008.

Akesson, K.B., Steffens, E.F.M., Strooisma, E. and Goossens, K.G.W. Real-time scheduling using credit-controlled static-priority arbitration. In 14th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications : RTCSA '08 : 25 - 27 Aug. 2008, Kaohsiung, Taiwan (pp. 3-14). Piscataway: IEEE, 2008.

Huang, J., Voeten, J. and Corporaal, H. Predictable real-time software synthesis. Real-Time Systems, 36(3), 159-198, 2007.

Hoes, R.J.H., Basten, T., Tham, C.K., Geilen, M.C.W. and Corporaal, H. Analysing Qos Trade-offs in Wireless Sensor Networks. In Proceedings of the 10 th ACM International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM 2007) 22-26 October 2007, Chania, Crete Island, Greece (pp. 60-69). New York, New York: ACM Press, 2007.

Stuijk, S. Predictable Mapping of Streaming Applications on Multiprocessors. Ph.D. dissertation, Eindhoven University of Technology.

Fatemi, H. Processor architecture design for smart cameras. Ph.D. dissertation, Eindhoven University of Technology.

3.16 Contribution of TUE-WI

Technische Universiteit Eindhoven, Department of Mathematics and Computer Science, Visualization Group
Prof.dr.ir. J.J. van Wijk, Prof.dr.ir. R. van Liere, Dr.ir. H.M.M. van de Wetering, Dr. M.A. Westenberg, Dr. A.C. Jalba

Research 2007-2008

The aim of visualization is to develop methods and techniques using interactive computer graphics such that most insight in large data sets can be obtained. The TU/e Visualization group is active in the following areas:
Information Visualization. We study how large amounts of abstract data, such as tables, trees, networks and combinations of these, can be visualized. We focus especially on Software Visualization, which is a challenging and attractive field. This has led to a variety of new methods and techniques, described in a large number of articles and several PhD theses, focussing on software evolution (PhD Voinea, 2007), state spaces (PhD Pretorius, 2008), and compound graphs (PhD Holten, 2009).
Scientific Visualization. Scientific visualization concerns data from simulations and measurements, defined over geometric spaces. Within this area we study feature extraction, flow visualization, mathematical visualization, and, in cooperation with the TU/e Department of Biomedical Engineering, medical visualization. An important focus has been feature based visualization. For very large data sets representation of the data at a higher semantic level is beneficial, with respect to ease of understanding by the user as well as faster processing. This has led to two PhD theses, on skeletonization of binary volumes (PhD Reniers, 2009) and on feature based visualization of mass spectroscopy data (PhD Broersen, 2009).
3D interaction. In cooperation with CWI we study how affordable desktop Virtual Reality systems (hard- and software) can be designed to simplify interaction with 3D data and objects for interrogation and navigation. This has led to new methods and techniques for the interactive measurements of 3D objects (PhD Kruszynski, 2010), and on a programmable display architecture for VR applications (PhD Smit, 2009).

Future Research Plans

In the next period we pay more attention to more fundamental issues, including methodological issues (how to develop systems and how to evaluate them) and a better foundation on supporting disciplines, including for instance statistics, HCI, perception, and cognitive psychology. In this respect, we intensified our cooperation with prof. J.-B. Martens, Dept. of Industrial Design, who is an expert in perception, usability studies, and statistical analysis, and we increasingly publish our work at ACM CHI.
Another important trend is Visual Analytics: The science of analytical reasoning, supported by interactive visual interfaces. Integration of other data analysis methodologies (statistics, data mining); heterogenous data; and consideration of the complete data analysis process, from collection to presentation, are key aspects. We participate in the EU VisMaster project, which aims at setting up a European research agenda for Visual Analytics; Shrinivasan focuses in his PhD project on support for navigation and knowledge capture in visualization. In the POSEIDON project we cooperate with Thales and prof. E. Postma, University of Tilburg, on visual analysis of vessel movement data; we have started cooperation with dr. M. Worring, University of Amsterdam, on interactive multimedia analysis; we have started cooperation with prof. W. van der Aalst on interactive process mining. We have cooperated with PRI, Wageningen, on visualization for bio-informatics; with the advent of dr. M. Westenberg to our group, this topic has obtained a strong boost.
Furthermore, with the advent of dr. A. Jalba, we have brought expertise on numerical methods for image analysis and visualization within our group, which strengthens our interest in scientific visualization. Cooperation with the groups of prof. L. Florack (TU/e Mathematics) and dr. A. Vilanova (Department of Biomedical Engineering) has started up.

Key Publications 2007-2008

D.H.R. Holten, J.J. van Wijk (2008). Visual comparison of hierarchically organized data. Computer Graphics Forum, 27(3), 759-766.

A.J. Pretorius, J.J. van Wijk (2007). Bridging the semantic gap : visualizing transition graphs with user-defined diagrams. IEEE Computer Graphics and Applications, 27(5), 58-66.

A.J. Pretorius, J.J. van Wijk (2008). Visual inspection of multivariate graphs. Computer Graphics Forum, 27(3), 967-974.
S.L. Voinea, A.C. Telea (2007). Visual data mining and analysis of software repositories. Computers and Graphics, 31(3), 410-428.

S.L. Voinea, A.C. Telea (2006). CVSgrab : mining the history of large software projects. In EuroVis 2006 (Proceedings Joint Eurographics - IEEE VGTC Symposium on Visualization, Lisbon, Portugal, May 8-10, 2006) (pp. 187-194). Aire-la-Ville, Switzerland: Eurographics Association.

D. Reniers, J.J. van Wijk A.C. Telea (2008). Computing multiscale curve and surface skeletons of genus 0 shapes using a global importance measure. IEEE Transactions on Visualization and Computer Graphics, 14(2), 355-368.

Y.B. Shrinivasan, J.J. van Wijk (2008). Supporting the analytical reasoning process in information visualization. In M. Czerwinski, A.M. Lund, D.S. Tan (Eds.), Proceedings of the 2008 Conference on Human Factors in Computing Systems (CHI 2008, Florence, Italy, April 5-10, 2008) (pp. 1237-1246).

A. Broersen, R. van Liere, R.M.A. Heeren (2007). Parametric visualization of high resolution correlated multi-spectral features using PCA. In K. Museth, T. Möller, A. Ynnerman (Eds.), EuroVis07 (Proceedings Joint Eurographics - IEEE VGTC Symposium on Visualization, Norrköping, Sweden, May 23-25, 2007).

F.A. Smit, A.J. van Rhijn, R. van Liere (2007). GraphTracker : a topology projection invariant optical tracker. Computers and Graphics, 31(1), 26-38.

F.A. Smit, R. van Liere, B. Fröhlich (2007). Non-uniform crosstalk reduction for dynamic scenes. In Proceedings IEEE Virtual Reality Conference 2007 (VR'07, Charlotte NC, USA, March 10-14, 2007).

3.17 Contribution of TUE-BMT

Technische Universiteit Eindhoven, Department of Biomedical Engineering Biomedical Image Analysis
Prof.dr.ir. B.M. ter Haar Romeny, Prof.dr. F.A. Gerritsen, Prof.dr. L.M.J. Florack, Dr. A. Vilanova Bartroli, Dr.ir. H.C. van Assen, Dr.ir. R. Duits, Dr.ir. B. Platel

Research 2007-2008

The group Biomedical Image Analysis started in 2001, and now comprises 6 staff members, 12 PhD students, 3 postdocs and around 20 MSc students. The group focuses on generic mathematical approaches to solve image analysis problems in cardiovascular and neurological applications, and advanced visualization. The group collaborates with the TU/e Magnetic Resonance Lab, Philips Healthcare, the FC Donders Institute, the University of Zürich, the Maastricht, Utrecht, Nijmegen and Leiden University Hospitals and the Epilepsy Center Kempenhaeghe. Education is an important aspect, being housed in the largest BME Dept. in the Netherlands, with 500 students. A full range of courses is given, from 1st – 5th year, and at PhD level (ASCI course a8).

A multi-scale framework has been established for doing high-order differential geometry on high-dimensional images, with applications as adaptive 'geometry-driven' edge preserving enhancement, multi-scale optic flow extraction, and deep structure analysis for content-based image retrieval (VICI Florack 2005). This visual perception-inspired framework is expanded to multi-orientation analysis in 2D and 3D, giving rise to powerful contextual operators. In collaboration with Philips Healthcare (Best) work on computer aided diagnosis focuses on dynamic contrast enhanced MRI of breast tumors, low-dose catheter tracking, and cardiovascular dynamics to study local dense ventricular optic flow and deformation for non-invasive ventricular infarct size estimation. The challenge is taken to segment the thin atrial wall, to optimize cardiac ablation procedures.
Much effort has been given to the GPU-based visualization and analysis of tensor fields of DTI (Diffusion Tensor Imaging) data (VENI Vilanova 2006), and High Angular Resolution Diffusion Imaging (HARDI), as well as 4D flow data. The clinical branch at the Maastricht University Hospital focuses on precise navigation for Deep Brain Stimulation, in close collaboration with the neurosurgery dept.

The mathematical analysis of (higher order) tensor fields, either for DTI, HARDI or strain tensor fields in cardiac deformation, includes methods from Finsler geometry, numerical methods for geodesic ray tracing, and a new Lie-group based theory on 3D orientation 'scores'. A new start-up company (InViso) is established to implement these methods in massive FPGA based hardware. This project is rewarded a 200 K€ STW valorization grant.
To inspect and interact with these complex tensor fields a sophisticated and flexible GPU-based visualization tool is developed, 'DTItool', which enables the interactive manipulation of all relevant parameters, 3D orientation glyphs, and tractography. The tool is used in many collaborating labs. The neurosurgical navigation for DBS towards the difficult to locate subthalamic nucleus (and sub-parts of it), which has to be stimulated in Parkinson patients to cure tremor attacks, is assisted by contextual tractography, HARDI glyph classification, and optimal stereotactic path planning.

Future Research Plans

The current and near-future focus of research is on:

- cardiovascular applications, primarily with X-Ray and MRI, exploiting multi-valued images for local strain analysis of ventricular deformation from tagged MRI sequences, and 4D flow visualization;
- neuro applications, primarily focusing on DTI and HARDI exploited in tractography for brain connectivity and muscle fibre orientation analysis.
- GPU-based visualization of multi-valued information. The hardly investigated visualization of uncertainties is an important research topic. Applications include the surgical preparation for safe epilepsy surgery with fiber visualization to spare the optic radiation, and brain tractography in relation to other brain imaging modalities (fMRI, EEG). The many parameters of cardiac function are integrated into a clinically more effective comprehensive visualization.
- generic and advanced mathematical methods for multi-valued image analysis, for segmentation, enhancement, morphological and contextual operations, and tractography of crossing and splitting fibers.

Key Publications 2007-2008

H.E. Bennink, H.C. van Assen, G.J. Streekstra, R. ter Wee, J.A.E. Spaan, B.M. ter Haar Romenij, A Novel 3D Multi-Scale Lineness Filter for Vessel Detection, Lecture Notes in Computer Science, 4792, 436-443, (2007).

E.J.L. Brunenberg, A. Vilanova Bartroli, V. Visser-Vandewalle, Y. Temel, L. Ackermans, B. Platel, B.M. ter Haar Romenij, Automatic Trajectory Planning for Deep Brain Stimulation: A Feasibility Study, Lecture Notes in Computer Science, 4791, 584-592, (2007).

S. Busking, A. Vilanova Bartroli, J.J. van Wijk, Particle-based non-photorealistic volume visualization., The Visual Computer, 24(5), 335-346, (2007).

R. Duits, M. Duits, M.A. van Almsick, B.M. ter Haar Romenij, Invertible orientation scores as an application of generalized wavelet theory , Pattern Recognition and Image Analysis, 17(1), 42-75, (2007).

E.M. Franken, R. Duits, B.M. ter Haar Romenij, Nonlinear diffusion on the 2D Euclidean Motion group, Lecture Notes in Computer Science, 4485, 461-472, (2007).

V. Prckovska, A.F. Roebroeck, W.L.P.M. Pullens, A. Vilanova Bartroli, B.M. ter Haar Romeny, Optimal Acquisition Schemes in High Angular Diffusion Imaging, Lecture Notes in Computer Science, 5242(2), 9-17, (2008).

G. Schoonenberg, P. Lelong, R. Florent, O. Wink, B.M. ter Haar Romenij, The Effect of Automated Marker Detection on In Vivo Volumetric Stent Reconstruction, Lecture Notes in Computer Science, 5242, 87-94, (2008).

G.J. Strijkers, A. Bouts, W. M Blankesteijn, T.H.J.M. Peeters, A. Vilanova Bartroli, M.C. van Prooijen, H.M. Sanders, E. Heijman, K. Nicolay, Diffusion tensor imaging of left ventricular remodeling in response to myocardial infarction in the mouse., NMR Biomed, 22(2), 182-190, (2008).

R. Duits, M.A. van Almsick, The explicit solutions of linear left-invariant second order stochastic evolution equations on the 2D Euclidean motion group, Quarterly of Applied Mathematics, 66(1), 27-67, (2008).

M. Termeer, J. Olivan Bescos, M. Breeuwer, A. Vilanova Bartroli, F.A. Gerritsen, E. Gröller, E. Nagel, Visualization of Myocardial perfusion derived from coronary anatomy, IEEE, Transaction on Visualization and Computer Graphics, 14(6), 1595--1602, (2008).

3.18 Contribution of UL-LUMC

Leiden University Medical Center, division of Image Processing, laboratorium voor klinische en Experimentele Beeldverwerking
Prof.dr.ir. Johan H.C. Reiber, Dr.ir. B.P.F. Lelieveldt, Prof.dr. R. Nelissen

Research 2007-2008

The main goal of the Division of Image Processing is the research, implementation and validation of image processing approaches, which allow the objective and reproducible assessment of objects in medical images. LKEB activities belong to the main research fields of the LUMC under the headings “Vascular Medicine”, “Neuro-science” and “Molecular Imaging”.

Part of the research involves computer vision research and algorithm development, whereas clinical applications play a very important role. Applications focus on Neuro-imaging, Pulmonology, Orthopaedics, cardiovascular, as well as molecular and cellular imaging. In 2007 and 2008, important research directions were:

Statistical shape modeling
Statistical shape models are widely used to integrate a-priori knowledge about shape and image appearance into segmentation algorithms. Research at LKEB is directed towards dimensional extension of statistical shape models. A 3D Active Shape Model has been developed, along with 3D and multi-view Active Appearance Models; these models have been applied to segmentation of cardiac MR, CT, echo and X-angiographic data. Apart from segmentation, we are developing statistical shape models for computer-aided diagnosis to detect cardiac shape- and motion abnormalities in MR images for patients with a cardiac infarction, and for quantifying local shape changes caused by brain diseases. We recently investigated shape differences of the brain ventricles in the aging population and Alzheimer disease. Because of the challenging shape of the ventricles, we developed a new method based on growing neural networks to fully automatically model complexly shaped objects. We also developed algorithms based on blind source separation methods to correct for motion artifacts in perfusion imaging.
Multi-agent image processing
The major objective of this research is to develop a general and adaptive learning multi-agent image interpretation system, which automatically learns how to interpret (medical) images from examples and user-interactions. The system should be flexible and easy to adapt to changes in patient context, expert preferences, or imaging devices, by the use of both low-level training / optimization and high-level rules.
The system will be applied to very difficult segmentation problems in images that cannot be solved with only traditional and/or model-based segmentation methods alone. To this end we are investigating how probabilistic models, reinforcement learning techniques, evolutionary algorithms, high-level (explicit) knowledge and low-level image processing may be integrated into our current multi-agent image interpretation system. The system will first be used for the interpretation of IntraVascular UltraSound (IVUS) and Computed Tomography Angiography (CTA) images.

Molecular image integration
In this project, we are investigating novel algorithms to combine complementary information in molecular, structural and functional imaging. We address novel image processing challenges brought on by new molecular imaging modalities such as bioluminenscence imaging and fluorescence imaging. We focus on whole body registration between optical and structural data in follow-up studies, detection of changes and abnormalities and on integrating information sources over the scale range from molecule to organism.

Clinical Image Analysis Applications
Much of the research at LKEB is driven by questions from clinical partners. To this end, we are developing algorithms and software for:

- Detection and quantification of pulmonary emphysema in CT Images
- Early detection of micro motion of prosthetic implants in stereo X-ray images
- Automatic analysis of coronary vessels in CT and intravascular ultrasound images
- Automatic analysis of coronary and left-ventricular angiograms
- Automatic analysis of cardiac function, perfusion, infarct imaging in MR and CT patient studies
- Automatic analysis of changes in brain structure with ageing and disease
- Automatic analysis of vascular MR data
- Automatic white matter lesion detection in MR images of the brain

Key Publications 2007-2008

J. Milles, R.J. van der Geest, M. Jerosch Herold, J.H.C. Reiber, B.P.F. Lelieveldt, “Fully Automated Motion Correction in First-Pass Myocardial Perfusion MR Image Sequences”, *IEEE Transactions on Medical Imaging*, vol. 27(11), pp 1611-1621, 2008.

H.C. van Assen, M.G. Danilouchkine, M.S. Dirksen, J.H.C. Reiber, B.P.F. Lelieveldt, “A 3D Active Shape Model driven by Fuzzy Inference: Application to Cardiac CT and MRI”, *IEEE Transactions on Information Technology in Biomedicine*, vol. 12(5), pp 595-695, 2008.

E. Angelie, C.R. Oost, D. Hendriksen, B.P.F. Lelieveldt, R.J. van der Geest, J.H.C. Reiber, “Automatic Contour Detection using Active Appearance Models: the Effect of the Composition of the Training Set”, *Investigative Radiology* 42(10), pp 697-703, 2007.

Ferrarini L, Olofsen H, Palm WM, Buchem MA van, Reiber JHC, Admiraal-Behloul F. GAMEs: growing and adaptave meshes for fully automatic shape modeling and analysis. Medical Image Analysis 2007; 11: 302-14.

Stralen M van, Leung KY, Voormolen MM, Jong N de, Steen AF van der, Reiber JH, Bosch JG Time continuous detection of the left ventricular long axis and the mitral valve plane in 3-D echocardiography. Ultrasound in Medicine and Biology, vol 34(2), pp 1389-1399, 2008.

M.Baiker, J. Dijkstra, I. Que, C.W.G.M. Lowik, J.H.C. Reiber, B.P.F. Lelieveldt, "Organ approximation in micro-CT data with low soft tissue contrast using an articulated whole-body atlas", paper nr 1191, proc. ISBI 2008, pp 1267-1270, 2008.

3.19 Contribution of UT-EWI-DACS

University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science, Design and Analysis of Communication Systems
Prof.dr. B. Haverkort, Dr.ir. G.J. Heijenk, Dr.ir. A. Pras, Dr.ir. P.T. de Boer, Dr.ir. G. Karagiannis

Research 2007-2008

DACS Mission

DACS focuses on the design and analysis of *dependable networked systems.* A system is called dependable, whenever reliance can justifiably be placed on the services it delivers. Tailored to communication systems, which can be wired, wireless, or embedded in other systems, this means that we aim:

> "to contribute to the design and implementation of dependable networked systems, as well as to methods and techniques to support the design and dimensioning of such systems, such that they are dependable, in all phases of their lifecycle."

We thereby interpret the term dependability as encompassing availability, reliability, performance (quality of service) and security.

DACS Strategy

Three phases in system design. We distinguish three phases in the design of dependable networked systems. In the first (exploratory) phase the emphasis lies on the development of system models and the model-based evaluation of key system characteristics. In the second, more concrete phase the emphasis lies on designing and/or standardizing systems architectures, protocols, and algorithms, as well as on implementing prototype systems. In the third, operational phase the systems have been implemented and the emphasis lies on operationally managing them. As an example of these three phases, consider the Internet backbone. During 1965-1985, research concentrated on a variety of networking models (circuit-switching vs. packet-switching, various local area access mechanisms). In the period 1980-1995, research focus shifted toward designing, implementing and standardizing protocols (TCP/IP, Ethernet). Since the beginning of the 1990s, an ever more important challenge is to keep the Internet up and running (dependability), and solve problems related to (performance) management, security and scalability. A similar transition in focus can be seen for wireless communication system. The figure below illustrates, through a horizontal structuring, the above three phases in system design.

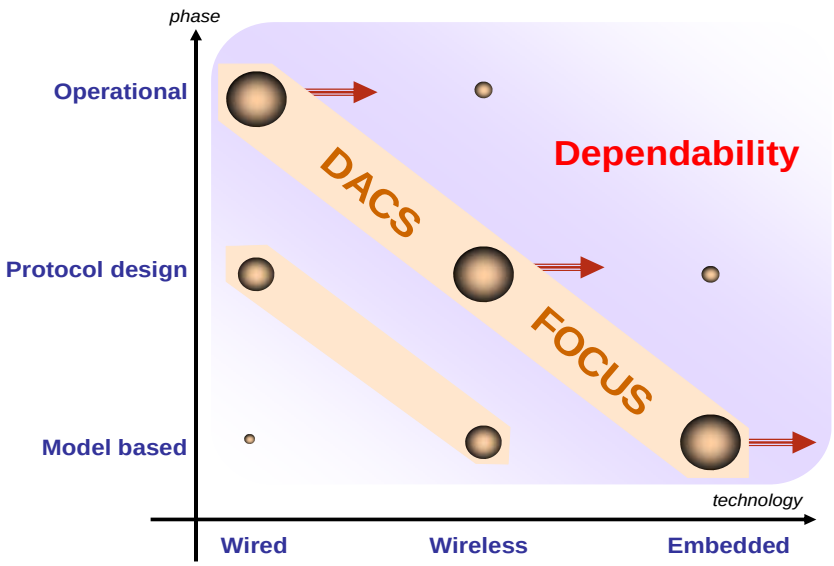

Three technologies for dependable networked systems. Research within DACS covers the whole spectrum of network technologies: from well-established technologies (like the wired Internet), via technologies that are under development (such as wireless networks) to emerging technologies (like embedded network systems).

In the case of well-established technologies, research concentrates on operational aspects, here, in particular, of the wired Internet. Specific topics include bandwidth allocation, accounting, self-management of lambda switches and protection against scans, denial-of-service attacks and phishing. Taking and interpreting measurements plays an important role in this research.

For technologies under development, research focuses on the design, evaluation, and prototype implementation of new protocols and algorithms for wireless and ad-hoc networks. Topics include algorithms for context- and power-aware routing in ad-hoc networks, and, lately, more and more on car-to-car communications and wireless sensor networks. These types of networking systems are exponents of embedded networking technologies as well.

The remaining research on embedded networking technologies focuses on system specification and evaluation techniques to describe such systems, and the resource constraints (performance, dependability, energy usage) they have to operate under. This includes the development of new stochastic model checking techniques and the application thereof to predict dependability and performance properties.

The figure above also shows these three technology domains (where a clear cut separation can, of course, never be made). When technologies mature over time, a shift in phase and type of activity is foreseen. Inherent to university research, DACS moves on the increasing side of the wave (to the right). Note that the presented structure is based on technical content, not an organizational sub-structuring of the group. There is strong collaboration between the various group members, and projects generally cover multiple issues.

Key Publications 2007-2008

Baier, C., Cloth, L., Haverkort, B.R.H.M., Kuntz, G.W.M., Siegle, M. Model Checking Markov Chains with Actions and State Labels. IEEE transactions on software engineering nr: 4, 33 (nr: 1), (2007, January 1). (pp. 209-224) Los Alamitos IEEE Computer Society Press ISSN: 0098-5589.

Boer, P.T. de, Scheinhardt, W.R.W. Alternative proof and interpretations for a recent state-dependent importance sampling scheme. Queueing systems nr: 2-3, 57 (nr: 67310A), (2007, January 1). (pp. 61-69) Dordrecht Springer ISSN: 0257-0130.

Foreest, N.D. van , Haverkort, B.R.H.M., Mandjes, M.R.H., Scheinhardt, W.R.W. Versatile Stochastic Models for Networks with Asymmetric TCP Sources. Performance evaluation 64 (nr: 1), (2007, January 1). (pp. 507-523) Amsterdam Elsevier ISSN: 0166-5316.

Meijerink, A., Taniman, R.O., Etten, W.C. van Coherence-Multiplexed Optical RF Feeder Networks. Journal of lightwave technology nr: 11, 25 (nr: 67310A), (2007, November 1). (pp. 3396-3406) Piscataway IEEE Computer Society Press ISSN: 0733-8724.

Pras, A., Schoenwaelder, J., Burgess, M., Festor, O., Martinez Perez, G., Stadler, R., Stiller, B. Key Research Challenges in Network Management. IEEE communications magazine nr: 10, 45 (nr: LNCS4549), (2007, October 1). (pp. 104-110) New York IEEE Communication Society ISSN: 0163-6804.

Remke, A.K.I., Haverkort, B.R.H.M., Cloth, L. CSL model checking algorithms for QBDs. Theoretical computer science nr: 1, 382 (nr: LNCS4549), (2007, August 28). (pp. 24-41) Amsterdam Elsevier ISSN: 0304-3975.

Bakshi, R, Cloth, L., Fokkink, W., Haverkort, B.R.H.M. Mean-field analysis for the evaluation of gossip protocols. ACM SIGMETRICS performance evaluation review nr: 3, 36 (nr: WP 08-02), (2008, December 1). (pp. 32-39) New York ACM ISSN: 0163-5999.

Dimitrova, D.C., Karagiannis, G., Boer, P.T. de Severe congestion handling approaches in NSIS RMD domains with bi-directional reservations. Computer communications nr: 14, 13 (nr: WoTUG-31), (2008, September 5). (pp. 3153-3162) Elsevier ISSN: 0140-3664.

Litjens, R., Berg, J.L. van den, Boucherie, R.J. Throughputs in processor sharing models for integrated stream and elastic traffic. Performance evaluation nr: 2, 65, (2008). (pp. 152-180) ISSN: 0166-5316.

Woudt, E.M., Boer, P.T. de, Ommeren, J.C.W. van Improving Adaptive Importance Sampling Simulation of Markovian Queueing Models using Non-parametric Smoothing. Simulation nr: 12, 83 (nr: 2), (2007, December 1). (pp. 811-820) SAGE Publications ISSN: 0037-5497.

3.20 Contribution of UT-EWI-CAES

University of Twente, Faculty of Electrical Engineering, Mathematics an Computer Science, Computer Architecture, Design & Test for Embedded Systems
Prof.dr.ir. G.J.M. Smit, Dr.ir. A.B.J. Kokkeler, Ir. E. Molenkamp

Research 2007-2008

The main emphasis of the group is on efficient architectures for dependable networked embedded systems.
Within this theme the chair performs research on related key areas:
1) efficient architectures for streaming applications,
2) architectures for efficient architectures,
3) wireless networked embedded systems (aka wireless sensor nodes) and
4) dependability issues of networked embedded systems.
Energy-efficiency and dependability are the main drivers of our research. Energy-efficiency is important for streaming applications found in battery powered mobile devices (e.g. smart phones and portable multimedia players), and is of crucial importance for wireless sensor networks. The effort on energy efficient architectures focuses on reconfigurable processors for streaming applications, e.g. found in battery powered mobile devices. However, within high performance embedded computing (medical imaging, radar processing), extrapolating the current trend of using general purpose processors for future systems predicts excessive power consumption and dramatically reduced reliability. MPSoC (Multi Processor Systems-on-Chip) devices for streaming applications are prime candidates for use in this application domain as well. Dependability plays an important role in sensor networks (nodes may fail or run out of energy unexpectedly). Due to the abundance of nodes distributed applications continue despite of failing nodes. In MPSoC (Multi Processor Systems-on-Chip) systems for streaming applications dependability techniques also play an important role.

Efficient Architectures for Streaming Applications
In 2007 we finalized the EU FP6 4S project. Within this project, a reconfigurable multi-processor has been developed. This chip has been produced and has shown to be functioning correctly. The chip has been integrated onto a platform which will be used to evaluate applications running on this platform. Research on techniques for analysis and synthesis of predictable multi-processor systems continued in cooperation with NXP. This resulted i.e. in new techniques for modeling run-time arbitration and fast calculation of buffer capacities for multi-core architectures. The CMOS Beamforming project was started at the end of 2006 and research questions have been formulated and modeling activities have started. Within the AAF project, efficient implementations of baseband processing algorithms on reconfigurable architectures for cognitive radio have been developed and published. A formal description of run-time mapping of tasks onto a (reconfigurable) processing platform and a first version of a run-time mapping algorithm was implemented and tested (RUM project).
In 2008 several new projects started: CRISP, EASY and NEST. The EU FP7 CRISP project (Cutting edge Reconfigurable ICs for Stream Processing) researches optimal utilization, efficient programming and dependability of reconfigurable many-cores for streaming applications. The EASY (NWO) project aims at the inclusion of non-functional requirements or mapping directives in the design of streaming applications. This is not supported by current programming paradigms. The NEST project is a nation-wide initiative (all three Technical Universities within the Netherlands, Leiden University, Thales, NXP, Philips Medical Systems and Océ) project to study reconfigurable architectures and to develop a common framework for design-time and run-time tools.
As of November 2007 the research on wireless networked embedded systems (WSN) continues in the new chair PS.

Architectures for Energy Efficiency
Besides the development of efficient architectures, the use of these architectures to increase energy efficiency in a more general sense, is a subject of research. The micro-CHP plays an important role in this research. A micro-CHP is expected to become the successor of the conventional high-efficiency boiler producing, besides heat, electricity with a comparable overall efficiency. Electricity is produced during peak load or during power failure. Because the electricity demands of multiple households show correlation, reducing the peak load of individual households will reduce the peak load of the power plant. Because of the reduced peak load this power plant can generate energy more efficiently. The research topics are: peak reduction within a single household, scheduling a fleet of micro-CHPs and islanded operation in case of power outage. This research is supported by HOMA software, Essent and E.ON, Gasterra and STW.

Efficient dependable networked embedded systems
The research in dependable systems, including mixed-signal IPs and sensors, is centered on the design and implementation of architectures for dependable networked embedded systems. The key research challenge is: to foster dependability (i.e. availability, reliability, integrity and maintainability) as embedded systems are business or safety critical in almost every human endeavour, In MPSoC (Multi Processor Systems-on-Chip) systems for streaming applications dependability techniques play an important role. One of the problems in manufacturing a MPSoC with millions of transistors using deep-submicron technologies (90 nm and below), is an increase in the probability of defects in silicon, which results in decreasing

manufacturing yield. To effectively deal with this increased defect density, we need efficient methods for fault detection, localization, and fault tolerant architectures implemented on-chip.

Key Publications 2007-2008

Wiggers, M.H. and Bekooij, M.J.G. and Smit, G.J.M. “Efficient Computation of Buffer Capacities for Cyclo--Static Dataflow Graphs” In: DAC '07: Proceedings of the 44th annual conference on Design automation, 4-8 June 2007, San Diego, California, USA. pp. 658-663. ACM Press. ISBN 978-1-59593-627-1.

Smit, G.J.M. and Kokkeler, A.B.J. and Wolkotte, P.T. and Hölzenspies, P.K.F. and van de Burgwal, M.D. and Heysters, P.M. “The Chameleon Architecture for Streaming DSP Applications”, EURASIP Journal on Embedded Systems, 2007. 78082. ISSN 1687-3955.

Kerkhoff, H.G. “Testing Microelectronic Biofluidic Systems”, IEEE Design and Test of Computers, 24 (1). pp. 72-82. ISSN 0740-7475.

Marin-Perianu, M. and Meratnia, N. and Havinga, P.J.M. and Moreira Sá de Souza, L. and Müller, J. and Spiess, P. and Haller, S. and Riedel, T. and Decker, C. and Stromberg, G., “Decentralized Enterprise Systems: A Multi-platform Wireless Sensor Networks Approach”, IEEE Wireless Communications, December 2007, 1536-1284/07.

Rauwerda, G.K. and Heysters, P.M. and Smit, G.J.M. (2008) “Towards Software Defined Radios Using Coarse-Grained Reconfigurable Hardware”, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 16 (1). pp. 3-13. ISSN 1063-8210.

Zhang, Q. and Kokkeler, A.B.J. and Smit, G.J.M. (2008) “Cognitive Radio Design on an MPSoC Reconfigurable Platform”, ACM/Springer Journal of Mobile Networks and Applications, 13 (5). pp. 424-430. ISSN 1383-469X.

3.21 Contribution of EUR-UMCR

Erasmus MC, Departments of Radiology and Medical Informatics, Biomedical Imaging Group Rotterdam
Prof.dr. W.J. Niessen, Dr. J. Veenland, Dr.Ir. H. Vrooman

Research 2007-2008

Through innovative fundamental and applied research BIGR aims at developing and validating advanced techniques for the processing and analysis of large, complex, and heterogeneous medical and biological image data sets. The research of BIGR is organized along five research themes: 1) Cardiovascular image analysis; 2) Neuro image analysis; 3) Cellular and molecular image analysis; 4) Image analysis in oncology; and 5) Image guided interventions.

1) The primary cause of cardiovascular disease is atherosclerosis. A large percentage of strokes are caused by plaque build-up in carotid arteries, which may lead to lumen narrowing and/or plaque rupture. The most important step towards carotid artery plaque quantification is the difficult task of detecting the vessel boundary, i.e. the outer vessel wall, which encloses both the vessel lumen and plaque. In order to solve the automatic outer vessel wall segmentation problem we used a combined approach of a machine learning technique and à deformable model fitting.
2) We developed and validated state-of-the-art classification tools for accurate and reproducible quantification of brain changes in healthy elderly subjects. We compared the conventional k-nearest neighbor classifier that requires manually labeled data for training, with a method that automates this training phase using atlas registration. For every voxel in the input data a decision is made to classify the voxel as cerebrospinal fluid, gray matter or white matter. Our tissue segmentation method is extended with a white matter lesion segmentation method. Several research questions studied in the Rotterdam Scan Study require segmentation of specific regions in the brain. In order to delineate brain structures, intensity models are generally not sufficient. We used manually labeled atlases to introduce the additional information necessary to solve this problem. DTI scans were incorporated in the imaging protocol. We have adopted a state of the art method for the voxel-wise analysis of micro-structural integrity measurements (Tract-Based Spatial Statistics).
3) Motion analysis of nano-scale intracellular objects, commonly studied using fluorescence microscopy imaging, requires tracking of large and time-varying numbers of spots in noisy image sequences. We developed new techniques for multiple-object tracking based on nonlinear Bayesian approaches. Since these better integrate available temporal information and application-specific prior knowledge, they can be shown to perform superiorly. Analyzing the motion and deformation of large numbers of cells in image sequences is a recurrent task in many biological studies. A new level-set based cell tracking algorithm has been developed. In order to compare the performance of this new algorithm with other algorithms of the same class and with human observers an extensive validation study has been set up.

4) To assess the malignancy of tumors within the liver and breasts, Dynamic Contrast Enhanced Magnetic Resonance Imaging (DCE-MRI) is a powerful technique. Time-intensity curves showing the contrast uptake within the tumor are used to characterize the tumor as being benign or malignant. To compensate for motion artifacts, images are registered to each other, so a perfect spatial alignment is reached. Our results on liver DCE-MRI sets show that the post processing step can reduce the average absolute value of the divergence by one to three orders in magnitude.
5) In the diagnosis trajectory leading to cardiovascular interventions, three-dimensional imaging data are frequently available. The radiologist or cardiologist usually relies on intra-operative monoplane X-ray projection imaging in interventional procedures. A 2-D/3-D registration method for the alignment of cardiac X-ray images to ECG gated CTA data of the coronary arteries has been developed. The method utilizes the graphical processing unit for rendering of digitally reconstructed radiographs from a 3-D CTA-derived coronary model. A standardized evaluation methodology and reference database for the quantitative evaluation of coronary artery centerline extraction algorithms has been developed.

Future Research Plans

In the cardiovascular field, state-of-the art imaging techniques have the potential to provide detailed information on the vessel wall, such as plaque composition, elastic wall properties, and even biochemical processes that take place in the plaque. Dynamic and perfusion imaging can provide functional information, e.g. for determining the perfusion or motion of the heart, or to study tumor activity. We will develop techniques for automatic extraction of coronary artery central lumen lines and methods for automatic segmentation of the coronary lumen.
For neuro-imaging, reproducibility and accuracy experiments will be conducted on several automatic brain tissue segmentation methods. The knowledge-based segmentation of brain structures, like hippocampus, ventricles and cerebellum will be further improved. Furthermore, next to measuring volumes of several brain structures, shape-based analysis will be further investigated. Finally, we will focus our work more and more on Diffusion Tensor Imaging (DTI) to assess the micro-structural integrity of the white matter in the brain and to construct connectivity networks in the brain.
Molecular imaging is a relatively new field of research, aiming at the visualization, characterization, and quantification of biological processes at the cellular and molecular levels. Combined with anatomical and functional imaging, this enables in-depth studies of the molecular origins of diseases, in relation to their effects at the organ and whole-organism level. These new imaging possibilities are having a significant impact on the basic life sciences as well as human healthcare, through a better understanding of disease mechanisms, the development of new biomarkers for early diagnosis, and enhanced preclinical validation of novel treatments in small-animal models as a first step towards clinical implementation. In the coming years, we will focus on a particular application of molecular imaging: the use of magnetic resonance imaging (MRI) to study the effects of stem-cell therapy in restoring myocardial tissue and function after a heart infarction.
For image-guided interventions, we will improve 3D anatomical guidance by relating pre-operative 3D imaging data or anatomical models to the sparse data acquired during the procedure. We will investigate the reconstruction of 3D shapes out of sparse projected 2D data by imposing shape constraints based on a-priori knowledge of the expected shape and/or intra-operative deformation.

Key Publications 2007-2008

E.F.S. van Velsen, W.J. Niessen, T. T. de Weert, C. de Monyé, A. van der Lugt, E. Meijering and R. Stokking, Evaluation of an Improved Technique for Lumen Path Definition and Lumen Segmentation of Atherosclerotic Vessels in CT Angiography, European Radiology, 2007.

H.A. Vrooman, C.A. Cocosco, F. van der Lijn, R. Stokking, M. A. Ikram, M. W. Vernooij, M. M. B. Breteler and W.J. Niessen, Multi-spectral brain tissue segmentation using automatically trained k-Nearest-Neighbor classification, NeuroImage, 2007.

S. Klein, M. Staring and J. P. W. Pluim, Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines, IEEE Transactions on Image Processing, 2007.

M. de Bruijne, M.T. Lund, L.B. Tanko, P.C. Pettersen and M. Nielsen, Quantitative vertebral morphometry using neighbor-conditional shape models, Medical Image Analysis, 2007.

M. A. Ikram, H. Vrooman, M. W. Vernooij, F. van der Lijn, A. Hofman, A. van der Lugt, W.J. Niessen and M. M. B. Breteler, Brain tissue volumes in the general elderly population The Rotterdam Scan Study, Neurobiology of Aging, 2008.

M. Schaap, A.M.R. Schilham, K. J. Zuiderveld, E. P. A. Vonken, M. Prokop and W.J. Niessen, Fast Noise Reduction in Computed Tomography for Improved 3D Visualization, IEEE Transactions on Medical Imaging, 2008.

Smal, E. Meijering, K. Draegestein, N. Galjart, I. Grigoriev, A. Akhmanova, M. E. van Royen, A. B. Houtsmuller and W.J. Niessen, Multiple Object Tracking in Molecular Bioimaging by Rao-Blackwellized Marginal Particle Filtering, Medical Image Analysis, 2008.

M. W. Vernooij, M. de Groot, A. van der Lugt, M. A. Ikram, G. P. Krestin, A. Hofman, W.J. Niessen and M. M. B. Breteler, White matter atrophy and lesion formation explain the loss of structural integrity of white matter in aging, NeuroImage, 2008.

F. van der Lijn, T. den Heijer, M. M. B. Breteler and W.J. Niessen, Hippocampus segmentation in MR images using atlas registration, voxel classification, and graph cuts, NeuroImage, 2008.

3.22 Contribution of RUN-UMCR

Radboud University Nijmegen Medical Centre, Radiology Department, Diagnostic Image Analysis Group (DIAG) Nijmegen
Dr.ir. N. Karssemeijer, Dr.ir. H.J. Huisman, Dr. M. Velikova

Research 2007-2008

Research has been focused on development of novel image analysis methods for screening and diagnostic radiology. Application areas were beast and prostate cancer. Development of computer aided detection (CAD) of abnormalities in mammograms to improve breast cancer screening remained a main topic of the group. New projects were initiated to develop computer aided diagnosis systems for detecting and diagnosing prostate cancer in MR imaging, breast cancer in MRI, and lung cancer detection in chest x-rays.

Mammographic image analysis is an active field of research and breast cancer screening is the first domain where computer aided detection became succesful in practice. However, computer programs still perform worse than trained expert radiologists, despite a strong increase in the size of databases and computational power to perform analysis and algorithm training. One of the key issues where computers fail is the analysis of context. Therefore, we investigated methods to improve the utilization of context in mammogram analysis. In particular we developed a system to correlate findings in different views, i.e. the two mammogram projections taken per breast, and the series of mammograms taken over time. Combination of information from different views is a problem in itself, and then potential of Bayesian Networks to deal with this aspect was investigated. Multi view analysis allowed us to significantly improve detection performance [1][2][3]. Also optimized display of temporal mammrams by registration was invesigated [4]. In another project we investigated transfer learning and image normalisation to optimize use of heterogeneous databases for training CAD system for new digital mammography units [5].

For localization and staging of prostate and breast cancer with MRI an advanced processing environment is being developed in the group. In the system parametric maps are computed to represent dynamic contrast enhancement (by pharmacokinetic parameters), diffusion weighted imaging, and MR spectroscopy. The system is used daily in the radiology department for data analysis and evaluation of prostate imaging in ongoing clinical studies. Research has been performed to automate parts of the analysis. This improves diagnostic performance and makes reading more efficient [7][8]. An annotated database is created to allow further development of the system with more intelligent pattern recognition techniques.

Key Publications 2007-2008

[1] Timp S, Varela C, Karssemeijer N. Temporal change analysis for characterization of mass lesions in mammography. IEEE Trans Med Imaging. 2007 Jul;26(7):945-53.

[2] van Engeland S, Karssemeijer N. Combining two mammographic projections in a computer aided mass detection method. Med Phys. 2007 Mar;34(3):898-905.

[3] Roelofs AA, Karssemeijer N, Wedekind N, Beck C, van Woudenberg S, Snoeren PR, Hendriks JH, Rosselli del Turco M, Bjurstam N, Junkermann H, Beijerinck D, Séradour B, Evertsz CJ. Importance of comparison of current and prior mammograms in breast cancer screening. Radiology. 2007 Jan;242(1):70-7.

[4] Snoeren PR, Karssemeijer N. Gray-scale and geometric registration of full-field digital and film-screen mammograms. Med Image Anal. 2007 Apr;11(2):146-56. Epub 2007 Jan 8.

[5] Kallenberg M, Karssemeijer N. Computer-aided detection of masses in full-field digital mammography using screen-film mammograms for training. Phys Med Biol. 2008 Dec 7;53(23):6879- 91. Epub 2008 Nov 12.

[6] Huisman H, Karssemeijer N. Chestwall segmentation in 3D breast ultrasound using a deformable volume model. Inf Process Med Imaging. 2007;20:245-56.

[7] Stoutjesdijk MJ, Veltman J, Huisman H, Karssemeijer N, Barentsz JO, Blickman JG, Boetes C. Automated analysis of contrast enhancement in breast MRI lesions using mean shift clustering for ROI selection. J Magn Reson Imaging. 2007 Sep;26(3):606-14.

[8] Vos PC, Hambrock T, Hulsbergen-van de Kaa CA, Fütterer JJ, Barentsz JO, Huisman HJ. Computerized analysis of prostate lesions in the peripheral zone using dynamic contrast enhanced MRI. Med Phys. 2008 Mar;35(3):888-99.