Advanced School for Computing and Imaging

Five Year Research Plan

2010-2014

August 2010

Advanced School for Computing and Imaging

ASCI
p/a P.O. Box 5031
2600 GA Delft
The Netherlands
phone +31 15 278 9660
fax +31 15 278 6632
email asci@ewi.tudelft.nl
www http://www.asci.tudelft.nl

Contents

ASCI in a nutshell

ASCI is a research school in the Netherlands at the level of PhD. It contains PhDs of ten universities in the Netherlands. Effectively it is one of three research schools in ICT, although it also covers some physics and medicine. 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 most valuable asset of ASCI is the education and guidance of its PhD students. The education is organized by the participating universities collectively through a set of multi-party PhD level courses, maintained and updated through the years. As the courses yield no (financial) reward to teachers, the course program is a sign of the enduring commitment of ASCI staff members to the development of their PhDs. An important aspect of the ASCI philosophy is that each of the courses will be taught by staff members from different organizations and schools of thought.

For more information about ASCI, please see our website: http://www.asci.tudelft.nl/

Arnold Smeulders
Scientific Director

ASCI and its Research

ASCI is a national research school on advanced computer and imaging systems. Research within the school can be characterized as applied, experimental and technical computer science, focused primarily on parallel and distributed systems and processing, as well as the processing of sensor data, image data and other media. With the emphasis on system development, integration of software and hardware, and the processing of sensory information, it directly addresses the needs in high performance computing and computing intensive applications, with a special emphasis on media oriented applications such as multimedia, medical imaging, computer vision, industrial automation and CAD/CAM. Other important topics are embedded systems and wide-area systems.

Participants in ASCI are groups from Delft University of Technology, Eindhoven University of Technology, University of Twente, University of Amsterdam, VU University Amsterdam, Leiden University, University of Utrecht, University of Groningen, Radboud University Nijmegen.

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.

Mission

The mission of ASCI is to train PhD students to perform high-quality research in the fields it is active in. These fields include the design, implementation, and application of advanced computer systems and the design, implementation, and application of computer imaging systems. In particular, we believe – like many other computer scientists – that parallel and distributed computer systems as well as embedded systems are the future. In addition, we observe a trend away from text-based information towards visual information of all kinds, which leads to research questions in image recognition, robotic vision, and multimedia systems. Furthermore, there is a trend for these areas to come together in advanced systems which will be both parallel and image oriented. These are all areas in which we are actively carrying out research.

The success or failure of a research school is determined by the quantity and quality of the students that are granted doctorates. With over 2000 refereed publications in renowned journals and conference proceedings, and 204 doctorates, we believe to have met this goal. We have also been successful in other ways, such as obtaining large grants and stimulating interdisciplinary work.

Within the area of computer systems, most of the research has focused on distributed and parallel computer systems. Distributed systems are those in which a large number of computers spread over a wide geographic area work together to provide communication and computing services to their users. The World Wide Web is an example of such a system, and some of our research is directed towards improving it. Parallel systems are those in which a collection of computers, possibly distributed over multiple sites, work together to solve a single large problem that is too large for any of them to solve individually. The Distributed ASCI Supercomputer, DAS, for which we received four grants from NWO, is a unique example of such a system, and one we have used for a variety of interesting research projects. An application is found in the infrastructure for compute-intensive e-science. This work has involved programming languages, models, operating systems, middleware and numerous applications, for example in computational science. A growing amount of research has been on computer architecture and embedded systems, including sensor networks.
The key properties of DAS are its physical distribution and hierarchical structure: The system consists of 4-6 cluster computers located at different universities, connected by a wide-area network. DAS-3 featured an advanced optical wide-area network provided by SURFnet, using multiple 10Gb/s dedicated optical interconnects between all sites. The new system DAS-4 uses a diversity of hardware accelerators (like GPUs) and is a platform for research on Green IT.

Within the area of imaging, we are performing research on a wide variety of topics including image sensors and acquisition, low-level image processing, feature extraction, multisensor fusion, scientific visualization and multimedia techniques. Some of the work in imaging is fundamental, but a great deal of it is applied, in which imaging techniques are applied to a variety of important problems, from understanding maps to providing the vision system to allow robots to play soccer. Advanced imaging applications put heavy demands on computers, which encourages interaction between the researchers designing embedded, distributed and parallel computer systems and the imaging researchers who need such systems. Some of the applications running on the Distributed ASCI Supercomputer involve imaging, for example. An application is found in data-intensive e-science. We have found this co-operation to be fruitful and expect it to increase in time as a result of ASCI. Before ASCI existed, such co-operation did not exist at all.

In the next five-year period we aim to continue producing top quality research in computing, in imaging, and in the area where the two intersect. Such topics as advanced distributed systems and digital sensor data processing are likely to remain important research themes and we intend to pursue them. Other important topics are (distributed) multimedia, interaction, social data mining and embedded systems. We expect that important parts of these ambitions can be fulfilled as part of the national program COMMIT for public - private research in ICT. We also plan to continue and refine the extensive educational program that ASCI has operated over the past fifteen years.

Arnold Smeulders
Scientific Director

Participating Groups

The following research groups participate in ASCI. They are represented together with their abbreviations. For each group the staff members are listed and the contact person is underlined (situation May 2010).

EUR-UMCR
Erasmus University Rotterdam, Departments of Radiology and Medical Informatics
Biomedical Imaging Group Rotterdam
http://www.eur.nl/fgg/mi/
Dr. M. de Bruijne, Dr. S. Klein, Dr. E. Meijering, Prof. Dr. W.J. Niessen, Dr. J. Veenland, Dr. H. Vrooman, Dr. Th. Van Walsum

RUG-CS-IS
University of Groningen, Faculty of Math. and Natural Sciences, Johann Bernoulli Institute
Intelligent Systems
http://www.cs.rug.nl/is/
Prof. Dr. M. Biehl, Prof. Dr. N. Petkov, Dr. M.H.F. Wilkinson

RUG-CS-SVCG
University of Groningen, Faculty of Math. and Natural Sciences, Johann Bernoulli Institute
Scientific Visualization and Computer Graphics
http://www.cs.rug.nl/svcg/
Dr. H. Bekker, Dr. T. Isenberg, Prof. Dr. J.B.T.M. Roerdink, Prof. Dr. A.C. Telea

RUN-UMCR
Radboud University Nijmegen, Department of Radiology, Diagnostic Image Analysis Group
http://www.diagnijmegen.nl/
Dr. Bram van Ginneken, Dr. Ir. HenkJan Huisman, Dr. Ir. Nico Karssemeijer, Dr. Clarisa Sanchez, Dr. Peter Snoeren

TUD-EWI-MM-CGCC
Delft University of Technology, Faculty of Electrical Eng., Math. and Computer Science
Computer Graphics
http://graphics.tudelft.nl
Dr. Ir. A.R. Bidarra, Dr. Ir. C.P. Botha, Dr. W.F. Bronsvoort, Prof. Dr. Ir. F.W. Jansen, Ir. F.H. Post

TUD-EWI-MM-ICT/1
Delft University of Technology, Faculty of Electrical Eng., Math. and Computer Science
Multimedia Signal Processing Group
http://www.ewi.tudelft.nl/
Prof. Dr. Ir. J. Biemond, Dr. A. Hanjalic, Prof. Dr. Ir. R.L. Lagendijk, Dr. M. Larson, Dr. J.C.A. van der Lubbe, Dr. P. Peris Lopez, Dr. P. Serdyukov, Dr. T. Veugen, Prof. Dr. Ir. A. de Vries

TUD-EWI-MM-ICT/2
Delft University of Technology, Faculty of Electrical Eng., Math. and Computer Science
Pattern Recognition and Bio-informatics
http://prlab.tudelft.nl, http://bioinformatics.tudelft.nl, http://visionlab.tudelft.nl
Dr. Ir. R.P.W. Duin, Dr. E.A. Hendriks, Dr. M. Loog, Prof. Dr. Ir. M.J.T. Reinders,
Dr. Ir. D. de Ridder, Ir. J. de Ridder, Dr. D.M.J. Tax

TUD-EWI-ST-PGS
Delft University of Technology, Faculty of Electrical Eng., Math. and Computer Science
Parallel and Distributed Systems
http://www.pds.ewi.tudelft.nl
Dr. S.O. Dulman, Dr. Ir. D.H.J. Epema, Dr. A. Iosup, Prof. Dr. K.G. Langendoen,
Dr. Ir. J.A. Pouwelse, Prof. Dr. Ir. H.J. Sips

TUD-L&R-FRS
Delft University of Technology, Faculty of Aerospace Engineering, Dept. of Remote Sensing
Optical and Laser Remote Sensing
http://www.lr.tudelft.nl/olrs/
Dr. Ir. B.G.H. Gorte

TUD-TNW-QI
Delft University of Technology, Faculty of Applied Sciences, Department of Imaging S & T
Quantitative Imaging
http://www.tnw.tudelft.nl/
Prof. Dr. Ir. L.J. van Vliet, Prof. Dr. I.T. Young, Prof. Dr. W.J. Niessen (Erasmus MC),
Dr. B. Rieger (FEI), Dr. F.M. Vos (AMC), Dr. S. Stallinga (Philips Research)

TUE-BMT
Eindhoven University of Technology, Department of Biomedical Engineering
Biomedical Image Analysis
http://w3.bmt.tue.nl/nl/onderzoek/divisies/biomedical_imaging_modeling/
Dr. Ir. H.C. van Assen, Dr. Ir. R. Duits, Prof. Dr. L.M.J. Florack, Dr. A. Fuster,
Prof. Dr. Ir. B.M. ter Haar Romeny, Dr. A. Vilanova,
Prof. Dr. Ir. M. Breeuwer (Philips Healthcare)

TUE-ET
Eindhoven University of Technology, Department of Electrical Engineering
Design Methodology for Electronic Systems
http://www.es.ele.tue.nl/
Prof. Dr. Ir. T. Basten, Prof. Dr. H. Corporaal, Dr. Ir. M.C.W. Geilen,
Prof. Dr. Ir. G. de Haan, Prof. Dr. Ir. R.H.J.M. Otten, Dr. Ir. S. Stuijk

TUE-WI
Eindhoven University of Technology, Department of Mathematics and Computer Science
Visualization group
http://w3.win.tue.nl/nl/onderzoek/onderzoek_informatica/visualization/
Dr. A.C. Jalba, Prof. Dr. Ir. R. van Liere, Dr. M.A. Westenberg,
Dr. Ir. H.M.M. van de Wetering, Prof. Dr. Ir. J.J. van Wijk

UL-LIACS
Leiden University, Faculty of Sciences, Leiden Institute of Advanced Computer Science,
HPC and Media Group
http://www.liacs.nl/research/
Dr. E.M. Bakker, Prof. Dr. Ir. E.F.A. Deprettere, Dr. D.P. Huijsmans, Dr. Ir. A.C.J. Kienhuis, Dr. M.S. Lew, Dr. Ir. T.P. Stefanov, Prof. Dr. H.A.G. Wijshoff, Dr. A.A. Wolters

UL-LUMC
Leiden University Medical Centre, Division of Image Processing (LKEB)
Clinical and experimental imaging
http://www.lumc.nl/con/1010/83058/87360/87377/
Prof. Dr. Ir. J.H. Reiber, Dr. Ir. B.P.F. Lelieveldt

UT-EWI-CAES
University of Twente, Faculty EWI/CTIT
Computer Architecture for Embedded Systems
http://caes.ewi.utwente.nl/
Prof. Dr. Ir. Marco Bekooij, Dr. Ir. Hans Kerkhoff, Dr. Ir. André Kokkeler, Dr. Ir. Jan Kuper,
Dr. Ir. Sabih Gerez, Prof. Dr. Ir. Gerard J.M. Smit

UT-EWI-DACS
Faculty EWI/CTIT, University of Twente
Design and Analysis of Communication Systems
http://utwente.nl/ewi/dacs/
Prof. Dr. Hans van den Berg, Dr. Ir. Pieter-Tjerk de Boer, Dr. Tiago Fioreze,
Prof. Dr. Ir. Boudewijn Haverkort, Dr. Ir. Geert Heijenk, Dr. Ir. Georgios Karagiannis,
Dr. Ir. Aiko Pras, Dr. Anne Remke, Dr. Ramin Sadre

UU-ICS-GMT
Utrecht Institute for ICT research, Dept. of Information and Computer Science, Faculty of Science, Utrecht University
Game and Media Technology
http://www.cs.uu.nl/groups/mg/
Dr. M. van Kreveld, Prof. Dr. M.H. Overmars, Dr. F. van der Stappen, Prof. Dr. R.C. Veltkamp

UvA-FNWI-CSA
University of Amsterdam, Faculty of Sciences, Informatics Institute,
Computer Systems Architecture Group
http://www.science.uva.nl/research/csa/
Dr. Clemens Grelck, Prof. Dr. Chris Jesshope, Dr. Andy Pimentel

UvA-FNWI-IAS
University of Amsterdam, Faculty of Sciences, Informatics Institute
Intelligent Autonomous Systems
http://www.science.uva.nl/research/ias/
Dr. Ir. L. Dorst, Dr. V. Evers (since 1-1-2010), Prof. Dr. D.M. Gavrila,
Prof. Dr. Ir. F.C.A. Groen, Dr. Ir. B.J.A. Kröse, Dr. A. Visser, Dr. S. Whiteson

UvA-FNWI-ISIS
University of Amsterdam, Faculty of Sciences, Informatics Institute
Intelligent Sensory Information Systems
http://www.science.uva.nl/research/isis/
Dr. J.M. Geusebroek, Dr. T. Gevers, Prof. Dr. M.L. Kersten, Dr. D.C. Koelma,
Prof. Dr. Ir. A.W.M. Smeulders, Dr. M. Worring

UvA-FNWI-SNE
University of Amsterdam, Faculty of Sciences, Informatics Institute
System and Network Engineering (SNE)
http://www.science.uva.nl/research/sne/
Dr. M. Cristae, Dr. Yuri Demchenko, Dr. P. Grosso, Prof. Dr. Ir. C.T.A.M. de Laat,
Prof. R. Meijer, Dr. A. Taal, Dr. Z. Zhao

VU-EW-CS
VU University Amsterdam, Faculty of Sciences, Division of Computer Science
Computer Systems
http://www.cs.vu.nl/en/research/computer-systems/
Prof. Dr. Ir. H.E. Bal, Dr. H. Bos, Dr. B. Crispo, Dr.-ing. Habil. T. Kielmann,
Dr. G. Pierre, Dr. M. Rieback, Dr. F.J. Seinstra, Prof. Dr. Ir. M. van Steen,
Prof. Dr. A.S. Tanenbaum, Dr. S. Voulgaris

Overview of ASCI Research

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 during the last six years. 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 dramatically, 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. Six years ago, ASCI´s research on embedded systems was described as *small but growing*, and indeed it has lived up to its expectations and 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. An important future research direction for Computing is Green IT; several ASCI groups already address aspects of green computing and communication, such as energy-efficient processor and network design. In addition, there is much interest in cloud computing in ASCI. 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.

	Methods & Algorithms	*Systems & Architectures*
Computing	A 1. High Performance Computing *Wijshoff (UL)* 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), Deprettere (UL)* 4. Sensor networks *Langendoen (TUD), Van Steen (VU)*
(Sensory) Information	C 1. Image and Multimedia Sensing *Van Vliet (TUD), Ter Haar Romeny (TUE)* 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), Lew (UL)* 2. Sensing and Learning Systems *Niessen (EUR), Lelieveldt (LUMC), Karssemeijer (RUN), Reinders (TUD), Groen (UvA)* 3. Acting and Visualization Systems *Gorte (TUD)*

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.

Achievements

The major scientific achievements of ASCI are its DAS infrastructure and the numerous externally funded 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-4 is scheduled to be operational in the fall of 2010. The successive systems were funded largely by four equipment grants from NWO. Each generation consists of four to six of clusters at different locations, connected by a wide-area network and integrated into a single, shared, distributed system. DAS-3 (2006) pioneered a novel flexible dedicated optical interconnect *Starplane* provided by SURFnet. DAS-4 will focus on diversity: it will contain a wide variety of accelerators, such as graphical processing units and large-scale multi-cores. It is also designed to allow research on Green IT, cloud computing, and (through participation of ASTRON) astronomy applications. It can also be used for routine work in scientific research for exploring and analyzing enormous volumes of data in multimedia archives.

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. During the past evaluation period, however, 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 was a 30 MEuro BSIK program that ran 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 was well connected to industry and non-profit organizations. The operational model was based on the work-table model where intensive co-operation at the lowest level of the organization took place and immediate transfer of knowledge and know-how was achieved. MultimediaN lives on in the national e-Science Research Center where it will contribute to the large-scale content processing, and in the FES2009 COMMIT program.

VL-e (Virtual Laboratory for e-Science) was a 40 MEuro BSIK program running from 2004 to 2009 that studied virtual laboratories for e-Science. About one third of this program consisted of Computer Scientists (mostly from ASCI) who designed generic methods and tools for scientific applications from high-energy physics, medicine, bioinformatics, biodiversity, and other areas. This program boosted our research on grid programming environments, workflow systems, problem solving environments, resource management, and networking. The work will be continued in the national e-Science Research Center to sustain e-Science research in the Netherlands as well as the FES2009 COMMIT program.

ICIS. The aim of the Interactive Collaborative Information Systems (ICIS) project is to develop techniques that support humans in complex information environments and that facilitate distributed decision-making capabilities. ICIS emphasizes the importance of building actor-agent communities: close collaborations between human and artificial actors that highlight their complementary capabilities, and in which task distribution is flexible and adaptive. To fulfill such a prospect, we need intelligent systems that observe their environment, interpret and fuse information, have learning and decision-making capabilities, and have the ability to work in teams. It moreover studies the interaction of humans with their artificial counterparts in such settings and how to meet the mutual information needs.

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 collaborated 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.

The BSIK projects MultimediaN, VL-e and ICIS have been formally evaluated. Their midterm reviews were performed by international assessment committees and the final evaluations were done by the 'Committee of Wise Persons'. All three projects received very positive reviews on all accounts. Many ASCI members played a leading role in the current FES2009 proposal COMMIT, especially profs. Smeulders and Lagendijk are main authors of the proposal and are on the board of the directorate. In July 2010, the Dutch government decided to invest about 50 MEuro in the COMMIT program.

Evaluation by a Peer Review Committee (Committee Zwaenepoel)

Dutch Computer Science research has recently been formally evaluated through a committee of internationally acclaimed scientists among the best in the field:

- Marta Kwiatkowska (Oxford University)
- Carlo Ghezzi (Politecnico di Milano)
- Frans Kaashoek (MIT)
- John Mylopoulos (University of Toronto and University of Trento)
- Thomas Ertl (University of Stuttgart)
- Wiebe van der Hoek (University of Liverpool)
- Willy Zwaenepoel (EPFL, chair).

This evaluation covered the period 2002-2008 following VSNU/NWO/KNAW's Standard Evaluation Protocol 2003-2009. The grades of the ASCI participants are given below.
(3 = good; 4 = very good; 5 = excellent).

	Research Program	Quality	Productivity	Relevance	Vitality and Feasibility
EUR-UMCR	Not evaluated				
RUG-CS-IS	4	4.5	4	3.5	4
RUG-CS-SVCG	4.5	4.5	4.5	4	4.5
RUN-UMCR	Not evaluated				
TUD-EWI-MM-CGCC	4.5	4.5	4	4.5	4.5
TUD-EWI-MM-ICT/1+2	4.5	4.5	5	5	5
TUD-EWI-ST-PGS	5	4.5	5	5	5
TUD-L&R-FRS	Not evaluated				
TUD-TNW-QI	Not evaluated				
TUE-BMT	Not evaluated				
TUE-ET	Not evaluated				
TUE-WI	5	5	4.5	5	5
UL-LIACS	4.5	4.5	4.5	4.5	4.5
UL-LUMC	Not evaluated				
UT-EWI-CAES	Not evaluated				
UT-EWI-DACS	4.5	4.5	4.5	4.5	4.5
UU-ICS-GMT	5	5	5	4	5
UvA-FNWI-CSA	3.5	3.5	3	3	4
UvA-FNWI-IAS	4.5	4.5	5	4.5	4.5
UvA-FNWI-ISIS	5	5	5	5	5
UvA-FNWI-SNE	3.5	3.5	4	4	4
VU-EW-CS	5	5	5	5	5

Some quotes from the evaluation report relevant to ASCI:

- *In general, computer science in the Netherlands is a vibrant enterprise. In each department the committee saw strong evidence of excellence and in many departments a distinct improvement over the course of the evaluation period. As a country, the Netherlands remains among the top nations in computer science research, and in the absolute top in a number of sub-areas.*

- *The PhD students and postdocs the committee met from the various departments are diverse, motivated and well-prepared to take on a research career.*

- *The research schools have come into their own, and have clearly had a salutatory effect on graduate education in the Netherlands.*

- *The research schools have clearly had a beneficial effect on the quality of graduate education in the Netherlands. All departments participate in one or more of these research schools. The PhD students benefit from exposure to researchers from other institutions and from courses that would otherwise not be available in their own institution.*

ASCI is proud to receive such unsolicited comments of the international review committee, and takes these words as the biggest motivation for its existence.

Group Contributions

Biomedical Imaging Group Rotterdam

Departments of Radiology and Medical Informatics, Erasmus MC, Erasmus University Rotterdam

Dr. M. de Bruijne, Dr. S. Klein, Dr. E. Meijering, Prof. Dr. W.J. Niessen,
Dr. J. Veenland, Dr. H. Vrooman, Dr. Th. Van Walsum

Background

The Biomedical Imaging Group Rotterdam (BIGR) is an initiative of the Erasmus MC (University Medical Center Rotterdam). It is a collaboration between the Departments of Medical Informatics and Radiology. 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 is organized in six main lines, viz. Neuro image analysis, Cardiovascular image analysis, Cellular and molecular image analysis, Oncological image analysis, Image guided Interventions, and Model-based medical image analysis. BIGR has strong collaborations with TU Delft, through joint appointments and a large number of joint research projects, and with Delft and Leiden, as part of the Medical Delta initiative.

Research plan

In 2010–2014, research in BIGR will be continued in all research lines. The Model-based medical image analysis research group, which was recently initiated, will develop novel methodologies in the fields of medical image segmentation, image registration, and machine learning, in order to support all other research lines, which are more application oriented. In the coming years, the focus within the fields of Cardiovascular and Neurovascular image analysis will be on the development and validation of quantitative imaging biomarkers for the early and differential diagnosis of cardiovascular, neurovascular and neurodegenerative disease. Methods will be implemented and validated by applying them in the context of large clinical or population cohorts. In the Oncological image analysis line, the focus will also be on imaging biomarkers, but in this case primarily focused on the prediction and/or monitoring of therapy effect. Integrated analysis of imaging data acquired with nuclear imaging techniques, and multi-modal MR will be employed in this research line. In the Cellular and molecular image analysis research line, the focus will be on the development of methods to quantitatively analyze processes at the cellular and molecular level, through the analysis of imaging time series. In this research line, collaboration with biologists will further be strengthened. The research within the field of Image guided interventions will focus on improved image guidance in minimally invasive interventions, with applications in interventional radiology, interventional cardiology, and radiotherapy. Hereto, techniques will be developed which in real-time map diagnostic quality pre-operative imaging data, to the per-operative situation.

In the longer-term, the research of BIGR into the extraction of relevant information from biomedical imaging data, will increasingly be complemented with information extraction from other domains (genomics, proteomics, metabolomics) in order to improve (early and differential) diagnosis and prognosis of a wide range of diseases, using integrated analysis of imaging and non-imaging data.

Intelligent systems

Johann Bernoulli Institute, Faculty of Mathematics and Natural Sciences, University of Groningen

Prof. Dr. M. Biehl, Prof. Dr. N. Petkov, Dr. M.H.F. Wilkinson

Background

The research area concerns interrelated topics from image processing and analysis, computer vision, pattern recognition, machine learning and brain-like computing. In biologically motivated image processing and pattern recognition (brain-like computing), models of the human visual system are developed and used in computer algorithms. Research on connected operators entails algorithm development (including parallelization), new classes of filters, applications to 2-D and 3-D medical images, and the new connectivity measures for these filters for increased robustness. In machine learning, the focus is on prototype-based learning schemes that involve automatic feature extraction by relevance learning. Further research concerns visual arts and health care applications.

Research plan

In biologically motivated image processing and pattern recognition (brain-like computing), we will focus on modeling the function of cortical areas responsible for shape processing and object recognition. We will apply the developed algorithms to practical problems, such as the recognition of vascular junctions in fundus images, handwriting and face recognition. In connected operators we making extensions to the new domain of hyperconnections and attribute space connections, which can deal with overlapping structures, both in image filtering, object detection and segmentation. We are developing theory and algorithms and applying these in applications ranging from astronomical imaging, through historic document processing to biomedical imaging. In machine learning, we will extend our studies of similarity based supervised and unsupervised data analysis. The focus will be on novel and adaptive similarity or distance measures and their use in prototype based clustering, classification, and visualization. We will apply our methods, for instance, in practical classifications problems from the biomedical domain and in image processing tasks.

In a joint project with the department of Dermatology we will develop a content-based image retrieval system that will support the diagnostic process. It makes use of the image database of that department which comprises 75,000 images and grows with 15,000 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 (165,000) people within the LifeLines study. We collaborate with the Ophthalmology department of Columbia University on the orientation analysis of retinal nerve fibers. In a joint project with researchers from the Medical School of the University of Birmingham, we apply machine learning is applied in the context of tumor classification and develop a novel practical tool for the diagnosis and monitoring of adrenal tumors. Further collaborations have been started in the areas of neurology and psychiatry.

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.

Scientific Visualization and Computer Graphics

Johann Bernoulli Institute, Faculty of Mathematics and Natural Sciences, University of Groningen

Dr. H. Bekker, Dr. T. Isenberg, Prof. Dr. J.B.T.M. Roerdink, Prof. Dr. A.C. Telea

Background
The research group Scientific Visualization and Computer Graphics is part of the Johann Bernoulli Institute for Mathematics and Computing Science, Faculty of Mathematics and Natural Sciences, Un. of Groningen. The group leader has a joint appointment with the University Medical Center Groningen. Website: http://www.cs.rug.nl/svcg.
The group 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), astronomy, and large-scale software engineering. Earlier work focused on multiscale visualization using wavelets and morphological pyramids with applications in functional neuroimaging, 3D shape analysis, surface rendering, GPU-based rendering, pathway and network visualisation with applications in bioinformatics and EEG analysis, perception-based visualization, and high-dimensional data visualization with applications in astronomy. In 2007 the group expanded its research by attracting new staff members, one in the area of information and software visualization, and another one in non-photorealistic rendering and innovative interfaces.

Research plan
In interactive data visualization, we focus on 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. We take advantage of knowledge about human perception to improve current visualization techniques, for example, in graph and network visualization.
In software visualization and program understanding, we study 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 engineering pipelines.
We develop non-photorealistic rendering (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.
We apply our methodological work in several application domains. In biomedical imaging, 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 very high-dimensional parameter spaces.
We plan to combine interactive network visualization with software visualization to understand complex software systems. Methods for quantitative graph comparison have to be developed for this purpose. Such methods will find wide application in other applied fields, such as biological networks in genomics or neuroscience. For developing new ways of interacting with large touch-sensitive displays we will focus on collaborative aspects. Inspiration from the artistic discipline is obtained through a recently initiated collaboration with the Minerva School of Arts in Groningen. Also various contacts with partners from industry or GTIs will be explored for joint projects.

There are two new international initiatives in which the group wants to extend its efforts: Visual Analytics and Neuroinformatics. Visual Analytics concerns the integration of visualization with other analytical methodologies, such as statistics, data-mining, and cognition. Neuroinformatics concerns the support of discovery and innovation in neuroscience.

Diagnostic Image Analysis Group (DIAG)
Department of Radiology, Radboud University Nijmegen Medical Centre

Dr. Bram van Ginneken, Dr. Ir. HenkJan Huisman, Dr. Ir. Nico Karssemeijer,
Dr. Clarisa Sanchez, Dr. Peter Snoeren

Background

In medical imaging, the accurate diagnosis and/or assessment of a disease depends on both image acquisition and image interpretation. The complexity of imaging data has increased tremendously with the introduction of digital technology, often making it more difficult and time consuming for clinicians to perform their tasks. Therefore, it is not surprising that in clinical practice more complex imaging techniques only become widely accepted, if their use is simple and supported by adequate image processing and analysis. The development of dedicated image analysis platforms for diagnostic radiology is a highly challenging task. The ultimate goal is to make computations unbiased with respect to the type of scanner and acquisition protocol, gradually moving radiology into the domain of exact quantitative imaging. The promise is enormous, as this may lead to definition of imaging biomarkers that can be used in diagnosis and in population studies aimed at screening and risk assessment.

In medical imaging, the image interpretation process is performed by humans. This interpretation is limited by the non-systematic search patterns of humans, the presence of structure noise (camouflaging normal anatomical background), and the presentation of complex disease states requiring the integration of vast amounts of image data and clinical information. Computer aided diagnosis (CAD) is a new field in medical imaging which aims at development of methods to aid radiologists with search and image interpretation. This technology has been adopted in breast cancer screening programs to help radiologists with the reading of mammograms, while applications are also developed for CT lung and colon cancer screening. The idea of CAD is that computers can be trained to detect and recognize abnormalities in imaging data automatically, and that CAD can be used to improve the quality of image interpretation by providing a second opinion. In addition, they can be used to build intelligent navigation aids to help readers to efficiently interpret high dimensional imaging data, and they can be used in tasks involving detection of changes over time in subsequent imaging studies, a crucial aspect of screening for cancer.

Research plan

In the coming years research will be focused on development of automated methods for quantitative (functional) image analysis in breast, lung, prostate, retinal, and neuro imaging. Methods will include automated detection of potential disease locations using novel image analysis and pattern recognition methods. In this way time consuming automated analysis can be carried out before the images are presented to the reader. To support development of robust and reliable computer aided detection systems, a significant investment will be made to build and maintain annotated structured databases for the disease patterns of interest. Huge databases for breast, lung and prostate cancer that were developed previously by the group form the basis of successful research today. Expertise in this area will be exploited by expanding application areas.

A new area of research that will be developed in the coming years is quantitative analysis of (perfusion) CT. This is related to increased interest in CT in the radiology department. Research in breast, prostate and lung imaging will be continued and expanded. In breast imaging, a leading theme will be integration of imaging modalities and breast cancer risk assessment. In chest imaging, automated detection of tuberculosis for application in developing countries will be an active area of research. A close collaboration with Fraunhofer MeVis (Bremen) is currently established by foundation of a MeVis CAD group in Nijmegen. This will be of great benefit for further development of the group.

The ambition of the group is to support multi-disciplinary clinical research with development of innovative image analysis methods and to strengthen its position as a world wide leading center for computer aided detection and diagnosis by extending the field of applications.

TUD-EWI-MM-CGCC C3

Computer Graphics

Department of Mediametics, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology

Dr. Ir. A.R. Bidarra, Dr. Ir. C.P. Botha, Dr. W.F. Bronsvoort, Prof. Dr. Ir. F.W. Jansen, Ir. F.H. Post

Background

Computer Graphics concerns the modeling and display of 3D objects and multi-dimensional data. Initially, main application areas were in product modeling for CAD/CAM and in scientific visualization for computational fluid dynamics (CFD). Recently, the emphasis has shifted towards Large Data Visualization, Medical Visualization and Game Technology.

The group has participated in several BSIK-programs (VL-e, BRICKS, LWI) and is currently participating in the FES-program GATE. For the Medical Visualization there is extensive collaboration with the LUMC-Leiden and EMC-Rotterdam. In the coming year, the chair (and program lead) will become vacant. A new profile ‘Interactive visualization’ has been drafted and recently approved.

Research Plan

In Large Data Visualization, we visualize environmental data (point clouds) and climate data (e.g. for extreme weather predictions). Real-time simulation and integrated visualization with computational steering is now feasible on desktop computers with GPU support. For the coming years we will further explore applications within environmental and climate research applications.

In Medical Visualization we focus on the visualization of multi-field data (CT and MRI) and higher order data such as diffusion weighted imaging (DWI). In addition, we develop visualization tools for minimally invasive techniques for instance for polyp detection in virtual colonoscopy, and for surgical planning and intra-operative guidance for orthopaedics, such as shoulder replacement surgery.

We have recently started projects on molecular imaging in order to visualize biological processes at cellular and molecular level over multiple time points and across multiple subjects, and also on the visual analysis of medical retrospective cohort studies, containing multi-modal multi-timepoint datasets of multiple patients.

In Game Technology, we focus on content generation, in particular the dynamic generation of virtual worlds. Ultimate goal is to effectively assist game level designers in expressing and consistently maintaining a model of the virtual world with all design intent maintained throughout the various iterations of the design process.

There is an explosion of new data sources for 3D and multidimensional data (laser point clouds, video, new medical imaging modalities). Visualization techniques, often in the form of interactive visual analysis, will play a crucial role in enabling humans to cope with the magnitude and heterogeneity of ever-growing datasets.

Multimedia Signal Processing Group
Department of Mediamatics, Faculty of Electrical Engineering, Mathematics and Computer Science,
Delft University of Technology

Prof. Dr. Ir. J. Biemond, Dr. A. Hanjalic, Prof. Dr. Ir. R.L. Lagendijk, Dr. M. Larson, Dr. J.C.A. van der Lubbe,
Dr. P. Peris Lopez, Dr. P. Serdyukov, Dr. T. Veugen, Prof. Dr. Ir. A. de Vries

Background

The mandate of the Multimedia Signal Processing (MSP) group, established in January 2010, is to develop new techniques for multimedia processing and retrieval that will meet user needs in the areas of multimedia content enrichment, search, and management. The *Information Security and Privacy research lab (ISP)* of the MSP group recognizes the critical role of trust and data protection in such multimedia applications. Its focuses on developing algorithms that assure content security and user privacy. We have gained experience in privacy-enhanced technologies via several past projects, in particular EU FP6 FET project SPEED (Signal Processing in the Encrypted Domain) and the Dutch PAW project (private computing). The *Delft Multimedia Information Retrieval lab (DMIR)* of the MSP group develops technologies designed to provide users with access to multimedia content that is situated within a specific social and functional context. The emergence of large of multimedia collections associated with networked communities has motivated us to focus our research effort on the exploitation of a 'Triple Synergy', consisting of content analysis, social annotation and user-network structures. Past projects have laid the groundwork for the currently ongoing EU FP7 Network of Excellence PetaMedia devoted explicitly to using the Triple Synergy for multimedia access and retrieval.

Research plan

Our research planned for the timeframe 2010–2014 will concentrate on exploiting the Triple Synergy for content enrichment and search engines, and on privacy protection of user-provided information in these processes. Projects within which this research will be carried out include PetaMedia, as well as other currently ongoing projects, in particular, the EU FP7 STREP project PuppyIR (developing information retrieval for children), Dutch STW project Kindred Spirits (privacy aspects of social networks), and the Dutch STW project Pearl (RFID security).

The combination of visual multimedia content analysis with social annotations will be explored in the areas of multimodal video retrieval, music retrieval and spoken-content-based multimedia retrieval. We will give particular attention to the intent and goals that motivate users to turn to multimedia search to satisfy their content needs. The combination of complex network structures and social annotations will be exploited in areas such as trust-aware recommendation, network-informed collaborative tagging, and user-generated metadata for P2P networks. A key aspect will remain the development of algorithms that are able to deal with the noisy annotations produced by automatic metadata generation techniques. Approaches such as re-ranking and fuzzy search have proved effective in the past and will continue to be a subject of investigation.

Two areas will be emphasized that are aimed at transcending traditional forms of information retrieval. First, improving multimedia information retrieval by integrating information orthogonal to topic (e.g., affect), and second, moving beyond the document to access of information objects. We will devote attention to developing tasks for benchmark evaluation, along with user-centered evaluation metrics. Specific effort will be devoted to the privacy aspects that arise in social networks and in the networked communities surrounding multimedia collections. A range of secure signal processing techniques for multimedia will be developed, such as privacy-enhanced recommender systems, location-based search systems, and systems for social annotation and automatic metadata generation (e.g., detection of faces).

The research will pursue algorithms for integrated solutions that address multimedia challenges within the social and functional context in which they arise, where users and users' needs are the source.
In the development of new cryptographic protocols, user needs will be central (social networks, paid-TV, recommender systems).

Pattern Recognition and Bioinformatics section (PRB)
Department of Mediametics, Faculty of Electrical Engineering, Mathematics and Computer Science,
Delft University of Technology

Dr. Ir. R.P.W. Duin, Dr. E.A. Hendriks, Dr. M. Loog, Prof. Dr. Ir. M.J.T. Reinders,
Dr. Ir. D. de Ridder, Ir. J. de Ridder, Dr. D.M.J. Tax

Background
Pattern recognition entails the processing of raw measurement data by a computer to arrive at a prediction which can then be used to formulate a decision or action to take. Problems to which pattern recognition is applied have in common that they are too complex to model explicitly, and therefore require algorithms to learn parameters in generic models from limited sets of examples. Pattern recognition practice is firmly focused on real-world, sensor-based applications. This places it at the core of the current process of scientific discovery, by allowing researchers to derive regularities in large amounts of data in areas as diverse as physics, biology and geology, but also psychology and neuroscience. Pattern recognition algorithms also find application in industrial and consumer settings, allowing machines to sense the environment and to decide on actions or support human decision making. The PRB section studies both aspects in three different research labs. One research lab (*pattern recognition lab*) focuses on the foundations of pattern recognition: representation and generalization, in which new ways of describing objects and learning from examples are studied. The other two other research labs (the *computer vision lab* and the *bioinformatics lab*) apply these techniques in the domain of images and in molecular biology.

Research plan
The Pattern Recognition Lab works following the classical trinity of representation, generalization, and evaluation, the core elements of every pattern recognition system. The principal focus is on developing tools and theories and gaining knowledge and understanding applicable to a broad range of general problems. Typically, this involves sensory data, e.g. time signals, images, video streams, or other physical measurement data. The lab has significantly contributed to the field, mainly on neural networks, the use of dissimilarities and one-class classifiers. Current major research directions are dissimilarity-based pattern recognition, multiple classifier systems, and multiple instance learning.
The Delft Bioinformatics Lab deals with developing novel computer models and algorithms to further fundamental biological knowledge and apply these models and algorithms to advance the state-of-the-art in health care and industry. The lab focuses on data-driven bioinformatics: creating algorithms to infer and exploit simple models of complex interactions, by coupling biological insights and available prior knowledge to high-throughput measurements. Recent contributions include the discovery of novel cancer genes by analyzing and modeling insertional mutagenesis data; proposing (combinatorial) cultivation dependent transcription factor activities based on a decomposition of transcriptomics data; and assessing gene therapy protocols by integrating viral insertions with gene expression data.
The Computer Vision Lab focuses on the segmentation and analysis of multidimensional sensor data (image sequences, multiple cameras, 3D/4D medical data like MRI and CT). Segmentation is used in the broadest possible sense, i.e. distinguishing any relevant image information from non-relevant information. The main research areas currently covered are 3D imaging (camera calibration, disparity estimation, 3D reconstruction, 3DTV/free viewpoint rendering), biomedical imaging (medical image segmentation, 2D/3D model reconstruction), social or human signal processing (pose and gesture recognition and tracking) and surveillance (video object detection, recognition and tracking).

The Pattern Recognition Lab will focus its research especially on alternative evaluation functions and semi-supervised and active learning. The Delft Bioinformatics Lab will study robustness in microbial systems, application of scale-space theory and the building of atlases. The Computer Vision Lab will intensify the collaboration with the Image Processing Division of the Leiden University Medical Center).

Parallel and Distributed Systems

Department of Software Technology, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology

Dr. S.O. Dulman, Dr. Ir. D.H.J. Epema, Dr. A. Iosup, Prof. Dr. K.G. Langendoen,
Dr. Ir. J.A. Pouwelse, Prof. Dr. Ir. H.J. Sips

Background

The mission of the Parallel and Distributed Systems (PDS) Group of TU Delft is to contribute to the scientific advancement in selected fields of parallel and distributed systems, both in a fundamental as well as in an experimental way – fundamental in the sense that we aim at the development and the evaluation of new *generic concepts*, experimental in the sense that we also develop and investigate *actual implementations* of these concepts in real environments (e.g., the DAS and a 100-node sensor network test bed that we have recently installed), to the point that these implementations are mature enough to be used by other users (e.g., the Tribler P2P system and the KOALA grid scheduler).

Research plan

The PDS group focuses its research on grid computing, multicore programming, peer-to-peer systems, and wireless sensor networks (WSNs). In grid computing, we investigate scheduling and resource management, and our research is centered around the KOALA grid scheduler. In multicore programming, we focus on stream processing, which stems from both consumer electronics applications and scientific applications where data streams are generated by a large numbers of sensors (e.g., radio telescopes). In peer-to-peer (P2P) systems, we research video distribution in the widest sense of the word, that is, file sharing, live streaming, and video on demand, but also content recommendations, keyword search, etc.; our research here is centered around the Tribler P2P client. In WSNs, the main research issues are resource management, in particular energy and memory footprint, and dependability.

In the area of grids and clouds, we plan to work on resource management and scheduling for specific application areas such as massive multiplayer online games and bioinformatics, on resource management and scheduling in hybrid systems, that is, systems with accelerators such as GPUs (the next-generation DAS-4 system will have such an architecture), and on failures (types, patterns of occurrence, etc.) in large-scale distributed systems. In multi-core programming, research will be directed to programming methods and frameworks and performance prediction for modern multi-core systems, like the Cell and GPUs. The focus will be on programming and optimization strategies for multi-cores and system-wide dimensioning techniques for complete applications that include multi-core components. In the area of P2P systems, we plan to work, among other things, on the measurement and the modeling of swarm-based P2P systems, on designing and testing trust and reputation mechanisms, and on improving the performance of the three download models of video. All of these elements will be incorporated into the Tribler P2P system and will be analyzed on the basis of real experiments. In WSNs, we are engaged in integrating our test bed in a European-wide test facility for WSN research. In addition, we are taking our software systems to a new challenge by deploying them on flocks of birds to study their behavioral patterns, which raises the required level of autonomy and robustness to new heights.

In the future, the internet will be dominated by clouds, data centers, and social networks, with huge numbers of wireless and mobile devices surrounding them. We will research the design of and the resource management in these infrastructures, including energy efficiency and multi-core technology.

Optical and Laser Remote Sensing (OLRS),
Department of Remote Sensing, Faculty of Aerospace Engineering, Delft University of Technology

Dr. Ir. B.G.H. Gorte, F. Karimi Nejadasl, A. Bucksch

Background

Over a recent period of years (2004–2008) the chair of OLRS, as well as two of the lecturer positions have been vacant. One assistant professor and two PhD candidates were the only ASCI members during that time. Since December 2008 the chair has been re-occupied. Subsequently, the lecturing positions were filled and the number of PhD projects increased drastically, with some more to start soon. This will result in an increasing ASCI-participation shortly.

Research plan

One of the projects executed within the ASCI context concerned the analysis of helicopter-recorded image sequences to measure and monitor driver behavior in traffic congestions on motorways. This was done in collaboration with the Faculty of Civil Engineering and Geosciences at TU Delft and funded by NWO. The resulting thesis was defended in February 2010. Another project was using terrestrial LiDAR to measure the structure of trees and other forest-parameters like wood volume and biomass. The resulting thesis is expected to be defended in the beginning of 2011. The study was performed in co-operation with the University of Goettingen and the German LiDAR manufacturer Z&F.

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 analyzed at OLRS. 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 equally great interest for very diverse applications in medicine, forestry, flood modeling, urban climate modeling and land use/land cover mapping.

The range of topics studied at OLRS is being expanded to include integrated physics-based retrieval of parameters of geo- and biodynamical processes.

TUD-TNW-IST-QI C1, D2

Quantitative Imaging Group
Department of Imaging Science & Technology, Faculty of Applied Sciences, Delft University of Technology

Prof. Dr. Ir. L.J. van Vliet, Prof. Dr. I.T. Young, Prof. Dr. W.J. Niessen (Erasmus MC), Dr. B. Rieger (FEI), Dr. F.M. Vos (AMC), Dr. S. Stallinga (Philips Research)

Background

Our mission is to invent 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 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, non-linear image filtering, multi-orientation representation as well as to the development of dedicated algorithms for image analysis..

Research plan

Basic work is also the backbone for successful applications in medical image analysis. We belong to the world-top in computer cleansing for Virtual Cononoscopy 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 will further improve our methodology by permitting limited patient preparations and low-dose CT recordings. Cardiovascular risk assessment is vital in second prevention and poses a multi-modal imaging challenge for the main cause of morbidity and mortality in the western world. We develop and validate advanced quantitative image analysis techniques to optimally exploit the rich information present in these complementary imaging techniques, to permit (early) detection, diagnosis, therapy planning, guidance and monitoring of cardiovascular disease by providing detailed information on patient anatomy, function and pathology. To differentiate in an early stage between physiological (aging) and pathological (e.g. Alzheimer's disease) conditions in the brain, we search for spatiotemporal biomarkers of neurological disorders using sophisticated image processing and pattern recognition methods in Diffusion Weighted MRI data.
Diagnostics in patient imaging relies heavily on molecular knowledge unraveled by microscopical techniques in pre-clinical research. In the coming years we will extend our efforts in quantitative microscopy by pursuing innovations in computation imaging. In electron tomography we aim for nanometer resolution by 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. In fluorescence microscopy, resolution can be gained by using multiple illuminations with spatially periodic illumination patterns and processing the raw images into a final high resolution image, and by exploiting the spatio-temporal characteristics of fluorophores.

Combining image information from multiple modalities across a range of scales will be a strongly emerging field. Computational imaging opens new ways to extract new information from a scene by redesigning the imaging system and data recording strategy.

Biomedical Image Analysis

Department Biomedical Engineering, Eindhoven University of Technology

Dr. Ir. H.C. van Assen, Dr. Ir. R. Duits, Prof. Dr. L.M.J. Florack, Dr. A. Fuster,
Prof. Dr. Ir. B.M. ter Haar Romeny, Dr. A. Vilanova, Prof. Dr. Ir. M. Breeuwer (Philips Healthcare)

Background

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 and advanced visualization, for cardiovascular and neurological applications. The group collaborates with the TU/e Magnetic Resonance Lab, Philips Healthcare, the FC Donders Institute, the University of Zürich, the University Hospitals of Maastricht, Utrecht, Nijmegen and Leiden, and the Epilepsy Center Kempenhaeghe. Education is an important aspect.

A multi-scale framework has been established for doing high-order differential geometry on high-dimensional images, with applications such as adaptive 'geometry-driven' edge preserving enhancement, multi-scale optic flow extraction, and deep structure analysis for content-based image retrieval (VICI Florack 2005). In collaboration with Philips Healthcare (Best), work on computer aided diagnosis has been performed focused on dynamic contrast enhanced MRI of breast tumors, low-dose catheter tracking, automated detection of pulmonary emboli, and the detection of polyps in virtual colonoscopy at strongly reduced X-ray dose of the CT scan. A toolkit has been developed for GPU-based visualization and analysis of tensor fields of DTI (Diffusion Tensor Imaging) data (VENI Vilanova 2006), and High Angular Resolution Diffusion Imaging (HARDI). This enables the interactive manipulation of all relevant parameters, 3D orientation glyphs, and tractography. The tool is used in many collaborating laboratories. The clinical branch at the Maastricht University Hospital focuses on precise navigation for Deep Brain Stimulation (DBS), in close collaboration with the neurosurgery department.

Research plan

In 2010 the research collaboration IST/e (Image Science & Technology / Eindhoven) was established, with a grant of 1 M€ of the TU/e board. Partners are the Depts. of Mathematics & Computing Science (Florack, Van Wijk) and Biomedical Engineering (Ter Haar Romeny, Nicolay).

The research plans for the period 2010-2014 focus on four themes. The first one is generic and advanced mathematical methods for *multi-valued* image analysis, for segmentation, enhancement, morphological and *contextual* operations, and tractography of crossing and splitting fibers. This visual perception-inspired 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'.

GPU-based visualization of multi-valued information is the second research issue. The visualization of uncertainties is an important research topic which has hardly been investigated. 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.

The third research subject is cardiovascular applications, primarily with X-Ray and MRI, exploiting multi-valued images for local strain analysis of ventricular deformation from tagged MRI sequences, 4D flow visualization, and cardiovascular dynamics to study local dense 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.

Finally, neuro applications, primarily focusing on DTI and HARDI exploited in tractography for brain connectivity and muscle fibre orientation analysis, form the fourth theme in our research plans.

Vision for longer-term research

Our longer term research will focus on bio-inspired mathematical algorithms for contextual analysis, on uncertainty visualization in biomedical images, on congenital heart disease, and brain connectivity. We strive towards collaboration with Philips Healthcare, Philips Research, Utrecht University, Maastricht University Hospital, and Kempenhaeghe.

Design Methodology for Electronic Systems

Department of Electrical Engineering, Eindhoven University of Technology

Prof. Dr. Ir. T. Basten, Prof. Dr. H. Corporaal, Dr. Ir. M.C.W. Geilen, Prof. Dr. Ir. G. de Haan,
Prof. Dr. Ir. R.H.J.M. Otten, Dr. Ir. S. Stuijk

Background

The mission of Design Methodology for Electronic Systems is to provide a scientific basis for design trajectories of digital electronic circuits and systems from (generalized) algorithm to realization. To identify the key problems and verify the validity, robustness and completeness of our results, we develop, implement and maintain consistent and complete design flows. This entails converting the 'art' of designing electronic systems into methodology. Due to flexibility and design-time demands, we observe a major trend in using processor platforms as design target, combining programmable execution engines with accelerator blocks, reconfigurable logic and a flexible memory hierarchy, all communicating through an interconnect network. The core of the design trajectory for these systems is the mapping of computational networks onto the platform. Video processing and vision applications are ideal vehicles for our activities: innovation in processing algorithms, improvement of design trajectories and challenging hardware solutions. An important enabling technology for both research directions mentioned above is compiler-technology for re-targetable code generation and re-configurability aiming at high performance and low power dissipation. The goal of a system-level design methodology is to cope with the increasing system complexity by raising the abstraction level at which systems can be designed and implemented. Therefore we have developed formal languages, analysis and predictable synthesis techniques and corresponding tools and we focus on model-driven development of electronic systems, aiming at predictability of the results. We have applied those concepts in design, synthesis and run-time management for multiprocessor systems-on-chip (MPSoC), wireless sensor networks (WSN), and, more recently, high-tech embedded systems such as professional printers and wafer steppers.

Research plan

We plan to continue our work on multiprocessor design flow based on dataflow analysis, focusing specifically on analyzing performance/resource usage trade-offs. We plan to further develop model-driven run-time adaptation techniques for MPSoC, aligned with the design flow. We further work on methodology and tools for design space exploration based on a sound mathematical framework and towards multi-objective optimization criteria. We will develop platforms and systems that support this approach. Composability of a platform will enable independent design, implementation, and verification of different applications that co-exist on a single system-on-chip(SOC). Avoiding interference between real-time applications in terms of performance (timing, throughput, latency, synchronization, etc.) is a novel aspect of composability. The essential weakness is its static nature. To overcome this, we explore the possibility of composable energy management, where the static worst-case dimensioning is accepted, but the run-time cost is reduced by using slack to modify the frequency and voltage operating points. A second, more speculative, direction is exploiting variability to avoid worst-case design using introspective or self-calibrating architectures. Both directions contribute to a single platform called CompSOC that serves as a proof of concept and a FPGA testbed for research and education.
Over the last ten years video processing research has focused on the development of algorithms and implementations for image-enhancement, format conversion and analysis for consumer electronics, but new applications in healthcare, lighting and lifestyle are of growing importance. We are therefore shifting our focus to include also lighting, health and wellbeing applications. We further investigate 2D to 3D-conversion and vision libraries for a GPU-platform.
To meet energy efficiency demands, we work at the circuit level on adaptive power supply control based on workload variations. Energy aware design addresses the need for more scalable systems with trade-offs between different system parameters as opposed to the typical low power design with fixed parameters. In a constantly changing environment, a scalable system is highly desired as it is capable of adapting its behavior to the constantly change in the operation conditions.

We study the automated design trajectories for future digital electronic circuits and systems with emphasis on video processing and embedded architectures, including energy efficiency, context-awareness reliability, and adaptability. Compositional reasoning, a model-driven approach and run-time adaptation are envisioned ingredients of the approach.

Visualization group
Department of Mathematics and Computer Science, Eindhoven University of Technology

Dr. A.C. Jalba, Prof. Dr. Ir. R. van Liere, Dr. M.A. Westenberg, Dr. Ir. H.M.M. van de Wetering,
Prof. Dr. Ir. J.J. van Wijk

Background

The aim of visualization research is to develop methods and techniques using interactive computer graphics so that insight in large data sets can be obtained. The key-idea is to exploit the unique capabilities of the human visual system to detect patterns, structures and outliers by translating data into images and animations. Interaction is another vital ingredient, as it enables users to focus on those aspects of the data that are the most relevant to them and to define appropriate visual mappings. The field of visualization has a long history and roots in disciplines like cartography and graphical statistics. With the advent of affordable and fast graphics hardware in the early nineties, visualization has developed into a flourishing field, with a number of specializations. The TU/e Visualization group is active in several areas. Information Visualization concerns the visualization of abstract data, such as tables, trees, networks and combinations of these. We have worked especially on Software Visualization, which is a challenging and attractive field. Scientific visualization concerns data from simulations and measurements, defined over geometric spaces. Within this area we study feature extraction, flow visualization, and, in co-operation with the TU/e Department of Biomedical Engineering, medical visualization. Furthermore, in co-operation with CWI we study how virtual reality techniques can be used for visualization purposes, and we also do work in mathematical visualization.

Research plan

Visual Analytics is a newly emerging research field, which will be our primary focus for the coming years. Visual Analytics is defined as 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. This area offers many challenges, both with respect to generic issues as well as application in practice.
Concerning generic issues, we are interested in methodological issues (such as how to design systems and how to evaluate them); a better foundation on supporting disciplines, including for instance statistics, HCI, perception, and cognitive psychology; and generic methods, for instance for navigation support and knowledge capture.
For application in practice, we aim at expanding our co-operation with experts from other domains who will serve as knowledge providers and as users. Areas of interest are mobile data (for instance on vessel, pedestrian or animal movement), multimedia data (searching and categorization), and event logs from business processes. With the advent of dr. M. Westenberg to our group, the topic of visual analytics for bio-informatics has obtained a strong boost.
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 position in scientific visualization. Co-operation on medical visualization and image analysis within TU/e has intensified.

Visualization and visual analytics can contribute significantly to a more effective and efficient analysis of data, thereby leveraging many activities in many different areas. How to deal with huge amounts of data, how to systematically develop such systems, and how to integrate methodologies from other fields are open questions.

UL-LIACS A1, A2, B3, D1

HPC and Media Group

Leiden Institute of Advanced Computer Science, Faculty of Sciences, Leiden University

Staff

Dr. E.M. Bakker, Prof. Dr. Ir. E.F.A. Deprettere, Dr. D.P. Huijsmans, Dr. Ir. A.C.J. Kienhuis, Dr. M.S. Lew, Dr. Ir. T.P. Stefanov, Prof. Dr. H.A.G. Wijshoff, Dr. A.A. Wolters

Background

Our mission is to improve the state of the art in systems and media analysis, and to guide new research areas which are clearly of importance to society. In our research, we rigorously demonstrate the effectiveness of the the newly developed techniques and contribute to leading evaluation and benchmarking projects. Our program 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 in bio-informatics. The research activities have 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.

Research plan

Our research activities have become more focused on data representation, processing data representation, and the reduction of data representation by employing (semantic) transformation and interpretation of media data. The general thought behind these activities consists of a data oriented vision of computer science whereby data and its representation is seen as starting and end point of any computational process. Currently, we will concentrate on several sources of data: data arising from sensor networks, grid computing, DNA sequence data, data reduction in astronomy, real time astronomy satellite data, imaging data, and on different data representation transforming systems, e.g. from physical model to code transformation, and from code to systems.

Among other research projects, the adopted data-centric view can be seen in the project 'Towards Data Structure Independent Computing (ToDSIC)'. In this project, we will extend previously developed restructuring techniques to pointer-based codes. Restructuring memory access patterns is the central theme in this research. By transforming pointer traversals codes into a regular intermediate, analysis of these codes is enabled and these codes can be optimized by taking many factors into account, e.g. the target architecture, the cache configuration and the number of cores. In fact, the work as proposed can be seen as a first step in realizing data structure independent programming whereby specific information on how data is stored and structured is hidden from the core application.

The state of the art shows a data explosion to be answered by the development of our paradigm of computational imagination for human-machine interactive searching for visual information. A computational imagination seeks to explore the ability of the human mind to create multimedia information which does not exist in reality in the form of synthesized imagery.

UL-LUMC D2

Clinical and experimental imaging

Division of Image Processing (LKEB), Leiden University Medical Centre

Prof. Dr. Ir. J.H. Reiber, Dr. Ir. B.P.F. Lelieveldt

Background

The main goal of the research at 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 one of the seven main research themes of the LUMC under the headings 'Vascular Medicine', 'Neuro-science' and 'Biomedical Imaging'. Part of the research involves computer vision research and algorithm development, whereas clinical applications also play an important role. Applications focus on Neuro-imaging, Pulmonology, Orthopaedics, Cardiology and molecular and cellular imaging.

Research plan

Algorithmic research at the division of image processing is organized along three lines.
Statistical shape modeling: Techniques 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.
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.
Multi-modal image fusion: Here we address novel image processing methods to combine complementary information in molecular, structural and functional imaging, concentrating on challenges brought on by new molecular imaging modalities such as bioluminescence imaging and fluorescence imaging.

In the next five years the research at LKEB will continue to focus on the development of generic methodologies for quantification, visualization and analysis techniques for biomedical images. Concrete clinical applications will involve pulmonology, particularly the analysis of pulmonary disease from follow up CT scans using local non-rigid registration methods. Another application is in orthopaedics, the pre-operative surgery simulation and post-operative evaluation of prosthesis fixation, quantification of joint degeneration in Rheumatism. Vascular image analysis (quantification and risk stratification of vascular disease from several imaging modalities) and cardiac image analysis (quantification and information fusion to support cardiac diagnostics from large, heterogeneous image data sets) are also application areas, as well as neurological image analysis (the development of shape comparison methods to detect pathologies compared to an atlas of normal subjects, and development of algorithms to monitor the progression of white matter lesions and ageing effects), molecular image analysis (fusion of and tracking in complementary imaging data from the whole-body scale to the molecular scale, with an emphasis on translational research) and generic multi-atlas registration for computer-assisted segmentation and diagnosis.

In principal, our research plans do not look further ahead than the five years described above. The longer-term research agenda will strongly depend on the developments at the biomedical front, as well as on future grant schemes.

Computer Architecture for Embedded Systems (CAES)
Faculty EWI/CTIT, University of Twente

Prof. Dr. Ir. Marco Bekooij, Dr. Ir. Hans Kerkhoff, Dr. Ir. André Kokkeler, Dr. Ir. Jan Kuper,
Dr. Ir. Sabih Gerez, Prof. Dr. Ir. Gerard J.M. Smit

Background
Energy efficiency and dependability are the main drivers of the research in the CAES group. Energy-efficiency is important for streaming applications found in battery powered mobile devices, as well as for high-performance systems. In addition to that, the group works on energy management for buildings or neighborhoods.

Research plan
The main emphasis of the group is on efficient architectures for dependable networked embedded systems. Within this theme CAES performs research on three related key areas: efficient architectures for streaming applications, architectures for efficient energy management, and dependability issues of embedded systems.
The effort on energy efficient architectures focuses on reconfigurable processors for streaming applications, such as found in battery powered mobile devices (e.g. portable multimedia players) or sensor networks. However, within high performance embedded computing (e.g. medical imaging, radar processing), extrapolating the current trend of using general purpose processors for future systems predicts excessive power consumption. MPSoC (Multi Processor Systems-on-Chip) devices for streaming applications are prime candidates for use in this application domain as well.
The focus of the energy management theme is to use ICT technology to improve the overall energy efficiency. By building so-called smart grids the overall energy efficiency can be significantly improved by choosing the best energy generator/source currently available. In this research, we have a three step approach based on prediction, planning and real-time control.
Dependability plays an important role in sensor networks (nodes may fail or run out of energy unexpectedly). In MPSoC (Multi Processor Systems-on-Chip) systems for streaming applications dependability techniques also play an important role. In the last five years the research has resulted in four start-up companies: Recore systems, Ambient Systems, HOMA Software and Smart Signs Solutions.

In the long run, the group will focus on two research areas: energy-efficient ICT systems and using ICT systems to make other systems more energy-efficient. As energy saving is becoming a crucial design factor, there will be ample opportunities to obtain funding.

Design and Analysis of Communication Systems (DACS)

Faculty EWI/CTIT, University of Twente

Prof. Dr. Hans van den Berg, Dr. Ir. Pieter-Tjerk de Boer, Dr. Tiago Fioreze, Prof. Dr. Ir. Boudewijn Haverkort, Dr. Ir. Geert Heijenk, Dr. Ir. Georgios Karagiannis, Dr. Ir. Aiko Pras, Dr. Anne Remke, Dr. Ramin Sadre

Background

DACS focuses on the design and analysis of *dependable networked systems.* A system is called dependable, when 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: "at contributing 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, so 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.

Research plan

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). DACS also covers the whole lifetime of networked systems: from initial model based analysis to protocol design and operational aspects. In the future, the research on embedded networking will focus more on truly embedded systems (and the evaluation of the performance and dependability of those systems), also in co-operation with ESI. We have already started activities in the area of model-based design for in-car automotive systems, thereby adhering to architectural design languages and systems such as Autosar and East-ADL. Furthermore, we will focus on the design and evaluation (security, dependability, survivability) of embedded networking technology as part of vital infrastructures, such as industrial plants, water and energy networks. The design of dependable wireless systems is shifting its focus towards vehicular networking. The new focus is on using communication between vehicles to improve traffic efficiency, safety, and environmental impact. Recently, some new externally funded research projects in this area have been started, and a number of proposals are being evaluated for funding. Finally, research on operational aspects will be extended towards wireless networks. For that extension, the focus will be on WLAN, since it is, as opposed to GSM and UMTS, relatively easy to obtain measurements data for such networks. Furthermore, with the increasing use of 10Gbps technologies for the core Internet, scalability of measurement systems becomes problematic. Therefore the focus will be shifted to flow-based measurement and analysis techniques (Netflow, IPFIX, sampling).

The research will focus on truly embedded systems, the design and evaluation of embedded networking technology as part of vital infrastructures and on communication between vehicles to improve traffic efficiency, safety, and environmental impact.

UU-ICS-GMT D1

Game and Media Technology

Utrecht Institute for ICT research, Dept. of Information and CS, Faculty of Science, Utrecht University

Dr. M. van Kreveld, Prof. Dr. M.H. Overmars, Dr. F. van der Stappen, Prof. Dr. R.C. Veltkamp

Background

The research area of Game and Media Technology focuses on the various aspects of algorithms and technology in interactive virtual worlds and multimedia. In particular, we research the motion in games and virtuals worlds, and the analysis and retrieval of music, 3D GIS, 3D models and scenes, video, and images. Our research topics include motion in games, path planning, crowd simulation, manipulation planning, animation, interaction, feature extraction, matching, recognition, and indexing of different types of media.
Characteristic to this approach is the emphasis on provable and rigorous properties of algorithms. In addition to fundamental modeling, design and analysis of algorithms, experimentation plays a crucial role in showing proof of concepts, and providing experimental verification.

Research plan

In the coming period we plan to extend the area of motion in games further. In particular, we want to extend our work to planning motions of other moving entities, like vehicles, bikes, and animals. Our focus will lie even more on computing visually convincing motion, which we plan to achieve by further multi-disciplinary collaboration with researchers from psychology, physiology, and traffic and transport studies. Also, we plan to extend the group expertise on physical simulation as this is another form of motion in games and virtual world with an important relation to our current work.
Utrecht University now forms a clear national center for research in game technology. We also have many international contacts and collaborations. In the coming period we have the ambition to establish the center for Advanced Gaming and Simulation (AGS) even more as one of the most important international research centers in this area. Game technology will remain an important research area with deep underlying fundamental problems, challenging technologies, and innovative applications, both in entertainment and training. The importance is expected to grow further over the coming period. The group will continue to be active in this domain and extend its expertise, collaborations and scientific, economic and societal impact even further.
Our research on multimedia technology will focus on the big challenge in multimedia research for the next ten years, namely the handling of multimedia information that is perceptually and semantically relevant in a way that is both guaranteed effective, robust, and efficient. To achieve this, two types of problems must be solved. (1) Multimedia processing methods must be made perceptually and cognitively relevant. (2) Multimedia techniques must provide guaranteed performance in terms of effectiveness and efficiency.
In order to bring multimedia science a significant step further, new algorithms must be designed that provide solutions to these unsolved problems. Our approach is to base solutions on geometric pattern recognition and geometric algorithmics.

We plan to put more emphasis on the interrelationship of path planning, crowd simulation, manipulation, animation and interaction. Also, we plan to investigate motion problems involving other entities than characters, such as for example vehicles and animals. The aim is to consider algorithmic robustness of similarity measures, invariance, and efficiency. A new direction of research will be the application in mobile interaction.

UvA-FNWI-CSA B3
Computer Systems Architecture Group
Informatics Institute, Faculty of Sciences, University of Amsterdam

Dr. Clemens Grelck, Prof. Dr. Chris Jesshope, Dr. Andy Pimentel

Background
The predictable but quite unprepared for shift from frequency scaling to concurrency in computer systems architecture has played into the strengths in the computer systems architecture group. We have extensive expertise both in concurrency management for systems-on-chip (SoC) design and in the design and use of high-performance computing (HPC) systems. This shift now puts concurrency firmly into mainstream computing and some of our recent research efforts have supported this direction (http://www.apple-core.info). The major motivation for this seismic shift in computer architecture is power dissipation and we see power management as a key motivator for our research over the coming five-year period. It will span a broad range of computing platforms, from SoC to HPC, although we anticipate a convergence of techniques and methodology used to optimize design across this range.

Research plan
In all areas we see a shift in the provisioning of computing away from a focus on achieving theoretical peak performance and more towards achieving optimal efficiency while a resource is being used. This in turn leads to a requirement for dynamic provisioning. Our strategy will be to match the specific requirements of the heterogeneous components of a computation to processing resources in order to minimize energy consumption. This in turn requires the characterization of both computations and targets for non-functional capabilities and requirements. A major emphasis here will be on maintaining spatial locality. We will therefore adopt appropriate concurrency models (with formal contacts) that expose these requirements and our approach will embrace both static and dynamic optimizations.

In the embedded system domain, one of the most important challenges is the development of design tools that allow for fully automated system-level synthesis and programming of multi-processor SoC architectures. Our joint work with Leiden University in developing the Daedalus system-level MP-SoC synthesis framework offers a fully integrated tool-flow in which the various steps of system-level MP-SoC synthesis, such as parallelization, design space exploration, hardware-software partitioning, application mapping and system prototyping, are highly automated. Future research will extend this technology in two main directions: 1) introducing technology awareness at the system level to e.g. improve the performance and energy-efficiency of the resulting designs, and 2) introducing adaptivity to handle changing workloads by dynamically (re-)mapping applications or even reconfiguring system components.

Our work on fine-grain threaded architectures with data-driven scheduling using the SVP concurrency model will continue but we will also explore software implementations of SVP on other emerging multi-core architectures such as Niagara and Intel's SCC. This will allow us to explore multi-grain architecture and develop an infrastructure to support such an approach. One of the major directions in this work will be the development of a coherent set of operating system services that support space sharing in these heterogeneous environments and yet provide a secure operating environment that can be scaled from chip-level micro-grids to globally distributed Grids.

One of the major challenges, especially in mainstream computing, will be in making these systems programmable without specialized concurrency knowledge and we have designed programming language support to express parallel computations and systems at a very high level of abstraction and developed compilation technologies that effectively map the abstract descriptions to concurrent computing environments. We have international collaborations developing the functional, data parallel language SAC (Single Assignment C) and the asynchronous co-ordination language S-Net. Our long-term vision is in the direction of a compilation infrastructure that automatically adapts running programs derived from high-level specifications to a heterogeneous and dynamically varying execution environment based on continuous reflection of execution parameters.

Intelligent Autonomous Systems

Informatics Institute, Faculty of Sciences, University of Amsterdam

Dr. Ir. L. Dorst, Dr. V. Evers (since 1-1-2010), Prof. Dr. D.M. Gavrila, Prof. Dr. Ir. F.C.A. Groen,
Dr. Ir. B.J.A. Kröse, Dr. A. Visser, Dr. S. Whiteson

Background

Autonomous systems can interact intelligently and naturally with a human-inhabited environment. They operate in a perception-action loop and represent the next phase of ICT: given the advent of inexpensive sensors and actuators, combined with powerful processors and high-bandwidth communication, distributed autonomous systems will soon be widely applied for smart services in health, public safety, well-being, security, comfort and intelligent transportation. Therefore, our main scientific challenge is the development of distributed autonomous systems for intelligent environments.

Through the years, the strategy of the IAS group has been to focus on the strong interrelation between the two critical aspects of real-world autonomous systems: perception and action. About a decade ago, the IAS group made a strategic shift from single-agent to multi-agent robotic systems (the group then introduced robot soccer in the Netherlands). This has led to renowned research on distributed perception, distributed learning and decision making. The next challenge is the collaboration of teams of artificial and human agents that take part in real world applications.

Research plan

The focus of the IAS group is currently on creating 'autonomous systems for intelligent environments' that are capable of interactions with humans. As these component technologies mature, future efforts will concentrate methodologies and theoretical frameworks to synthesize them into complete and robust autonomous systems.

A fundamental problem is to develop *cognitive sensors* that enable machines to interact intelligently and naturally with a human-inhabited environment. This requires the modeling of this human-inhabited environment, typically from distributed or moving sensor systems, to build a map or a 3D model of the environment and track the moving objects through that environment recognizing their activity and interacting with them in a natural way. We will continue to address these issues in collaborations with HvA, TNO and Daimler.

Completing the perception-action loop requires mechanisms for *intelligent decision making*. Agents must be able to select a series of actions that efficiently achieve long-term goals, despite uncertain and dynamic environments. To this end, we are developing state-of-the-art methods for decision-theoretic planning and learning: for single agents, methods that learn efficiently from scarce data; for multiple agents, methods that allow teams of agents to coordinate their behavior despite noisy sensors and limited communication. In the future, we aim to augment these methods with functionality critical for successful real-world application. Specifically, we will extend single-agent methods to work robustly in multi-task settings; we will extend multi-agent methods to exploit simplifying problem regularities to reduce computational costs. To demonstrate the advantages of these approaches, we will apply them in settings such as robot control, gaming, and information retrieval.

Our research includes sensor systems that model *human-inhabited dynamic environments*, in order to understand observed human behavior, and act on it. We need to understand human responses and attitudes toward intelligent (semi-)autonomous multi-agent systems operating in such an ambient manner. We are developing a user lab to perform interaction experiments and tests on the needs and format of effective information exchange for human/machine collaboration. We will continue our funded activities in automated aggression detection; activity monitoring for the elderly with sensor networks; and distributed information acquisition and decision-making for environmental management (with Thales).

The group works on intelligent (semi-)autonomous systems in a human-inhabited dynamic environment. Research topics are model-based reconstruction, distributed learning, and human/agent interaction. Application areas are safety surveillance, intelligent cars, home automation and smart environments. The group aims at building demonstrators that operate in realistic, real-world environments.

Intelligent Sensory information systems (ISIS)

Informatics Institute, Faculty of Sciences, University of Amsterdam

Dr. J.M. Geusebroek, Dr. T. Gevers, Prof. Dr. M.L. Kersten, Dr. D.C. Koelma,
Prof. Dr. Ir. A.W.M. Smeulders, Dr. M. Worring

Background

The ISIS management style encourages an informal but competitive environment for the group, in which teamwork is often considered more important than individual success. Many of our successes can be attributed to collaboration among ISIS-members. PhD students with proven qualities are encouraged to conduct their own research.
The prime scientific target is to understand the content of images by learning from large repositories. The research area is that of semantic computer vision, cognitive vision, and interactive visualization of large picture datasets. These are connected in the sense that human cognition determines what we perceive in images and this is the basis of interaction and semantics.

One research line is *semantic computer vision*, directed to the understanding of arbitrary images by discovery of concepts in the image. We strive towards a very large thesaurus of concepts – currently, 700 concepts (sea, outdoor, person, bicycle, sunny etc.) – detected by one carefully designed algorithm machine-learned from example collections of realistic images. The other research line is cognitive vision, where we investigate how computer vision compares with human vision. And there is a third research trend in the *interactive visualization* of large picture datasets. The three strands are connected in the sense that human cognition determines what we perceive in images and this is the basis of interaction and semantics.

Research plan

Since the complexity of content based analysis technology is ever increasing, while at the same time the size of multimedia collections to be handled is growing significantly, we closely collaborate with the CWI database group. Our focus is on bridging the gap between distributed database architectures and the need to scale algorithms for multimedia information retrieval beyond a few thousand elements. The topic has appeared to be scientifically challenging, causing progress of integrated media-database mining solutions to be slow. Furthermore, large-scale concept learning yields many opportunities for rigorous high-performance computing, for which we maintain a strategic collaboration with the VU computer systems group of Prof. Henri Bal. Here, the focus is on high-performance distributed multimedia computing, and we are making steady progress. Another important aspect in semantic computer vision is statistical machine learning from large-scale data, which has shown to be very effective in computer vision over the past five years. Additional investment in rigorous machine learning is an important asset for the future. Our developing strand into cognitive vision is supported by a strong collaboration with the group of Prof. Victor Lamme (UvA psychology, ERC Advanced Researcher) and of dr. Frans Cornelissen (Groningen brain cognition). These two strategic collaborations study visual perception. Another strategic opportunity is the collaboration with media sciences of humanities within the UvA. Collaboration between science and humanities is often hard to achieve and the impact in media sciences develops slowly.

The research goal on longer term is to continue the strive to a recognition or even cognition of realistic scenes. And to that end: to align - if possible - these algorithms with the working of the human brain; to understand the details of color in computer vision; and to understand ways to visualize large datasets.

System and Network Engineering (SNE)
Informatics Institute, Faculty of Sciences, University of Amsterdam

Dr. M. Cristae, Dr. Yuri Demchenko, Dr. P. Grosso, Prof. Dr. Ir. C.T.A.M. de Laat,
Prof. R. Meijer, Dr. A. Taal, Dr. Z. Zhao

Background
The System and Network Engineering group (SNE) innovates and researches experimental methods for e-Science that use advanced and new ICT infrastructures. The group focuses its research on emerging new local and wide area optical networks and the associated models, systems and protocols. The group builds tools and proof of concept applications that 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. In the last eight years we developed models for hybrid (multi technological) networks for NREN's. We made a tool to check and compare different Internet transport protocols on trans Atlantic optical connections. In recent years we have started studies of the applicability of semantic web methods and tools to describe multi technology interconnected networks and the algorithms to find working paths in them. We selected data intensive e-Science applications to study how data processes can be mapped on a cluster of computers connected by photonic networks. We worked with the Dutch NREN to study grid application controlled dynamically changeable topology of the national photonic interconnect network. We researched the use of programmable network components for sensor grids. We study security aspects of networks and distributed data processing. In particular, we work on resource usage authorization schemes and privacy when processing data in the grid.

Research plan
In the next four years we will study GN3, Geysers, DAS4, NOVI, gigaport, cinegrid. Networking technology has evolved much faster than computing and storage; the interconnecting network performance should no longer be the limiting factor in the development of applications. New photonic technology combined with the availability of dark fibre has led to a model in which end users and their applications get direct access to high capacity wave- lengths in fibres. This led to the new paradigm of hybrid networking, a combination of packet and circuit switched functionality. SNE-group at the University of Amsterdam together with SURFnet is among the leading institutes in this paradigm shift. Major parts of the architectures, procedures, programming models and protocols in use were created in the timeframe before this network paradigm shift. Events like iGrid2002 in Amsterdam, iGrid2005 in San Diego, annual GLIF workshops, the annual SuperComputing conference and projects as GigaPort-NG, CineGrid and OptIPuter have demonstrated a completely new wealth of applications if effort is spent on carefully engineering all the systems involved.
System and Network Engineering aims at designing systems, protocols, and grid middleware to empower applications to optimally allocate and use the new emerging local and wide-area network photonic infrastructures. In this way applications can become location and distance independent except for the unavoidable limit of the speed of light. To achieve this goal, the group carefully delves into the factors that inhibit the optimal use of the emerging new local and wide are photonic networks, builds prototypes of the new technologies and embeds them in international testbeds utilizing the network resources and international connectivity present in Amsterdam. SNE develops middleware and proof of concept applications as necessary; to understand the behavior of the architecture, SNE creates simulation models to study the scaling of the technologies. Very important new topics in such dynamic environments are infrastructure security and data integrity and the carbon footprint of infrastructure.

We will study the applicability of semantic web methods and tools to describe multi technology interconnected networks. Other research topics are the mapping of data processes on a cluster of computers connected by photonic networks and the security aspects of networks and distributed data processing.

VU-EW-CS B1,2,4

Computer Systems

Division of Mathematics and Computer Science, Faculty of Sciences, Vrije Universiteit Amsterdam

Prof. Dr. Ir. H.E. Bal, Dr. H. Bos, Dr. B. Crispo, Dr.-ing. Habil. T. Kielmann, Dr. G. Pierre, Dr. M. Rieback, Dr. F.J. Seinstra, Prof. Dr. Ir. M. van Steen, Prof. Dr. A.S. Tanenbaum, Dr. S. Voulgaris

Background

The Computer Systems group of the VU (VU-CS) does fundamental and experimental research on Secure and Reliable Computer Systems, High-Performance Distributed Computing, and Distributed Computer Systems. The group studies the underlying systems software, develops realistic applications by collaborating with domain experts in different fields, makes software prototypes and distributions, and develops the infrastructure that is needed to do state-of-the-art experiments (e.g., DAS, sensor networks). Many of its software systems (e.g., Minix, Ibis, JavaGAT, Argos, Streamline, Tribler) are used extensively by other researchers. The group has its own educational program (the topmaster Parallel and Distributed Computing Systems) to train future PhD students.

Research plan

Our work on Secure and Reliable Computer Systems will focus on improving the reliability, availability, and security of system software, especially operating systems. Among other things, it will focus on how to transparently recover from crashes of stateful components, such as file systems, as well as how to update major portions of the operating system on the fly, without having to reboot. Also we are investigating how multicore chips can be best exploited to enhance reliability. We are looking at how to allow grid users to not only select machines based on CPU speed and RAM size but also on their security properties. Other work is related to how to permit third-party applications to run on mobile devices with enhanced security and trust. Additional work relates to improving the security of RFID systems. In particular, we are trying to devise a system for monitoring and countering RFID surveillance and also provide tools for doing security audits of RFID systems before and after they are deployed.

We will continue our work on High-Performance Distributed Computing (HPDC) in several directions. We will study additional high-level programming frameworks, for applications like distributed model checking, image processing, and graph-based bioinformatics algorithms and we will look at data-intensive applications, such as scalable reasoning over semantic web data and e-Science applications. We will study how to exploit accelerators like Graphical Processing Units (GPUs), using DAS-4 as hybrid testbed. Other work is how to use Ibis and JavaGAT to integrate cloud computing and mobile computing (e.g. with smart phones). We will investigate providing scientific compute clusters (like DAS-4) as service platforms, mostly as Infrastructure as a Service (IaaS) and Platform as a Service (PaaS). This will provide 'virtual clusters' based on tightly-coupled cluster hardware with flexible allocation and accounting policies. In addition, we will expand our research in system security – partly by offloading security checks to distributed computer resources, and partly by analysis of malware and goodware.

Our efforts on Distributed Computer Systems will concentrate on two related themes. First, research continues in the area of managing large-scale Internet-based distributed systems, including grid and cloud computing environments. Systems management issues include allocation of resources as well as easing the accessibility of compute clouds as back ends. This work will be done in collaboration with the HPDC group. Second, we will be putting significant efforts in large-scale wireless distributed systems, with an initial emphasis on systems for ubiquitous computing. Part of this research will involve applying techniques from computational intelligence (Eiben's group at the VU). Much of this work will be carried out with other groups (Chess, Philips Research, Social Sciences).

As longer-term plans we want to study how the software we developed for GPUs, clouds, etc. can also be applied for intelligent mapping and (multi-objective) self-optimization, for example to reduce energy consumption of distributed computing systems like DAS-4. Also, we aim to develop fully self-managing large-scale distributed systems, be they wired or wireless.