CWI

Centrum voor Wiskunde en Informatica

ANNUAL REPORT '97

Kruislaan 413, 1098 SJ Amsterdam, the Netherlands
P.O.Box 94079, 1090 GB Amsterdam, the Netherlands

CWI is the National Research Institute for Mathematics and Computer Science. CWI is administered by the Stichting Mathematisch Centrum (SMC), the Dutch foundation for promotion of mathematics and computer science and their applications. SMC is sponsored by the Netherlands Organization for Scientific Research (NWO). CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics. CWI participates in the

Telematics Institute.

General Director

G. van Oortmerssen

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
Kruislaan 413, 1098 SJ Amsterdam, The Netherlands
Telephone +31 20 592 9333
Telefax +31 20 592 4199
URL: http://www.cwi.nl

CONTENTS

This Annual Report is complementary to the Jaarverslag CWI (in Dutch), which concentrates on the Institute's management including financial and social aspects. A complete overview of CWI's research activities is also available.
This Annual Report and the other reports can be ordered at
Mrs. D.C.M. Amende; Phone +31 20 592 4128, e-mail: Thea.Amende@cwi.nl

OVERVIEW

As judges of wine would say: 1997 was for CWI an excellent year, with a rich harvest of high quality. This was not only true for fundamental research – the institute's basic material –, but also for the second part of CWI's mission: dissemination of achieved results in the society. Success in this latter field is not at all self-evident, as a recent evaluation of the EU Framework Programmes shows: the conversion of basic knowledge into competitive products has still not satisfactorily got off the ground, all well-meant attempts notwithstanding. Viewed in that light CWI has extra reason to look back onto the report year with some satisfaction. On the basis of a healthy financial position CWI not only produced its usual high output of scientific results, but also reinforced its research acquired externally in competition – with 30% of the total budget a share already considered rather high, given CWI's profile as a fundamental research institute. As a rule, trade & industry is strongly involved in this external research, which is expected to further increase by CWI's participation in the newly founded national Telematics Institute. New projects in genetic algorithms and neural networks, and in mathematical finance show CWI's alertness to go along with rapid developments in the research world, enabled by its flexible organization. Added to its strong, ever renewing ties with the academic world, this proves CWI's ability to combine fundamental research with a business-like approach, and also shows that societal involvement can be a very stimulating source of studying new fundamental problems. The way in which CWI has dealt with its challenges in the report year are detailed below.

Two important developments marked the year 1997 for CWI. First, the planned transition to a theme-oriented organization of its research was effectuated. Now CWI's research themes are grouped into four clusters, which have replaced the previous six discipline-oriented departments. In the new structure, which met the expectations, interdisciplinary collaboration finds a natural place in the research plans. Two new pilot themes were added during the year: *Evolutionary Computation and Applied Algorithmics* (theme leader: J.A. La Poutré), and *Mathematics of Finance* (theme leader: J.A. Schumacher). Both themes are excellent examples of research where input from several disciplines in mathematics and computer science is indispensable. Another pilot theme, *Quantum Computing and Advanced Systems Research*, developed so favourably that it passed the trial phase already after one year. On the other hand the activities of the *Dynamical Systems Laboratory*, which stood apart from the theme structure, came to an end. During the five years of its existence pioneering results have been achieved in the study of bifurcation phenomena in systems of ordinary differential equations. The research led, among other things, to the construction of the widely used software package CONTENT for computer-assisted bifurcation analysis.

The new research structure gives the possibility to appoint a researcher for a period of several years as CWI Fellow, enabling him to concentrate completely on the research of his liking. The first researcher to become a CWI Fellow was P.J. van der Houwen, who previously headed the department of Numerical Mathematics.

The second important event was the governmental decision to create – after extensive pre-selection – four Technological Top Institutes, in which industry and research world join hands. CWI participates in the research programme of the Telematics Institute, led by the Twente-based Telematics Research Centre, and with Twente University, TNO (the Dutch Organization for Applied Research), and Delft University of Technology as other participating knowledge institutions. In addition, seventeen companies are involved in the Telematics Institute, including IBM, KPN, Lucent Technologies, ING and Rabofacet as participants. In the report year CWI submitted several research proposals, three of which have been granted so far.

These developments can be seen as consequences of CWI's policy, initiated in the early 1990s, to focus on frontier research to be carried out in the framework of multidisciplinary themes derived from societal demands. Optimal knowledge transfer is achieved by a flexible organization of the research in an ever expanding network of cooperation agreements with governmental institutions, universities and companies.

Several other activities reported on below bear

New numerical methods for CFD applied to ship hull design
Computational Fluid Dynamics (CFD) is applied in a wide variety of industrial settings, for example computation of water flow around ship hulls or air flow around airplane wings. CWI has considerable experience in the research and development of numerical techniques to be used in CFD methods. Present CWI research in this field is focused on the computation of fluid flows for all kinds of industrial applications. The fluids can be gases, liquids, or combinations of these (multi-phase flows). Research topics are: advanced discretization methods for systems of nonlinear conservation laws, multigrid and sparse-grid solution methods, local grid adaptation and distributed computing. Emphasis lies on the development of sparse-grid algorithms for 3D flow problems. A recent 4-year joint project with the Dutch Maritime Research Institute MARIN concerns the design of ship hulls using more realistic conditions. In particular CWI develops improved Navier-Stokes methods for ship hydrodynamics, in order to increase the reliability and efficiency of MARIN's viscous CFD technology, which is presently restricted to steady, simplified Navier-Stokes equations and the corresponding computational domains to ship sterns and flat water surfaces. An extension will be made to the unsteady, full Navier-Stokes equations, with application to complete ships (bow and stern) and free-surface wave making. The new CFD methods will be applied in flow topology studies around ship hulls (flow separation, re-attachment, vortex formation, etc.) at various sailing conditions (speed, trim, sinkage). Preliminary tests with flows around aircraft wings indicate that sparse-grid algorithms developed at CWI may also be effective when applied to ship hydrodynamics. http://dbs.cwi.nl/cwwwi/owa/cwwwi.print_projects?ID=8

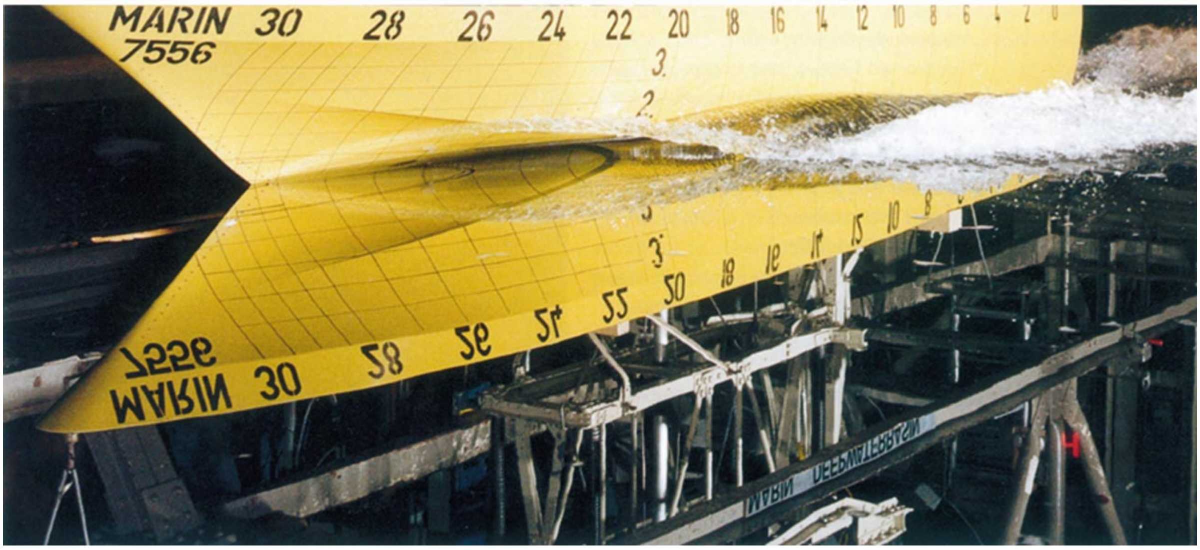

Model test at MARIN on a motor-yacht with breaking free-surface waves. (MARIN Report 54, 1995.)

witness to CWI's endeavour in realizing its mission: carrying out frontier research and transferring the resulting knowledge to society, in particular trade & industry. Of course, for its fulfilment an adequate research staff is required, in quality as well as in quantity. During the report year CWI had some difficulty in appointing qualified researchers in sufficient numbers (other scientific disciplines suffer from similar problems). As a consequence some projects started somewhat later than was anticipated. CWI will keep paying due attention to the recruitment of research staff in the near future.

Relations with Academia
Obviously CWI's development keeps pace with those in the academic world, given the many close ties of old between the two. The relations still evolve dynamically, as is witnessed by an initiative of CWI to come to a more structured cooperation with the Dutch universities. Particular attention was given to the exchange of researchers and cooperation with the Graduate Schools. An agreement was concluded with Eindhoven University of Technology about mutual secondment of researchers. CWI was also involved in three research proposals of Graduate Schools to be financed by NWO's 'depth strategy' budget. Furthermore CWI plays an active role in some nation-wide research networks – structures which enhance the coherence of academic research in a certain field – , for example the Operations Research network.

A major part of CWI's collaboration with academic researchers takes place in projects of the NWO foundations SION and SWON, for computer science and mathematics, respectively. Both foundations and CWI interact about their research programmes. For example, CWI was involved in preparing the National Research Agenda for Computer Science. CWI participates in 18 SION projects, with researchers coming from 9 universities. Since SWON's foundation mid 1996 CWI acquired participation in several projects (13 so far), with 8 universities involved. Here the subject Mathematics of Finance deserves special attention. Being SWON's 1997 Year Theme, research in this field received additional support. CWI fully benefitted from this support, which boosted preparations for an own full-grown research theme in the mathematics of finance. Apart from projects funded by SION and SWON, researchers from universities and CWI collaborate, e.g., in European projects and in several commissions from industry, totalling up to 36 academic partners in 21 projects.

CWI's close ties with Academia are also borne out very well by the high number of Ph.D. degrees received by young CWI researchers at a Dutch university, usually supervised by a professor who also leads a research group at CWI. The report year saw the completion of 13 Ph.D. theses (11 in 1996, 8 in 1995). Several of these qualified young researchers continue their careers with companies (including banks) or governmental institutions, which is one of our most effective forms of knowledge transfer.

Finally, CWI's library aims at closer cooperation between the Dutch academic libraries for mathematics and computer science. To this end it started a few years ago a semi-annual consultation meeting which has grown in importance ever since. Specific plans exist for close cooperation with the University of Amsterdam in developing a digital library for mathematics and computer science.

Difficult Dutch

Several languages, including Dutch and Latin, are far more difficult to automatize than English. CWI researcher Annius Groenink described in his Ph.D. thesis how to deal with some of these difficulties by applying techniques used in mathematically defined languages to the description of natural languages. Automatic translation still shows some substantial gaps, partly due to the technicians' concentration on English, which in a mysterious way is more related to computer languages than Dutch. The difficulties in Dutch (for example the word order) present themselves even stronger in Latin. To fit a language for automation its grammar should be described with such precision that on this basis a computer can automatically parse sentences. Existing methods, while successful for French and English, turn out to have a rather poor performance when applied to Dutch. One cause lies in the difficulties with the construction of an underlying 'tree structure' in Dutch. Contrary to English and computer languages, many Dutch sentences do not have such a tree structure. The solution proposed by Groenink was known already for some time among formal language theory researchers, but was never applied in this context. A part of a sentence, which corresponds to a branching point in the tree, is viewed not as just one string of words within the sentence, but as two or more separate parts. Such extensions of the existing methods provide in principle a description of Dutch which can serve as a basis for automation. The proposed method turns out to work also for a language like Latin. Groenink also proved that his proposed extensions with regard to parsing do not cause an exponential growth in computing time and he made plausible that more realistic implementations of his software, where for example checks on cases and singular/plural are included, also remain efficient.
http://www.cwi.nl/∼avg/

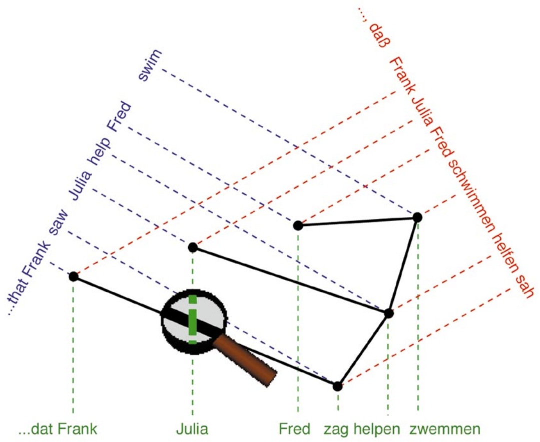

The universal underlying 'tree structure' (thick lines) of the same sentence in Dutch, English and German. For English and German straight lines (the dotted lines) can be drawn from the basic structure to the words. As for the Dutch sentence, the dotted lines will always either cross each other, or cross the thick lines at least twice.

CWI – Academia Relations
Secondment of CWI staff at universities:
15 professors, 4 researchers

Secondment of university staff at CWI:
10 researchers, 7 advisors, 5 graduate students

13 Ph.D. theses by CWI staff in 1997

Cooperation agreements with Graduate Schools:

- Thomas Stieltjes Institute for Mathematics
- Euler Institute for Discrete Mathematics and Applications
- Dutch Graduate School in Logic
- Institute for Programming Technology and Algorithmics
- Dutch Institute of Systems and Control
- School for Information and Knowledge Systems
- Mathematics Research Institute (to be concluded in 1998)

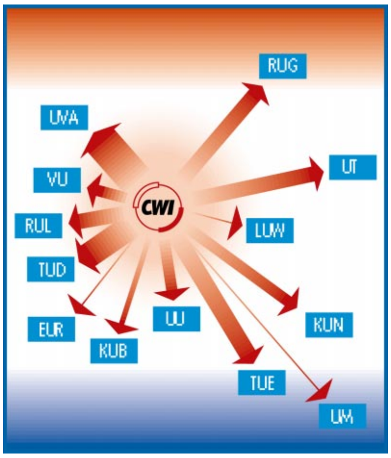

84 projects (18 SION, 13 SWON)

Relations with Trade & Industry
CWI's contacts with Dutch companies go back to the early years of its existence. For example, the Mathematical Centre (CWI's name prior to 1983) pioneered in the building of computers in the fifties. In particular in the current decade CWI made the active approach of industry a deliberate policy, thus aiming at the optimal utilization of its fundamental knowledge acquired over the years. Apart from direct contacts leading to various, sometimes long-term commissions, CWI also proceeded more indirectly. A fine example of the latter is provided by *CWI in Bedrijf*. At this annual event, always on the first Friday in October and first organized in 1992, CWI presents its research potential to the Trade & Industry sector, of which on an average hundred representatives attend. The 1997 event focused on the newly created Telematics Institute, and CWI's participation in it. In addition CWI organized, jointly with the University of Amsterdam and the Industrial and Applied Mathematics foundation, the workshop *The Mathematical Working-floor* in April, where 65 participants heard presentations of several open problems in industrial practice and discussed their possible solution. Finally, CWI was present at the technological knowledge market *Techknowledge* end May in Utrecht, organized by the employers organization VNO-NCW, where the worlds of (academic) research and industry met. Demonstrations of CWI research included the verification of software systems, visualization and parameter estimation.

CWI's continuing participation in projects financed by NWO's Technology Foundation STW, in which industry is committed to an active role, is another proof of the institute's success in linking up its fundamental research with societal demand. The conclusion of 'Parallel codes for circuit analysis and control engineering' in 1997 was followed by a (successful) request for a continuation project 'Parallel software for implicit differential equations'. And the termination of 'Parameter identification and model analysis for non-linear dynamical systems' was amply compensated by the start of no less than three new projects: MOBILECOM (mobile communication networks), Wavelets – analysis of seismic signals, and FASE (animation of facial expressions).

Recently CWI's many contacts with the Trade & Industry sector have extended from traditional 'technical' research areas such as computational fluid dynamics and transport phenomena to sectors like infrastructure and services, in particular financial services. In the latter field CWI, e.g., carries out contract research into the renovation of legacy systems (the SOS Resolver project, including ABN AMRO as a participant) and studies application of (quasi-) Monte Carlo methods to financial derivatives such as options. Infrastructural research at CWI includes railway traffic (an ongoing commission of Dutch Rail to develop tools for the design of a new timetable, advisory work concerning traffic control, and the safety of railway-yards) and the control of traffic flow in motorway networks, such as ring roads around large urban areas (the EU project DACCORD).

The generation of spin-off companies is an increasingly important method for institutes like CWI to convert fundamental knowledge to applications and at the same time create high-level employment. Close links with the mother institute during the initial years form a crucial success factor here. CWI researchers successfully managed to establish spin-off

Specification and analysis of embedded systems
In order to increase the quality of software components, e.g., in telecommunication and embedded systems, CWI studies specification techniques for the unambiguous description of these systems, and analysis techniques to establish that their implementations exhibit the expected functionality. The advent of distributed computer systems and parallel computation created a need for new specification techniques. A promising approach is the use of process algebra in which concurrent processes can be formally described, and of the specification language μCRL (micro Common Representation Language) based on it. CWI made important contributions to both and currently studies a real-time extension of μCRL. Applications include software validation for Philips' future generation of IC's and system compliance of its Digital Video Broadcast source decoder (still under development). Furthermore, CWI tries to increase the efficiency of current symbolic techniques to verify requirements on processes by a fundamental understanding of proof search in simple logical systems (*in casu* propositional logic), and develops proof checking methods in order to establish the correctness of programmed systems 'beyond any reasonable doubt'. A recent application concerns the safety of a VPI (Vital Process Interlocking) system, which is a kind of programmable controller, in operation at the train station of Heerhugowaard. The system has to comply with several safety requirements, which can be formulated as statements in propositional logic. The check on all requirements being satisfied is an NP-complete problem. CWI constructed a workable prover, Heerhugo, which was also successfully applied to consistency checks on two merged databases of automobile components (NedCar and Volvo). http://www.cwi.nl/~jfg/

By using formal methods CWI could prove the correctness of the safety system used at a local railway-yard.

companies in recent years. Some of these, for example Digicash and General Design, now have some dozens of employees. Recently CWI's youngest spin-off, Data Distilleries (operating in the datamining field), received a considerable capital injection from a major Dutch investor, in order to grow into a size considered as necessary for a successful operation at the global scale. During the report year CWI intensively searched for possibilities to create new spin-offs. Experience shows that in the preparatory phase carefulness is more important than speed. CWI now considers allocating special funds to the stimulation of developments in this direction.

Relations with governmental and other institutions
Several large governmental institutions with a research mission can very well play the role of an intermediary between fundamental research and concrete applications (the actual making of a product). Hence, CWI cherishes long-standing good relations with many of these institutions and is continuously searching for extension of these relations. A substantial part of CWI's research in cooperation with such institutions concerns large-scale computing. This research is reviewed in a separate chapter of this Annual Report.

CWI's relations with the ministry of Public Works

.

Recent CWI spin-off companies

Company	*Profile*	*Startup year*	*# Employees*
NLnet	Internet provider	1994	100
General Design	Web design	1994	35
Data Distilleries	Datamining	1995	8
CAN Diensten	Mathematical software	1995	4

(the RWS department) date back to the fifties, when the institute started with advisory work in statistics – a tradition continued up to the present. In the 1980s CWI also received commissions for research in the traffic area, and presently cooperates with RWS in the European project DACCORD. The start in 1997 of the new research theme *Evolutionary Computation and Applied Algorithmics* brought additional cooperation in the field of Neural Vision, in particular the use of neural networks for the visualization and clustering of high-dimensional data, on behalf of exploratory data analysis in a decision making process.

Cooperation with the Dutch organization for applied scientific research TNO (and other parties) takes place in some projects concerning porous media, focusing on the modelling of transport phenomena in the underground as a system of partial differential equations. One application is the removal of organic contaminants from the soil by pumping or air sparging. As for CWI's work on TNO's ozone model LOTOS we refer to this Report's chapter on Large-scale Computing.

Several other large institutions cooperate with CWI: the Royal Netherlands Meteorological Institute KNMI in wavelets, sparse grid methods for time-dependent partial differential equations, and regional 3D long-term ozone models (LOTOS); the Dutch Maritime Research Institute MARIN in improved Navier-Stokes methods in ship hydrodynamics and in wavelets; and the National Institute for Public Health and Environmental Protection RIVM in sparse grid methods and the LOTOS model. Finally, CWI's year-long cooperation with the National Energy Centre ECN in the field of Virtual Reality is still continuing.

International cooperation

CWI continues playing a very active role in the European Research Consortium for Informatics and Mathematics ERCIM, and in several of its international cooperative projects one or more consortium members are involved. An important part of the collaboration within ERCIM proceeds in a number of

Genetic algorithms and neural networks

Genetic algorithms are computation methods based on the principles of biological evolution theory, like 'survival of the fittest' and recombination of species. They thus yield a good collection of solutions for an optimization problem to be solved. Since the algorithm itself implicitly builds up knowledge about good solutions, it is ideally suited to deal with complicated or vague criteria, and with varying or inaccurate data or modelling. Neural networks are adaptive computation structures inspired by biological brain models, consisting of neurons and axons. Because of their capability of (autonomous) learning and extraction from a data set, neural networks are often very well suited to application within decision processes. CWI's basic research in these fields is applied to industrial and technological management situations and to commercial and economic strategy development. Projects include: an evolutionary planner for large-scale air traffic management, enabling a more efficient use of overflight facilities (joint work with NLR); visualization and clustering of high-dimensional data sets by means of data projection as part of a strategic decision process (cooperation with Nijenrode University and the ministry of Public Works); clustering techniques applied to information on sales slips, in order to determine customer behaviour (cooperation with a wholesale house and the University of Leiden); extraction of information on disaster areas from remote sensing data through data clustering and unsupervised (i.e., without additional expert information or ground measurements) classification (cooperation with the University of Amsterdam and the Cold Regions Research and Engineering Laboratory, USA); on-line management related to decision problems in advanced information technology, including computer networks and multimedia systems (cooperation with the NWO foundation SION and Philips Research Laboratories).

H.A. Lauwerier (1923 – 1997)

Professor Hans Lauwerier passed away on November 21, 1997, at the age of 74. When he joined in 1956 the Mathematical Centre – CWI's name prior to 1983 – he had been stationed there already for two years by his employer Shell, to work on problems connected with the North Sea. These studies originated in the large flood disaster of February 1953, with almost two thousand people killed, and were part of the following governmental Delta Plan to create a better protection of the Dutch coastal area. Lauwerier's work in this field considerably contributed to the fame of the Mathematical Centre in The Netherlands and abroad. From 1959 until his retirement in 1988 Lauwerier headed CWI's department of Applied Mathematics and held at the same time a professorship at the University of Amsterdam. During the 1970s he introduced biomathematical research in The Netherlands, which since then has grown into a flourishing mathematical discipline. In his last years at CWI Lauwerier became fascinated by the new (actually largely rediscovered) mathematics of Julia and Mandelbrot sets, and fractals. This led to a series of widely appreciated popular-scientific books and articles on these topics, on symmetry, and on order and chaos. His book on fractals appeared in several languages. All these books and articles were accompanied by simple computer programs, by which the reader could acquire hands-on experience. Until the very last Lauwerier remained active in writing popular accounts of mathematical subjects, showing their relevance as well as their beauty.

working groups. CWI coordinates three out of the thirteen current working groups: Constraints (K.R. Apt), Control and System Theory (J.H. van Schuppen), and Parallel Processing Network (H.J.J. te Riele). In several of the other working groups CWI researchers also play an active role. The same is true for ERCIM's Digital Libraries Initiative, for instance through the DELOS working group. In the report year CWI has decided to start an active search for bilateral contacts with ERCIM institutes, whose research programme closely links up with and is complementary to CWI research in the same field. The German ERCIM partner GMD was the first with whom such a contact was established, and cooperation started on themes including embedded systems, environmental modelling, and datamining. The cooperation was prepared in two workshops held in Bonn and Amsterdam. Similar consultations with the French ERCIM partner INRIA are scheduled for early 1998. (Indeed, CWI's contacts with both GMD and INRIA stem from long before the foundation of ERCIM in 1989, when the consortium started with these three members.) Later in 1998 bilateral collaboration with the Hungarian ERCIM partner SZTAKI will be considered. This initiative is also based on preceeding scientific contacts, for example in the framework of the Dutch-Hungarian cooperation agreement concluded by the ministry of Education, Culture and Science.

As for CWI's participation in European projects, the slight decrease in 1996 – mainly due to the increasing lack of tuning between European research programmes and CWI's research profile – did not develop further. On the contrary, a new increase is expected for the coming year. Even with its present participation in eleven research projects and the same number of HCM/TMR projects, CWI still scores high, taking into consideration the institute's size. Notwithstanding its less favourable positioning with respect to present European programmes, CWI still tries very hard to play a full-grown role here, as is clearly shown by the fact that end 1997 no less than thirteen proposals were under consideration in Brussels.

CWI has played a historic role in the development of Internet. For many years the institute acted as node for the connections between the USA and Europe. No wonder that with the emergence of the World Wide Web CWI soon acquired a prominent position here. The institute plays an important role in some working groups of the world-wide W3C consortium, led by MIT, INRIA and Keio University (Japan), which was founded in order to stimulate and coordinate the further development of the Web. (CWI's research in one working group is highlighted in this Report.) Meanwhile the W3C consortium has decided to open in some countries a separate office, which acts as a central point for all matters concerning the Web. The Dutch office is located at CWI.

Last, but not least, right before Christmas CWI co-organized a video conference on the theme 'Mathematics and Water', an initiative of the European Mathematical Society. From three locations: Madrid, Venice and Amsterdam, several experts could listen to each other's presentations and participate in the discussions. The conference also received input from society (government, media, environmental organizations).

Looking back over the past year, a feeling of satisfaction is certainly justified. The institute's financial position is sound and its strong embedding in society guarantees that the performed frontier research finds its way through various channels to applications. However, to remain at the top is often still more difficult than to get there. Maintaining the high standards of our research not only requires researchers of the highest quality, but they should also be sufficient in number. We are convinced that mathematics and computer science will remain such a challenge for young talent that they will attract many, and that CWI is an excellent place for an optimal development of this talent. This conviction we will propagate with all means at our disposal.

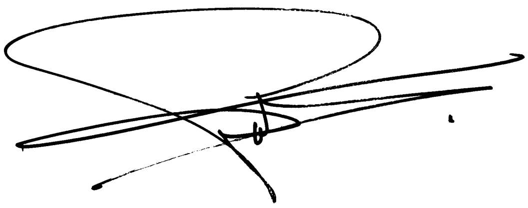

G. van Oortmerssen

General Director

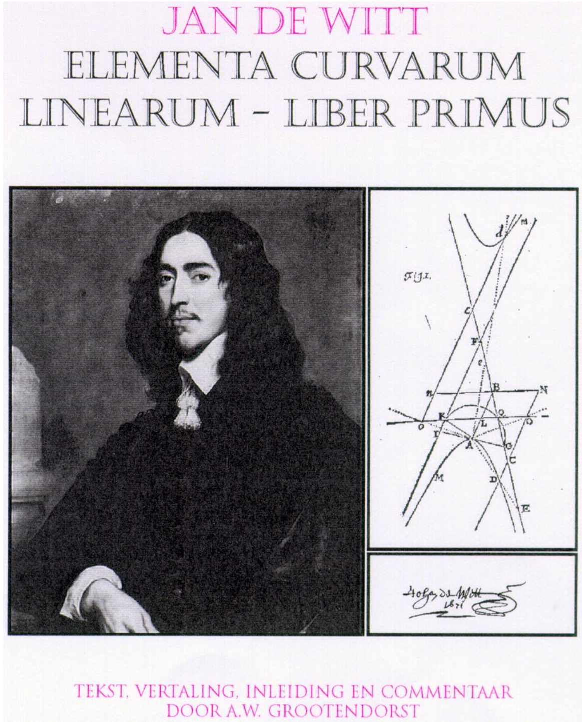

CWI published the first translation from Latin into Dutch (by A.W. Grootendorst) of a work on Geometry by the 17th century statesman (and for almost twenty years de facto *ruler of the Dutch Republic) and mathematician Jan de Witt (1625 – 1672). De Witt is also generally considered as the founder of life-insurance mathematics. (Painting by Adriaen Hanneman, Museum Boijmans Van Beuningen, Rotterdam.)*

ORGANIZATION

CWI (Centre for Mathematics and Computer Science) is the research institute of the Foundation Mathematical Centre (SMC), which was founded on 11th February 1946. SMC is funded mainly by the Netherlands Organization for Scientific Research (NWO). The organizational structure of CWI is shown on the opposite page. CWI's mission is twofold:

- to perform frontier research in mathematics and computer science;
- to transfer new knowledge in these fields to society in general, and trade & industry in particular.

CWI's research is carried out in theme-oriented research groups, which are grouped in four clusters.

Cluster - Theme	Cluster leader Theme leader
Probability, Networks and Algorithms	**O.J. Boxma**
- Networks and Logic – Optimization & Programming	A.H.M. Gerards
- Traffic and Communication – Performance & Control	J.H. van Schuppen
- Stochastics	M.S. Keane
- Signals and Images (pilot)	H.J.A.M. Heijmans
Software Engineering	**J.W. de Bakker**
- Interactive Software Development and Renovation	P. Klint
- Specification and Analysis of Embedded Systems	J.F. Groote
- Coordination Languages	J.J.M.M. Rutten
- Evolutionary Computation and Applied Algorithmics (pilot)	J.A. La Poutré
Modelling, Analysis and Simulation	**C.J. van Duijn**
- Environmental Modelling and Porous Media Research	J.G. Verwer
- Industrial Processes	P.W. Hemker
- Mathematics of Finance (pilot)	J.M. Schumacher
Information Systems	**M.L. Kersten**
- Data Mining and Knowledge Discovery	A.P.J.M. Siebes
- Multimedia and Human-Computer Interaction	D.C.A. Bulterman
- Interactive Information Engineering	P.J.W. ten Hagen
- Quantum Computing and Advanced Systems Research	P.M.B. Vitányi

The researchers are supported by state-of-the-art computer facilities, e.g., an internal high bandwidth glass fibre ATM network, and a well-equipped library of national importance. Hence, CWI is well-prepared to handle the dynamic and interdisciplinary demands of present day research.

SMC is governed by a Governing Board, chaired by L.A.A.M. Coolen (director KPN Research). The other Board members are: P.M.G. Apers (University of Twente, chairman SION Board), J.H.A. de Smit (University of Twente, chairman SWON Board), K.M. van Hee (Eindhoven University of Technology, director Bakkenist Management Consultants), and H.A. van der Vorst (University of Utrecht). The chairpersons of the SION and SWON Boards are members *ex officio*, in order to keep SMC's research policy well-geared to the programmes of SION and SWON – the NWO foundations for research at Dutch universities in computer science and mathematics, respectively.

SMC's management is delegated to the General Director, who is supported by a Management Team consisting of CWI's four research cluster leaders and the controller. The set-up of an Advisory Council, with representatives occupying leading positions in science and industry in The Netherlands, will be completed in 1998.

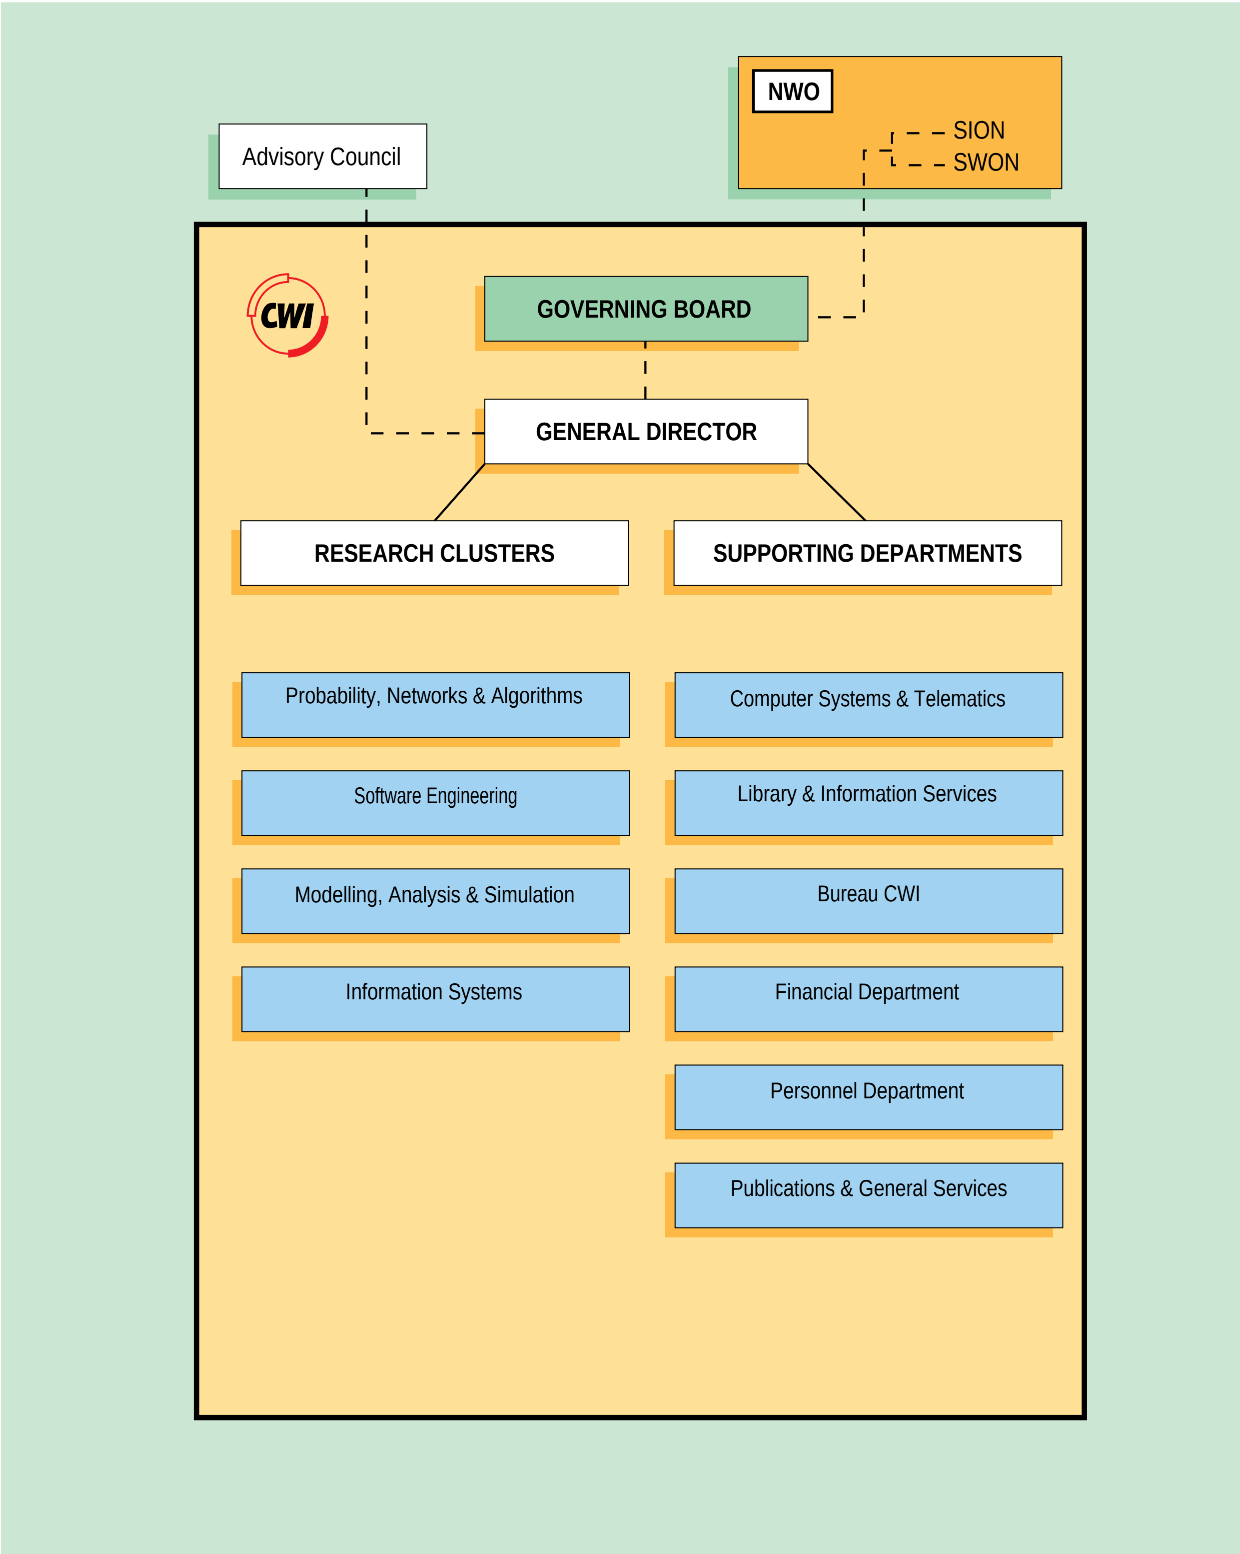

CWI organizational chart.

The Use of Wavelets in Seismics and Geophysics

Research Project	:	Wavelets – Analysis of Seismic Signals
Project leader	:	N.M. Temme
E-mail	:	Nico.Temme@cwi.nl
URL	:	http://www.cwi.nl/cwi/projects/wavelets.html

Introduction

Wavelets are wave patterns of small size, that can be used to analyze rapid changes in a signal, sharp contrasts in an image, and structures at different scales. The wavelet method, developed in the eighties and now provided with a sound mathematical basis, has become a powerful tool in the field of data compression, noise suppression and processing of images and signals. At CWI wavelets are studied in a project financed by the Technology Foundation STW which focuses on the application of wavelets in seismology and geophysics.

Research on wavelets and its applications in The Netherlands takes place for instance at KNMI (structure of the clouds), Delft University of Technology (soil structures, for oil exploration), Groningen University (image processing in medicine) and CWI (analysis of geophysical/seismic signals). CWI's research is focused on the development and use of directional time-scale analysis methods. The combination of the Radon transform (also called X-ray transform – a well-known technique in computerized tomography and seismic exploration geophysics for reconstructing a 3D object from a number of cross-sections) and the wavelet transform is a key issue here. A fast algorithm for this 'Wavelet X-ray Transform' is also under development. The reason for applying this method is to remove unwanted signal parts from measurements in geophysics. Secondly, research is conducted jointly with KNMI on the problems of polarization in seismic signals, where an estimation of the spectrum is made by preprocessing the signal data.

Problem formulation

The CWI wavelets group is working on mathematical wavelet techniques in order to filter two-dimensional seismic data sets known as 'shot records'. These images contain information about the subsurface of the earth such as the localization of geological interfaces. For the oil industry, information about the deep underground is pertinent for the exploration of oil reservoirs. Unfortunately, the seismic images do not only contain information about the deep subsurface, but they also contain irrelevant parts due to the near subsurface, for example contributions from waves coming directly from the explosion (that is used to generate the signals) to the detection point.

Such irrelevant parts disturb the extraction of the useful information, especially if they have a strong amplitude compared to useful parts of the image. The aim of filtering is to get rid of these disturbing events prior to extraction of the relevant parameters. The main idea behind the wavelet techniques studied at CWI is that relevant and disturbing parts in an image can be separated based on distinction in position (or time), scale (or frequency) and direction. The classical Fourier method to analyze signals or images in terms of sine- and cosine-waves, may fail in detecting certain features such as sharp edges or multi-scale structures. Although Fourier (and Radon) techniques have become a whole industry in seismic data processing, the Fourier analysis of shot records gives rise to frequency-wavenumber filtering techniques which are only global in nature. Essential local features might be missed. The Radon transform (known as 'slant-stack' in this application area) helps to overcome this problem. Both techniques are used to distinguish directions. Wavelet techniques, combined with the Radon (or X-ray) method, are expected to be successful in cases where Fourier and Radon techniques have failed.

Activities and results at CWI in 1997

Wavelet techniques are tailor-made to separate parts which are distinct in position-scale (or time-frequen-

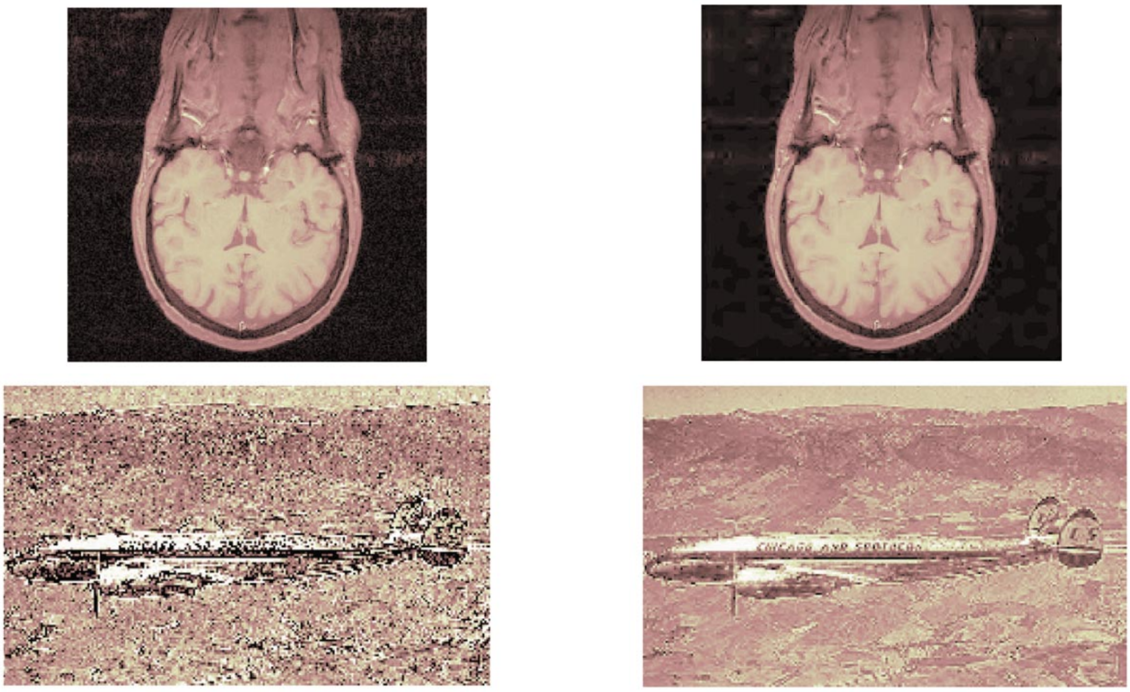

A compressed image of a brain-slice (top right) is stored six times as efficient as the original (top left). The resolution is retained. This is achieved by ignoring negligible wavelet coefficients. A corrupted image of a plane (bottom left) is cleared (bottom right). First a wavelet decomposition of the image is made. Then the coefficients corresponding to the noise part are muted. The final reconstruction of the image is based on the modified wavelet coefficients. In both cases MATLAB Graphical User Interface tools were used.

cy). In order to establish separation based on direction, the slant-stack technique mentioned above is being used extensively in exploration geophysics.

At CWI, wavelet and X-ray techniques are combined in the so-called wavelet X-ray transform. From a mathematical analysis point of view, this integral transformation has the same pleasant properties as the wavelet transform. A rigid mathematical formulation of the transform, its basic properties and its discretization have been obtained.

Moreover, the discrete version of the wavelet X-ray transform allows for fast implementation. The computer code has been established in a Matlab environment. The code has been applied to synthetic and real seismic datasets, and the first results are promising.

Time-frequency analysis

A second research topic is connected with seismograms: registrations of the ground motion at a certain point at the earth's surface. This motion can be caused by an explosion, a passing vehicle or probably an earthquake. There exist many kinds of earthquakes, which differ from each other by the power and the position of its source. To analyze the different kinds of seismic events, we have to study the waves coming from the source and arriving at the Earth's surface at different arrival times. In a seismogram these waves appear as so-called phases. These phases can be seen as the ground motion due to the arrival of a particular wave coming from the source. All these phases are measured during a short time period. An aim of research is to find the time periods in a seismogram, where these phases appear. Furthermore we want to analyze these phases accurately.

A powerful common method for analyzing these segments is spectral estimation: we analyze the frequency components of a segment contained in the spectrum of its Fourier transform. Within the field of spectral estimation many techniques have already been developed, which work well for the analysis of long records of stationary processes. However, in analyzing the spectrum of a short segment in a longer record, we have to deal with the uncertainty introduced by the stochastic portion of the data series, i.e.

Wavelets – a Mathematical Microscope

Wavelets are wave patterns of small size, with which mainly rapid changes in a signal or sharp contrasts in an image, and structures at different scales can be studied. The wavelet method, developed in the eighties and now provided with a sound mathematical basis, has become a powerful tool in the field of data compression, noise suppression and processing of images and signals. The most powerful mathematical tool for analyzing and processing signals and images is the Fourier method. Here a function is considered as a superposition of sine and cosine waves with different frequencies. The Fourier method, however, is not very good at detecting rapid changes in signals, such as the onset of a spoken syllable or sharp edges in images. Furthermore, many signals such as speech and video display structure at many different scales. The Fourier method 'misses' most of this multi-scale structure.

Jean Morlet, a French geophysicist employed by the oil company Elf Aquitaine, attempted early 1980s to overcome the deficiencies in the Fourier method when analyzing seismic signals. While building on earlier work by others, e.g. Dennis Gabor, a Nobel laureate in physics, he invented 'wavelets'. A wavelet is a wave-like function with short extension: its graph oscillates only over a short distance, or damps very fast; its mean value over the whole domain (its integral) equals zero. A wavelet is localized both in frequency and position. From this 'mother wavelet' a whole family of other wavelets is derived by displacement and scaling). One usually chooses $\mathbf{\Psi}_{j,k}(x) = 2^{j/2}\mathbf{\Psi}(2^j x - k)$, with j and k integers (k controls displacement, j scaling). In this family every scale is represented for every wanted position. Only a few mild mathematical conditions are needed when 'mother' $\mathbf{\Psi}$ is chosen. A wavelet family is comparable to the set of sine and cosine functions of all frequenties in Fourier analysis.

The FBI has compactly stored its 200 million fingerprints by using wavelets.

In the wavelet approach, a signal is scanned by fitting each member of a wavelet family at each time as a mould to the signal, measuring how good the fit is. The graphical representation of wavelet coefficients resembles a mountain landscape above the time-frequency plane. Any drastic change in the signal is clearly visible as a peak at the high frequencies at the time of change. Wavelet analysis operates as a miscroscope: as a wavelet family contains 'building blocks' of arbitrary small scale, we may zoom in at any time on any detail of the signal. A noteworthy difference with Fourier analysis is that the wavelet coefficients of a one-dimensional signal form a two-dimensional array (depending on time and frequency). Precisely this additional degree of freedom (flexibility) makes possible the success of wavelets in the analysis of sharp transitions and multi-scale structures.

A problem is that wavelets are generally more complicated functions than Fourier's sine waves. Morlet's 'simple' wavelets lack an important factor: they are not orthogonal. As is well-known, orthogonality is an almost indispensible tool in Fourier analysis. The Belgian mathematician Ingrid Daubechies, now working in the USA, constructed in 1988 orthogonal wavelets with finite domain. With this theoretical breakthrough a whole series of practical applications were now close at hand. These wavelets, however, are 'recursive' (there is no explicit formula), but the computer can very well deal with them. When designing efficient wavelet algorithms the choice of the mother wavelet is important for a particular application. In particular France, the USA and Japan are very active in developing such algorithms.

Wavelets operate on the same 'market' as the Fourier method and their range of application is in principle very wide. As we observe images and acoustic signals mainly by contrasts, wavelets can be very effective in storing and transmitting such images and signals by data compression. By applying wavelets the FBI managed a couple of years ago to compress its file of 200 million fingerprints (over a thousand terabytes) to a few percent of its original magnitude, without any loss of the prints identity. Wavelets are also sucessfully applied for removing noise in old sound recordings and restoring images from MRI and CT scans. Since a wavelet family forms a set of self-similar functions, the method is able to detect the presence of fractal components in observational data. The wavelet method has recently been applied to e.g. the study of fluid flows (turbulence), air-condition in offices, tracing very short irregularities in a signal (e.g. an ECG), the proper functioning of an assemby line, and speech. Certainly worth mentioning is the wavelet study of a new golf ball with a 'sporty' sound by Japan (the usual golf ball sounds 'poor' to a Japanese ear and is out of tune in their often very expensive golf-clubs).

the spectrum cannot be determined exactly from the samples in the segment if the data process is stochastic. Neither can it be determined if the data process is a noisy deterministic signal.

In signal analysis with the Fourier transform it is a common technique to preprocess such data segments before analyzing them, by multiplying the data segments with some window function, for example a cosine function. Through this operation the data are smoothed, so that the role of edge effects will decrease. This technique is called 'tapering' in the literature (the window is also called a taper). Further, a seismogram is a non-stationary data process. Therefore an estimate of the time-varying spectrum is more appropriate than an estimate of the spectrum, involving a Fourier transform. Due to the transient character of the seismic data, we have chosen to use the wavelet transform to build time-varying spectra for the analysis of seismic data. Since at higher scales the wavelet transform also uses information of the data outside the short segment, it also cannot deal with the uncertainty introduced by the short segment. Therefore we introduced a method which preprocesses the short segment with the tapering algorithm and then take the wavelet transform of the new data. The first mathematical results of our research contain error and convergence estimates of the new algorithm. Experiments with synthetic data have been done and in the near future also the algorithm will be tested with real seismic data in collaboration with KNMI.

Recently also research on a second topic in analyzing seismograms has been started, namely the detection of time points, where the several phases appear in the seismogram. The most significant phases to detect are the P-phase and the S-phase. To illustrate the importance of this research, the existence of a P-phase is an indication whether the seismic event is a (nuclear) explosion in the subsurface or an earthquake. Furthermore, these specific time points can be used to locate the source of the seismic event. Nowadays methods to detect these phases mainly consist of statistical methods and assumptions on the measured seismic signals. These kinds of methods are also used in The Netherlands at KNMI. At CWI we started to build an algorithm, based on the wavelet transform. Together with the seismology department of KNMI, physical properties of the seismic signals, especially their scaling behaviour are used to refine the algorithm.

What Do You Mean, Coordination?

Research Project	:	Coordination – Experimental Testbed and Applications
Project leader	:	F. Arbab
E-mail	:	Farhad.Arbab@cwi.nl
URL	:	http://dbs.cwi.nl/cwwwi/owa/cwwwi.print_projects?ID=57, and ID=58

Introduction

Electronic computers were first used almost exclusively as number crunchers. With their integration in society other types of use rapidly emerged, and nowadays they are routinely applied in such areas as databases, graphics, etc., with the number of non-numerical applications still rapidly growing. In several applications computers work in cooperation, and their interaction, mutual as well as with human users, is pivotal for the application's proper functioning. Because this cooperation is usually a complex process, special coordination languages have been designed, in which one can specify, analyze and control the cooperation between the various system components. An important application of such languages is in the parallelization of computation intensive sequential programs, as they occur for example in fluid dynamics, matching of DNA strings, molecular synthesis, monitoring of medical data, analysis of financial data integrated into decision support systems, and game playing (chess).

Experimental research on coordination started at CWI in 1990. The major result of these efforts so far has been the development of the coordination language MANIFOLD, based on a novel model for control-oriented coordination (IWIM). The language has found several licensees in Europe as well as the USA.

Present activities include the development of semantic models to create debugging and visualization tools for coordination languages, enhancements to and experiments with MANIFOLD and its visual environment, and its application to a variety of real-life problems in areas such as numerical computing, distributed constraint satisfaction, and shallow water modelling.

Computing

The formal notions of computing and computability were introduced by Church and Turing, as a response to Hilbert's challenge to derive all mathematical truths from a set of axioms by 'mechanical operations'. In order to define such operations, Church introduced his lambda-calculus and Turing what became known as the Turing Machine. As is well-known, Hilbert's programme was brought to a premature end by Gödel's incompleteness theorem, which showed that machines cannot find all mathematical truths. As it turned out, Turing's and Church's formalizations could be used as a reasonable mathematical definition of what we intuitively mean by a *mechanical operation*: any operation that can be performed by a Turing Machine (lambda-calculus, Turing Machines, and all other formalizations so far have been shown to be equivalent). A mathematical definition of computing in a strictly technical sense was provided by Church's thesis: a Turing Machine (or its equivalents) mathematically defines the concept of an algorithm (or an effective, or recursive, or mechanical procedure).

Interaction

Contrary to Turing's abstract computing device, real computers do much more than mere technical computing. Among other things, their physical presence gives rise to heat and noise. More interestingly, computers also act as facilitators, mediators, and coordinators enabling the collaboration of other computers (or computer programs), sensors and actuators that involve their real world environment, or human beings. Performing technical computations and acting as an instrument for the collaboration of other agents are two quite different things. In the latter case *interaction* with these agents is required. This distinction was expressed in the concept of Interaction Machine, proposed by P. Wegner in 1995. An Interaction Machine is an extension of a Turing Machine that can interact with its environment with new input and output primitive actions. Thus, contrary to a Turing Machine, an Interaction Machine is an *open system*, constituting an essential change in the behaviour of the machine. Whereas Turing Machines are restricted to *computable functions* only, Interaction Machines can express *computable non-functions*, which represent the behaviour of many interesting real systems. An Interaction Machine receives input from its environment, whose behaviour cannot necessarily be described by a computable function. Moreover, the output of the Interaction Machine in turn may influence this environment, thus making its next input dependent on its earlier interaction. No input tape can encode such a 'computation' for any Turing Machine.

It is the ability of computers (as Interaction Machines) to interact with the real world, rather than their ability (as mere Turing Machines) to carry on ever-more-sophisticated computations, that is having the most dramatic impact on our societies. In the emerging models of human-computer interaction users increasingly become active components of their running applications, where they examine, alter, and steer on-going computations. It is realized that in many contemporary applications realistic results can be achieved only if human intuition and commonsense is combined with raw, formal reasoning and computation. Interaction Machines are suitable conceptual models for describing such applications.

Interaction Machines are both useful and interesting for formal studies, although they interact unpredictably with the environment, and their open-ness can make their behaviour underspecified, or even ill-defined. The components of real-life systems exhibit this property. Considered in isolation, i.e., without imposing any constraint on its interaction with the environment, each component behaves in an ill-defined way. Only when such open systems come together as components of a larger system, the topology of their interactions forms a context that constrains their mutual interactions and yields well-defined behaviour. The behaviour of an Interaction Machine can be studied as a computation between each pair of its successive interactions. More interestingly, one can abstract away from all such computations, regarding them as internal details of individual components, and embark on a formal study of the constraints, contexts, and conditions on the interactions that ensure and preserve well-behavedness. And this material is the thread that weaves the fabric of coordination.

Coordination

Coordination languages, models, and systems constitute a new field of study, with the goal of managing the interaction among concurrent programs. By studying the dynamic topologies of interactions among Interaction Machines, one aims to construct protocols to realize such topologies that ensure well-behavedness. Analogous to the way in which topology abstracts away the metric details of geometry and focuses on the invariant properties of (seemingly very different) shapes, coordination abstracts away the details of computation in Interaction Machines, and focuses on the invariant properties of (seemingly very different) programs. As such, coordination focuses on *program patterns* that specifically deal with interaction.

Interaction and concurrency are closely related concepts. Concurrency means that computations in a system overlap in time. They may actually run in parallel or be interleaved on a single processor. Parallel computations may be geographically distributed. An interactive system differs from other concurrent systems in that it has unpredictable inputs from an external environment that it does not control.

Coordination is relevant in the design, development, debugging, maintenance, and reuse of all concurrent systems. The primary concern in the design of a concurrent application must be how its various active entities are to cooperate with each other. Most present cooperation models are essentially a set of *ad hoc* templates that have been found to be useful in practice. There is no paradigm wherein we can systematically talk about cooperation of active entities, and wherein we can compose cooperation scenarios such as, or as alternatives to, client-server or workers pool models. Consequently, programmers must directly deal with the lower-level communication primitives to realize their desired cooperation models. This hampers the development of working concurrent applications containing large numbers of active entities with non-trivial cooperation protocols. Since the protocol is often only implicit in the behaviour of the rest of the concurrent software, its maintenance and modification is much more difficult than necessary, and reuse is next to impossible.

A number of software platforms and libraries are presently popular for easing the development of concurrent applications. Such systems, e.g., PVM, MPI, CORBA, etc., are sometimes called *middleware.* Coordination languages can be thought of as their linguistic counterpart. One of the best known coordination languages, Linda, is based on the notion of a shared tuple space: a centrally managed space containing all pieces of information that processes want to communicate. Linda uses a data-oriented coordination model and is as such suited to applications using a substantial shared body of data, for example database and transaction systems used in banking and airline reservations. Linda's underlying model provides primitives that must be incorporated *within* a computation for its coordination. For certain applications this may be a natural procedure, A MANIFOLD program receives two process type definitions (A and M), recursively creates instances of itself (processes x_i), A (processes a_i), and M (processes m_i), and (re)connects the streams routing the flow of information among them. This abstract coordination protocol can be compiled separately, and linked with the object code of other processes to build an application. It is immaterial what exactly processes A and M do, and process instances x_i, a_i, and m_i, that are created at run time, can run on various hosts on a heterogeneous platform. The depth of the recursion (i.e., the exact number of x_i's) depends on the number of input units, and the number of units each instance of a_i decides for itself to consume. What matters is the pattern of communication among, and creation of, these process instances.

Entirely different applications can use this same MANIFOLD program as their coordination protocol. For example, we have used this protocol:

1. to perform parallel/distributed sorting, by supplying:
   - as A, a sorter, each instance of which takes in $n > 0$ input units, sorts them, and produces the result as its output; and
   - as M, a merger, each instance of which produces as its output the merged sequence of the two sorted sequences of units it receives as its input.
2. to perform parallel/distributed numerical optimization, e.g., of a function as depicted below, by supplying:
   - as A, an evaluator, each instance of which takes in an input unit describing a (sub)domain of a function, and produces its best estimate of the optimum value of the function in that (sub)domain; and
   - as M, a selector, each instance of which produces as its output the best optimum value it receives as its input.

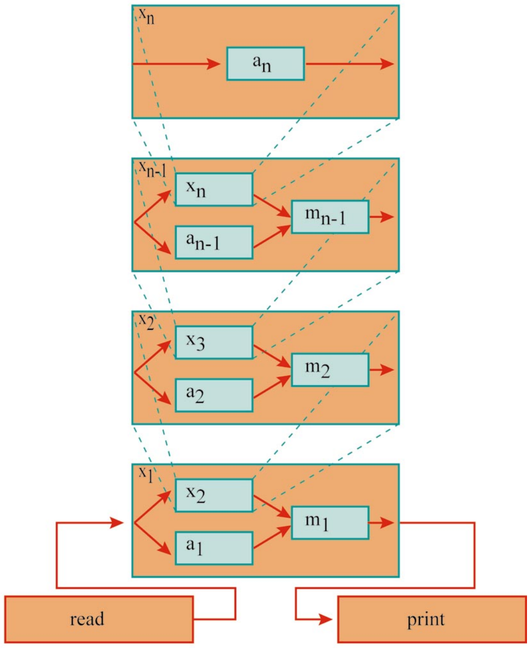

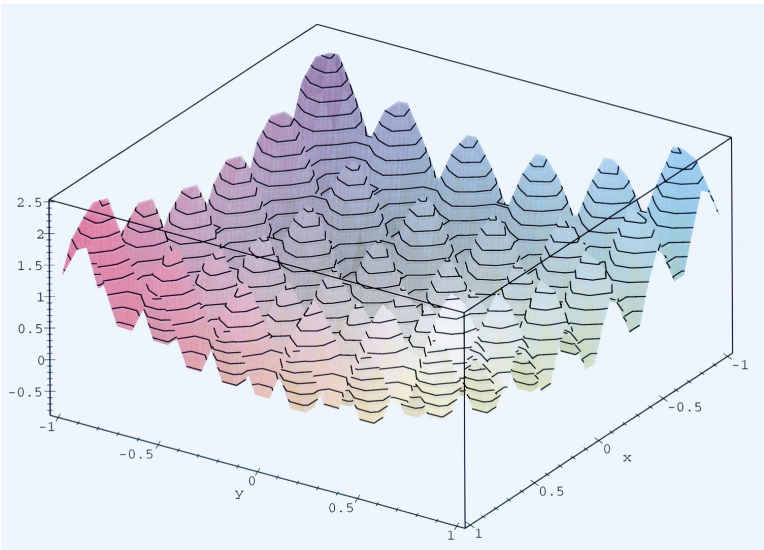

but the intermixing of coordination primitives with computation code has the consequence that no piece of source code is identifiable as the cooperation model or the coordination protocol of an application, that can be designed, developed, debugged, maintained, and reused, in isolation from the rest of the application code.

MANIFOLD

The coordination language MANIFOLD, developed at CWI during the 1990s, follows a different approach. It is control-oriented, directed to aspects such as processing or the flow of control, rather than large-scale data handling, and is suited to, e.g., applications involving work-flow in organizations, or multi-phase applications where the content, format, and/or modality of information substantially changes from one phase to the next. Its primitives support coordination of entities from *without*, so that the coordination modules are separated from the computation modules and, hence, can be reused in other applications. MANIFOLD manages complex, dynamically changing interconnections among sets of independent, concurrent, cooperating processes.

Whereas computation processes can be written in any conventional programming language, coordinator processes are written in the MANIFOLD language. The language is based on the IWIM model (Idealized Worker Idealized Manager), which is a generic, abstract communication model that separates computation (worker) from communication and cooperation (manager) modules. Moreover, in IWIM communicating parties (i.e., modules or processes) need not know each other.

MANIFOLD is a strongly-typed, block-structured, event-driven language, in which, semantically, there is no need for 'computational' entities or constructs like integers, floats, strings, arithmetic expressions, sequential composition, conditional statements, loops, etc. MANIFOLD only recognizes processes, ports, events, and streams (which are asynchronous channels), and its only control structure is an event-driven state transition mechanism. Programming in MANIFOLD is a game of dynamically creating process instances and (re)connecting the ports of some processes via streams, in reaction to observed event occurrences. Since computation and coordinator processes are absolutely indistinguishable from the point of view of other processes, coordinator processes can, recursively, manage the communication of other coordinator processes, just as if they were computation processes. Thus any coordinator can also be used as a meta-coordinator, to build a sophisticated hierarchy of coordination protocols. Such higher-level coordinators are not possible in most other coordination languages and models.

The formal semantics of MANIFOLD is defined in terms of a two-level transition system which reflects the inherent dichotomy of computation vs. coordination in MANIFOLD, and represents a novel application of transition system in formal semantics. The key concept here is that the second level abstracts away the (computational) semantics of the first level processes, and is concerned only with their externally observable behaviour. Traditionally, *initial algebras* have been used as mathematical models for finite data types, such as finite lists. Their counterparts, *final coalgebras*, are used as mathematical models for infinite data types and for the semantics of object-oriented programming languages. More generally, coalgebras can serve as models for dynamical and transition systems. Coalgebraic models seem very appealing candidates for a mathematically sound foundation for the semantics of MANIFOLD. For example, the strict separation of computation from communication in MANIFOLD corresponds directly with the inherent 'strict information hiding' property of coalgebras. An alternative approach is to describe the essence of the semantics of a coordinator process in MANIFOLD as transitions between states, each of which is reminiscent of an (asynchronous) electronic circuit. Category theoretical models have been used to describe simple circuit diagrams. Extending such models to account for the dynamic topological reconfiguration of MANIFOLD is a non-trivial challenge.

MANIFOLD's coordination modules construct and maintain a dynamic data-flow graph where each node is a process. These modules only change the graph topology, by changing the connections among various processes in the application. The computation modules cannot possibly change the graph topology. This separation makes both sets of modules easier to verify and more reusable. For example, the same MANIFOLD coordinator program that was developed for a parallel/distributed bucket sort algorithm, was used at CWI in a quite different research field to perform function evaluation and numerical optimization using domain decomposition (see Box).

The MANIFOLD system runs on multiple platforms and consists of a compiler, a run-time system library, a number of utility programs, and libraries of built-in and predefined processes of general interest. Presently, it runs on IBM RS6000 AIX, IBM SP1/2, Solaris, Linux, and SGI IRIX. A MANIFOLD application consists of a (potentially very large) number of processes running on a network of heterogeneous hosts, some of which may be parallel systems. Processes in the same application may be written in different programming languages and some of them may not know anything about MANIFOLD, nor the fact that they are cooperating with other processes through MANIFOLD in a concurrent application. A number of these processes may run as independent operating-system-level processes, and some will run together as light-weight processes (preemptively scheduled threads) inside an operating-system-level process. None of this detail is relevant at the level of the MANIFOLD source code, and the programmer need not know anything about the eventual configuration of his or her application in order to write a MANIFOLD program.

MANIFOLD has been successfully used at CWI to implement parallel and distributed versions of a semi-coarsened multi-grid Euler solver algorithm. This represents a real-life heavy-duty Computational Fluid Dynamics application where MANIFOLD enabled restructuring of existing sequential Fortran code using pure coordination protocols that allow it to run on parallel and distributed platforms. The results of this work are very favourable: no modification to the computational Fortran 77 code, simple, small, reusable MANIFOLD coordination modules, and linear speed-up of total execution time with respect to the number of processors (e.g., from almost 9 to over 2 hours). Other applications of MANIFOLD include its use in modelling cooperative information systems, coordination of loosely-coupled genetic algorithms on parallel and distributed platforms, coordination of multiple solvers in a concurrent constraint programming system, and a distributed propositional theorem checker.

Numerical Test Set of Real-life Problems

Research Project	:	Parallel codes for the simulation of processes in circuit analysis, etc.
Researchers	:	J.J.B. de Swart, W.M. Lioen
E-mail	:	{Jacques.de.Swart,Walter.Lioen}@cwi.nl
URL	:	http://www.cwi.nl/cwi/project/IVPtestset/

Introduction

For many processes in industry the solution of differential equations is indispensable. Examples include the simulation of electrical circuits, chemical reactions, or the behaviour of a car on a bumpy road, and the steering of robots. Normally the differential equations are far too complicated to solve analytically, and one has to resort to numerical integration techniques implemented on a computer to obtain an approximation to the solution. Numerical analysts are challenged to come up with faster algorithms because problems become ever complexer (for example chip design) and in several cases a decrease in computer run time is required, for example in robotics, where it is not desirable to steer the motion of a robot by software that needs several seconds for computing the next move. To test such codes it is important to have a representative collection of test problems. Such a test set can decrease the effort for the code developer to test his software in a reliable way, and cross the bridge between the application field and numerical mathematics. CWI has built up and maintains such a test set with financial support of the Technology Foundation STW.

What does sensible testing mean?

There is a wide-spread attitude to take the testing of numerical algorithms for solving differential equations not very serious. What one typically encounters in scientific presentations is that the speaker uses 99% of his time to describe his beautiful new method, and then, when the chairman has already coughed a couple of times, he quickly puts up a slide with two curves, which are supposed to explain that his new method works better than some existing one, without explaining how he has performed his tests. Quite often, it turns out that the testing involved only a very simple, academic problem. In scientific literature the attitude is basically the same.

It is easy to justify the need for a precisely defined test protocol that tests on large, real-life problems. Although reliable software for solving differential equations has been around for the past few decades, new codes are necessary because the existing software often performs too slowly on large problems, or cannot handle nasty practical complications, such as discontinuities in the differential equations or the phenomenon that some components of a problem are

The CWI Test Set as a bridge between science and industry.

more sensitive for perturbations than other (*higher index character*). Of course, it does not make sense to promote a new algorithm on the basis of small, 'easy' test problems that could be solved satisfactorily by existing codes. Besides, a new important issue in code development is to make codes suitable for implementation on parallel computers with more than one processor. Especially for large problems it is sometimes necessary to resort to such computers in order to reduce the computing time to reasonable values. Again it is clear that one should not test these parallel algorithms on small problems.

Another justification of performing well-defined real-life tests is that poor testing may lead to not considering methods that perform badly in such tests, although they might work well for certain problems not represented in the tests. A probably even more frequent occurrence is an author's reliance on methods that perform well in a test, although they suffer from all sorts of disadvantages, not revealed by the test.

A good testing procedure should contain the following ingredients. A complete specification of the problem should be given, including initial values, parameter settings, integration interval, etc., in order to ensure that everyone really uses the same test problem. Furthermore, the mathematical results should be interpreted in terms of the original real-life form of the problem. For example, it does not make sense to test the method by computing the behaviour of a car on a smooth road, because that is not the 'hard' part and probably any method will pass this test. What matters is a test on a bumpy road. Along with a test problem, recommendations should be given, e.g., 'notice that a solver might have difficulty with solving this problem, if it does not pay proper attention to discontinuities'.

A test set should be representative for many application areas. Small test sets create inbreeding , in the sense that if codes are developed which work well on a non-representative test set, then there is a risk that only a few problem features are addressed, whereas features absent in the test problems are not addressed at all.

Several problems in the chemical processing industry can only be attacked with advanced mathematical methods, by which the problem is translated into a complicated set of equations. Their solution requires sophisticated numerical techniques. The AKZO Nobel concern contributed two industrial chemical problems to the CWI Test Set. (Photo: Paul Hoedemaekers. Courtesy AKZO Nobel Chemicals bv, Delfzijl.)

Example of a test problem

Here we describe a fairly large test problem which may serve as an example of how to model any chemical reaction scheme. The problem originates from plant physiology and describes how light is involved in morphogenesis. To be precise, it explains the 'High Irradiance Responses'(HIRES) of photomorphogenesis on the basis of phytochrome, by means of a chemical reaction involving eight reactants. The reaction scheme is given in Figure 1. The species $\mathrm{P_r}$ and $\mathrm{P_{fr}}$ refer to the red and far-red absorbing form of phytochrome, respectively. They can be bound by two receptors X and X$'$, partially influenced by the enzyme E. Every reaction scheme of the type shown in Figure 1 can be described by a set with species $\mathcal{S}$, a set of nodes $\mathcal{N}$, a transition matrix $\mathcal{M}$ and a source term vector τ. The set $\mathcal{S}$ contains all species that influence the chemical process. For problem HIRES, $\mathcal{S}$ is given by

$$\mathcal{S} = \{\mathrm{P_r}, \mathrm{P_{fr}}, \mathrm{P_rX}, \mathrm{P_{fr}X}, \mathrm{P_rX'}, \mathrm{P_{fr}X'}, \mathrm{P_{fr}X'E}, \mathrm{E}\}.$$

X, X' and P'_{fr} are not in S, the receptors because they are assumed to abound, and P'_{fr} because its amount can be derived from the conservation of the total amount of phytochrome. The set $\mathcal{N}$ contains all possible mixtures that react, or are formed by a reaction. If we view a reaction scheme as a graph, $\mathcal{N}$ consists of its nodes. The arcs of the graph in Figure 1 can also be represented by a matrix $\mathcal{M}$. Its non-zero elements $\mathcal{M}_{ij}$ are the reaction constants governing the reaction from the ith to the jth node in $\mathcal{N}$, for example $\mathcal{M}_{12} = k_1$. It may happen that certain species of the set $\mathcal{S}$ are continuously added to the chemical process. These added quantities are called source terms and form the non-zero entries of a vector τ. For problem HIRES, there is only a source term for $\mathrm{P_r}$: $\tau_1 = O_{k_s}$. Once $\mathcal{S}, \mathcal{N}, \mathcal{M}$ and τ are available, it is possible to write down the differential equations describing the change of the concentrations of the species in time. The concentration $[S_k]$ of S_k, the kth species in $\mathcal{S}$, decreases as a result of every reaction evolving from a node that contains S_k, and increases whenever S_k is in a node, to which a reaction leads. The rates of decrease and increase equal the velocities v_{ij} of the corresponding reaction. This velocity equals the product of the concentrations $[S_k]$ of all species in the node from which the reaction evolves, times the reaction constant $\mathcal{M}_{ij}$. This leads to the following differential equation for $[S_k]$:

$$\mathrm{d}[S_k]/\mathrm{d}t = \sum_i \sum_{\{j|\mathcal{M}_{ij}\neq 0 \ \wedge S_k \in \mathcal{N}_j\}} v_{ij} - \sum_j \sum_{\{i|\mathcal{M}_{ij}\neq 0 \ \wedge S_k \in \mathcal{N}_i\}} v_{ij} + \tau_k.$$

We have chosen realistic values for the reaction constants, and started the process by mixing 1 mol of $\mathrm{P_r}$ with 0.0057 mol of E. Solving the problem with PSIDE (Parallel Software for Implicit Differential Equations), a package developed at CWI, we obtained the behaviour of all chemical species involved. As an example we plotted $[P_{\mathrm{fr}}X']$ in Figure 2, which shows that this concentration initially grows and then tends to zero.

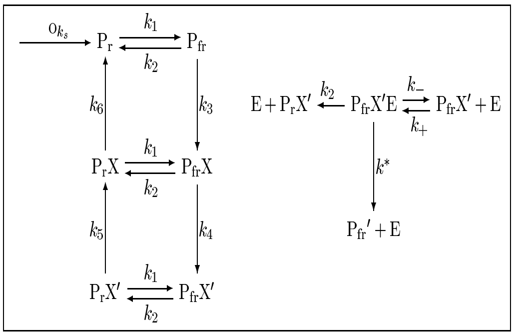

Figure 1: Reaction scheme for problem HIRES.

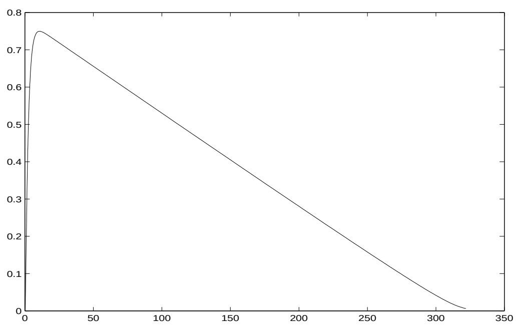

Figure 2: Concentration in mols of $P_{fr}X'$ plotted as a function of time in seconds.

Test problems should be of interest for people from industrial application fields. If so, they can help bridge the gap between industrial engineering and the numerical analysis community. Using industrial real-life problems as test problem assures that the solvers developed can deal with practical issues.

Using a good test set can save a lot of time and effort for the developer of software for differential equations, because the test problems need not to be programmed and described anymore in his article, nor does he need to produce the test results for reference solvers.

The CWI Test Set tries to fulfill these demands. It is available via the World Wide Web and is used by a lot of researchers all over the world.

The structure of the CWI Test Set

The CWI Test Set consists of a descriptive part and a software part. The first part describes test problems and reports on the behaviour of a few state-of-the-art solvers on these problems. The software serves as a platform to test the performance of a solver on a particular test problem yourself. This test test is available via the WWW page http://www.cwi.nl/cwi/projects/IVPtestset/, or via anonymous ftp at site ftp.cwi.nl in the directory pub/IVPtestset/ .

The descriptive part is divided into chapters, each describing one test problem. Every chapter contains four sections. The first section presents general information, such as problem dimension, index, and contributor. The second section contains the mathematical description of the problem, with all necessary ingredients for implementation given in mathematical formulae. The third section describes its origin, in order to enable a physical interpretation of the

problem, and gives references to the literature for further details. This section is important for bridging the gap between the application field and numerical mathematics. It describes the modelling process and gives a feeling for the important characteristics of the problem. The final section contains the numerical solution of the problem. Here one can find a reference solution, the characteristics of runs, in which several solvers solve the test problem, and a work-precision diagram. To obtain such a diagram, a range of input tolerances is used to produce a plot of the accuracy of the result obtained by a specific solver, measured as the number of correct digits in the numerical solution, against the number of CPU seconds needed for the run on a specific computer. Thus one gets an impression of the work needed to obtain a certain precision.

The software part of the test set offers a number of Fortran 77 sources, enabling the user to perform a test run himself. There are four categories: sources with the problem definitions; sources of well-known solvers; a source to handle the input of tolerances and output of the integration run; and sources with so-called drivers. These drivers contain global declarations and form an interface between the problem defining routines and a specific solver. Here all solver options are set as well. The availability of these drivers considerably decreases the effort of performing runs, because their design is a cumbersome and error-prone job.

Compiling a specific solver and linking it to its driver, a problem and the aforementioned I/O source yields a computer program that, for a certain input tolerance, delivers a table of run characteristics.
The author of an article, in which a new algorithm is presented, can use these results as reliable reference. He can easily adjust a driver so that it links the test problem to his solver, and produces numerical results which may be compared with those of preceding tests.

Authoring Multimedia for the Web

Research Project	:	Distributed multimedia authoring systems and applications
Researchers	:	Dick C.A. Bulterman, Lynda Hardman
E-mail	:	{Dick.Bulterman, Lynda.Hardman}@cwi.nl
URL	:	http://www.cwi.nl/Chameleon/ and http://www.w3.org/AudioVideo/

Introduction

The World Wide Web, generally seen as *the* embodiment of the infrastructure of today's information age, is currently unable to handle documents containing continuous media such as audio and video in a user-friendly way. Java extensions to HTML, known as Dynamic HTML, provide synchronization support into Web documents, such that the author keeps full control over the interactions within a document. However, defining even simple synchronization relationships becomes a relatively difficult task for the vast majority of Web users, who have little or no programming skills.

In early 1997, the W3C synchronized multimedia working group (SYMM) was established to study the definition of a declarative multimedia format for the Web. The first system to propose such a format was CMIF, developed at CWI. The SYMM working group has developed a declarative specification language SMIL (Synchronized Multimedia Integration Language). SMIL concentrates on the components of the language but does not specify any particular playback or authoring environment functionality. CWI, a SYMM participant, has developed an authoring environment GRiNS, a graphical interface for SMIL which, in combination with a browser, will enable users to explore SMIL's potential.

Below we review the SMIL language and briefly discuss typical applications and runtime environments, followed by an overview of the GRiNS authoring interface.

A Typical SMIL Document

As an example of the type of applications that SMIL can express, we present a generic Web application: a network newscast. Figure 1 shows two views of such a newscast, taken at different times in the presentation. On the left side, we see a portion of the introduction of a story on the growth of the World Wide Web. In this portion, the local host is describing how sales of authoring software are expected to rise sharply in the next six months. Figure 1(b) shows a point later in the presentation, when the local host is chatting with a remote correspondent in Los Angeles, who describes that many Hollywood stars are already planning their own audio/video homepages on the Web. In a concluding segment the local host gives a wrap up, which leads into the trailing theme music. Finally, each correspondent also has a Web homepage that can be accessed from the story. The ability to link to various homepages makes the semantic content of the document dynamic. As shown in Figure 2, the information content can be augmented during the story.

A Brief SMIL Overview

SMIL is a declarative language for describing Web-based multimedia documents that can be played on any SMIL-compliant browser. Such browsers may be stand-alone presentation systems that are tailored to a particular user community or they could be integrated into general purpose browsers. In SMIL-V1.0, language primitives have been defined that

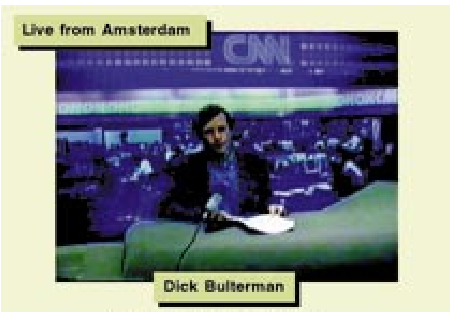

(a) Early in the Story

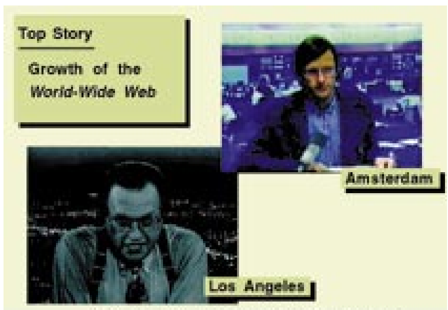

(b) During the Remote Report

Figure 1: Two views of the Web newscast.

allow for early experimentation and (relatively) easy implementation; this has been done to gain experience with the language while the protocol infrastructure required to optimally support SMIL-type documents is being developed and deployed.

The basic structure of SMIL is presented in Figure 3. Architecturally, SMIL is an integrating format: it describes how various components are to be combined temporally and spatially to create a presentation. It also defines how presentation resources can be managed, and it allows anchors and links to be defined within the presentation. Note that SMIL is not a replacement for individual formats (such as HTML for text, AIFF for audio or MPEG for video); rather, it takes information objects encoded using these formats and combines them into a presentation.

SMIL describes four fundamental aspects of a multimedia presentation:

- *temporal specifications:* primitives to encode the temporal structure of the application and the refinement of the (relative) start and end times of events;
- *spatial specifications:* primitives provided to support simple document layout;
- *alternative behaviour specification:* primitives to express the various optional encodings within a document based on systems or user requirements; and
- *hypermedia support:* mechanisms for linking parts of a presentation.

GRiNS

GRiNS is an authoring and presentation system for SMIL documents. It is a part of the CHAMELEON multimedia document processing suite developed in the ESPRIT project CHAMELEON. Where the

Figure 2: Augmenting information during the news.

authoring tools in the CHAMELEON suite can convert documents to a variety of presentation formats based on the (dynamic) needs of the authors, GRiNS consists of an authoring interface and a runtime player which can be used together to create and play SMIL documents. The GRiNS authoring tool is based on a structured approach to document creation. The GRiNS presentation tool provides a flexible presentation interface that supports all of the semantics of the SMIL V1.0 specification. Below we give examples of how GRiNS can be used to create the Web News example and discuss some of the issues involved in providing support for the GRiNS presenter.

The GRiNS authoring environment supports creation of a SMIL document in terms of three views:

- *the logical structure view*, where coarse-grain definition of the temporal structure takes place;
- *the virtual timeline view*, where fine-grain temporal interactions are defined; and

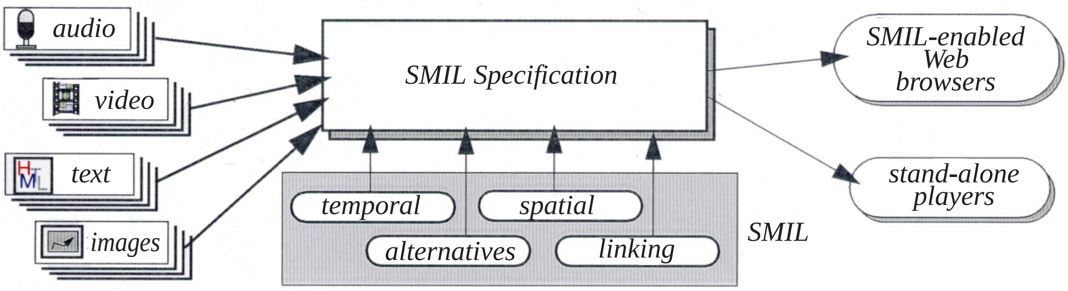

Figure 3: SMIL architecture.

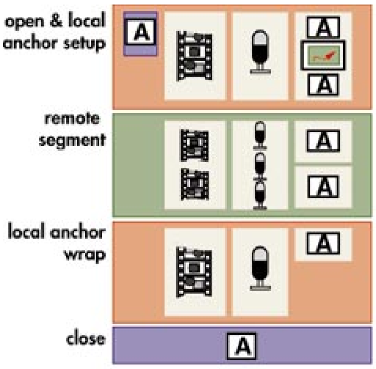

Figure 4: The structure of the Web Growth story. (Time flows from top to bottom.)

- *the playout view*, where layout projections are defined.

Hyperlink definition and specification of alternative data objects can occur in either the logical structure view or the virtual timeline view.

The Logical Structure View

If we were to define the Web Growth story in terms of its over-all structure, we would wind up with a representation similar to that in Figure 4. The story contains an introduction by the local host, followed by a report by the remote correspondent, and then concluded by a wrap-up by the local host. This 'table of contents' view defines the basic structure of the story and at this level may be reusable.

The GRiNS logical structure view allows an author to construct a SMIL document in terms of a nested hierarchy of media objects. The interface provides a scalable view of the document being created. Note that only structure is being edited here: much of this creation takes place before (or while) actual media objects have been created. The logical structure view has facilities for cutting and pasting parts of the presentation, and it allows individual objects or sub-structure to be previewed without having to play the entire application.

During design and specification, the values of object attributes can be entered by the author. In practice, duration of an object or a group will be based on the enclosing structure, which will be calculated automatically. Note that while the basic paradigm of the logical structure view is a hierarchical structure, the author can also specify loop constructs which give (sub-)parts of the presentation a cyclic character.

The Virtual Timeline View

Fine-grain timing relationships in a presentation are often difficult to visualize with only a logical structure view. For this reason, GRiNS also supports a timeline projection of an application. The GRiNS timeline is virtual: it displays the timing relationships as calculated from the logical structure view. This means that the user is not tied to a particular clock or frame rate, or to a particular architecture. (The actual timing relationships will only be known at execution time.)

Figure 5 shows a virtual timeline of the Web Growth story, with all components and their relative start and end times. The timeline view provides

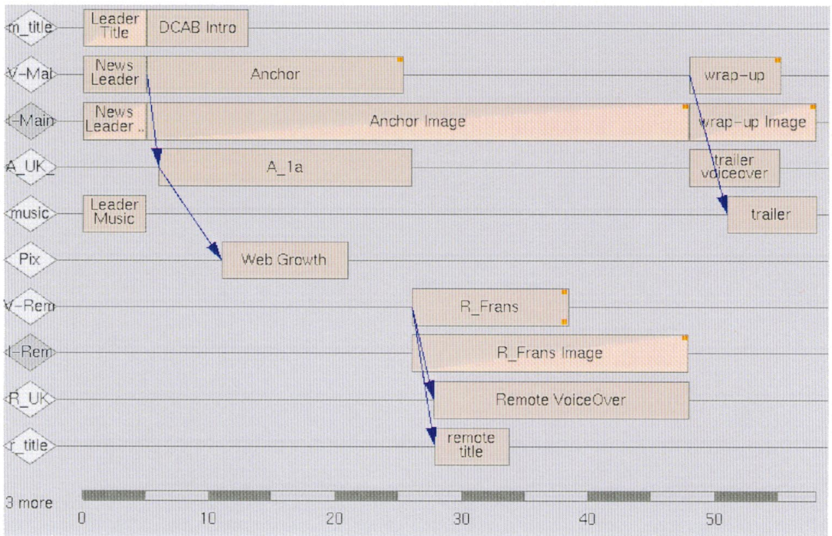

Figure 5: The virtual timeline view of the WebNews. (Time flows from left to right.)

Figure 6: The (a) presentation view and (b) associated channel map.

an insight into the actual temporal relationships among the elements in the presentation, showing the declared length and start times of objects whose duration or start offset is pre-defined, as well as the derived length of objects whose duration depends on the document structure. The timing view shows pre-defined durations as solid blocks and derived durations as blocks with a diagonal line. Events on the timeline are partitioned by the various output channels defined by the user. In the figure, channels which are darkened are turned off during this (part of the) presentation. This allows the author to see the effect of various alternate settings during the presentation.

When working with the virtual timeline view, the user can define exact start and end offsets using declarative mechanisms (sync arcs), shown as arrows in the figure. These specifications of timing relationships can be evaluated at runtime or by a scheduling pre-processor.

If changes are made to the application or to any of the attributes of the media objects, these are immediately reflected in the virtual timeline view. When the presentation view is active (see below), the virtual timeline view also displays how the playout scheduler activates, arms and pre-fetches data during the presentation.

Note that both the logical structure and the virtual timeline views are authoring views; they are not available when a document is viewed via the playout engine only. Both views shield presentation authors from the SMIL syntax.

The Presentation View

The presentation view provides during authoring a WYSIWYG view of the document under development, and an interactive mechanism for laying out the associated output channels. Figure 6(a) shows part of the presentation being developed, while Figure 6(b) illustrates the associated channel map. During playout, the values of the layout attributes can be dynamically changed as required. When the SMIL language becomes accepted as a W3C recommendation, players of the language will become available. In the meantime, the GRiNS authoring environment includes a playout engine conforming to the SMIL working draft.

Conclusions

The development of SMIL enables the creation of multimedia content by large numbers of users that is guaranteed to be playable on a wide range of software and hardware. This will encourage and stimulate both professional and amateur authors. While documents can be edited by hand, an authoring environment enables manipulation of the different aspects of a presentation in terms of the structure, timing and/or layout of a presentation rather than the underlying SMIL syntax. A future advantage of SMIL is that it allows multimedia documents to become part of a larger information processing environment. Source material can be stored in presentation independent forms and only at a later stage converted to formats for final presentation. Work in this direction is continuing within the research group.

REFERENCE

Dick C.A. BULTERMAN, Lynda HARDMAN, Jack JANSEN, K. Sjoerd MULLENDER, Lloyd RUTLEDGE (1998). GRiNS: A GRaphical INterface for creating and playing SMIL documents. *Proceedings WWW7, the 7th World Wide Web Conference*, Brisbane, Australia.
http://www7.conf.au/

Large-scale Computing at CWI

Introduction

A substantial part of CWI's research programme requires large-scale computing, where the capacity and processing speed of the most advanced computers are used to their limits. Moreover, CWI is also prominent as an organizer in this field. For example, the institute co-founded the Platform HPCN (High Performance Computing and Networking) – an initiative of the national HPCN Foundation – , which aims at promoting the lasting use of computational knowledge, methods and techniques, thus reinforcing the knowledge infrastructure and market position of the Dutch trade & industry sector. Through this Platform CWI contributes to the HPCN knowledge transfer on the European scale. Research in this field gained substantial momentum through the national HPCN Programme, an initiative of the ministry of Economic Affairs, in which an amount of 75 million guilders from the natural gas revenues was allotted for this purpose. CWI acquired participation in four out of the six projects granted in the first round of calls (presently there are thirteen projects). Some other projects at CWI requiring large-scale computing are supported by the NWO Programme on Massively Parallel Computing. The research described below gives an impression of the diversity of CWI's activities in the field of large-scale computing.

HPCN-EZ: HPV

The High Performance Visualization (HPV) project aims at substantially enlarging the accessibility of HPCN applications through general purpose, user-friendly visualization tools. To this end CWI develops an interactive visualization environment for high performance computing. The research is focused on data handling of large data sets and comparative visualization techniques. Concrete applications included the use of high performance techniques in the visualization of flows and of global smog models. Work continued on new 3D interaction techniques for intuitive specification of visualization. Another work concerned the visualization of abstract data structures, for example the internal structure of compilers, the content of a Web site, or data resulting

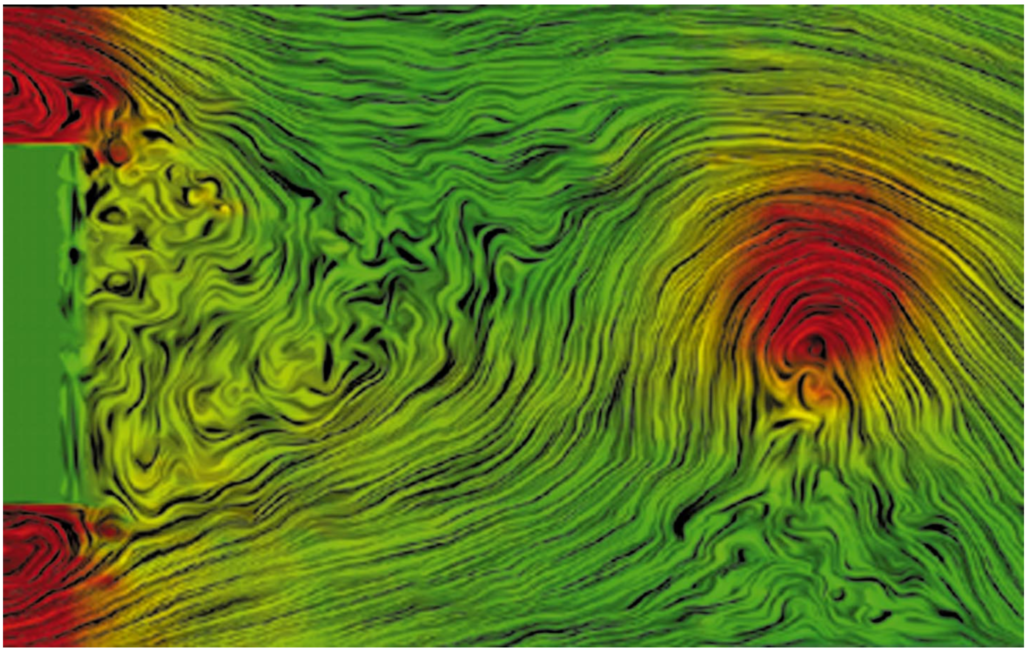

A cross-section of turbulent flow around a block. Colour indicates pressure. Such large-scale HPCN problems require research in novel computation, data management, and visualization methods. (Data: RUG, visualization: CWI.)

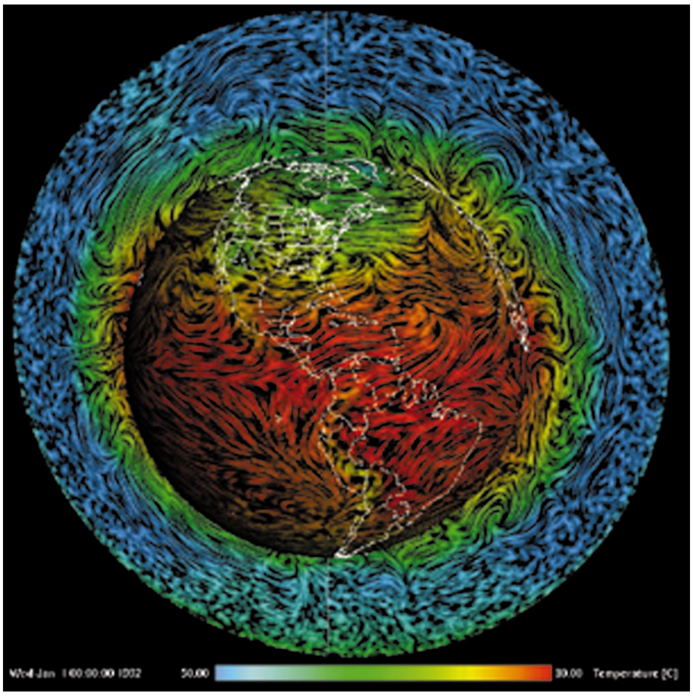

Global climate modelling. Colour indicates temperature, the flow pattern indicates the wind field. CWI researcher Edwin Spee completed a Ph.D. thesis on numerical methods in global transport-chemistry models.

from a search in a digital library. The visualization tools to be developed should be usable in generally available environments (PC's, WWW, etc.). Current attention concentrates on graphical representations of abstract graphs (cooperation with the University of Bordeaux). A Java implementation of a 2D visualization of tree structures turned out to work very well.

HPCN-EZ: TASC

The TASC project (Transport Applications and Scientific Computing) develops algorithms and parallel software for the simulation of 3D transport of pollutants in the atmosphere and in surface water, as well as fully integrated models and their application to environmental problems. CWI is a member of the TASC consortium, which was created for this purpose, and carries out research jointly with TNO to develop a regional 3D long-term ozone simulation model, replacing the current LOTOS model used by TNO. CWI contributes with the design of the mathematical model for a hybrid (terrain following and pressure based) coordinate system and, in particular, of tailored numerical algorithms and implementations on parallel and super-computers. Other research concerns the design of parallel numerical methods for the simulation of water pollution, the marine eco-system, dispersion of river water, sediment transport, etc. In connection with this a 3D transport model was studied, in particular the iterative solution of the equations resulting from their implicit time-discretization, and domain decomposition with domains of varying grid resolutions.

HPCN-EZ: NICE

The NICE project is concerned with numerical methods for flow simulation. Here CWI contributes with restructuring the models in use at Delft Hydraulics for the simulation of shallow water systems. These models, which are programmed in sequential Fortran code, are being adapted to be used in a parallel and distributed environment by applying MANIFOLD, a coordination language developed at CWI, which ensures control over the restructuring of complex programming processes.

HPCN-EZ: IMPACT

The financial world expresses a rapidly growing need for quick analysis of large sets of financial data and a user-friendly availability of the results. In the IMPACT project applications of HPCN techniques in a parallel and distributed environment are studied for this purpose. CWI contributes with the development of algorithms for datamining of financial time series. As it turns out, the discovery of characteristics in such a time series with wavelet analysis is a

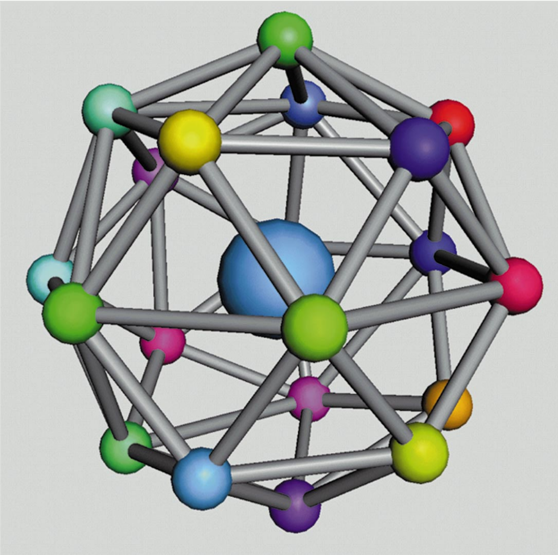

Solution of the Fekete problem: 20 balls optimally spread over a sphere. The large cental ball is added to facilitate 3D perception. The problem belongs to a test set solved by applying novel parallel computing methods developed at CWI.

promising approach. Furthermore, the architecture of the database management system Monet (developed at CWI in collaboration with the University of Amsterdam) was adapted to experiment with query optimization in a massively parallel environment. In a collaboration with Data Distilleries – one of CWI's spin-off companies, specialized in datamining – it was shown that this system performs here better than current database management systems.

Plasma physics simulation

Confronted with computational problems exceeding the power of present super- (vector-) computers, the obvious question is whether these could possibly be solved on parallel computers. To this end NWO initiated some years ago its Massively Parallel Computing programme. One of CWI's contributions to this programme concerns parallel numerical methods for problems in plasma physics, in particular magneto-hydrodynamics. This research is carried out jointly with the FOM Institute for Plasma Physics, Utrecht University and Delft University of Technology, with additional funding from a Cray Research Grant, obtained through the NWO foundation NCF. It focuses on the development, implementation and analysis of parallel iterative methods for the computation of eigenvalues and eigenvectors of generalized eigenvalue problems, with large, sparse and (complex) non-Hermitean matrices. Such eigenvalue problems emerge, e.g., in the study of the stability of tokamak plasmas. The developed computation methods were successfully tested on massively parallel systems, including a Cray T3E with 80 processors at Delft University of Technology and a Cray T3E at Eagan (USA) with 512 processors.

Number theory and data security

CWI cherishes a long tradition of research into number-theoretic problems, where the use of fast computers (in particular vector and parallel systems) can considerably contribute to their solution. Since many of these problems are extremely suitable for parallelization and require large-scale computations, this is an excellent testbed for the application of HPCN-techniques. In particular the advance of public-key cryptography stimulated the study at CWI of algorithms for factorization and primality testing, for computing discrete logarithms, and for the solution of large, sparse systems of linear equations over finite fields. An appealing result was the factoring world record, established at CWI on September 3, by the computation of the factors of the 180-digit number $N = (12^{167} + 1)/13$, with the so-called Special Number Field Sieve algorithm. The factorization took only twelve days of computation on 85 SGI/Cray computers at CWI. A special step in the computation, which requires a lot of memory, was carried out at the Amsterdam academic computing centre SARA on its Cray C90 supercomputer.

Modelling of biological processes

In another project of the NWO Programme on Massively Parallel Computing CWI worked on the modelling of biological systems based on individual characteristics. Such models are closer to reality than models without physiological structuring, but they are also far more difficult to compute. The complexity of the computations forces one to use parallel algorithms. With the algorithms developed at CWI continuation and bifurcation computations can be carried out for various problems. Such computations yield more insight into the asymptotic behaviour of the models in question when certain model parameters are varied. All algorithms are implemented in the BASE software package and can be downloaded from the CWI ftp server.

FINANCES, PERSONNEL, PH.D. THESES

FINANCES 1997

In 1997, SMC spent Dfl. 24.11 million. The expenses were covered by a subsidy from NWO (Dfl. 17.60 million), other subsidies and grants (Dfl. 0.44 million), and from the international programmes (mainly EC programmes, e.g., ESPRIT and HCM) (Dfl. 1.02 million). Finally, an amount of Dfl. 6.31 million was obtained as revenues out of third-party-services and other sources.

Income CWI

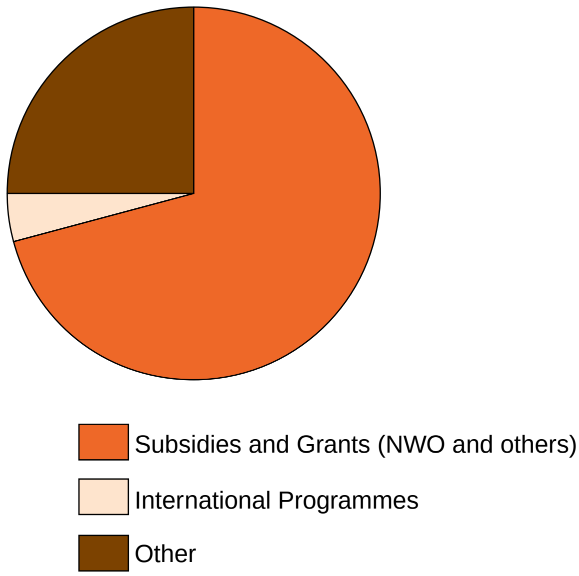

CWI Personnel:
145 fte + 40 fte seconded

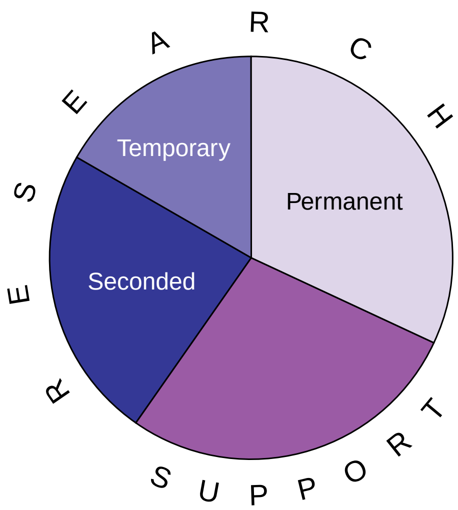

Expenses CWI

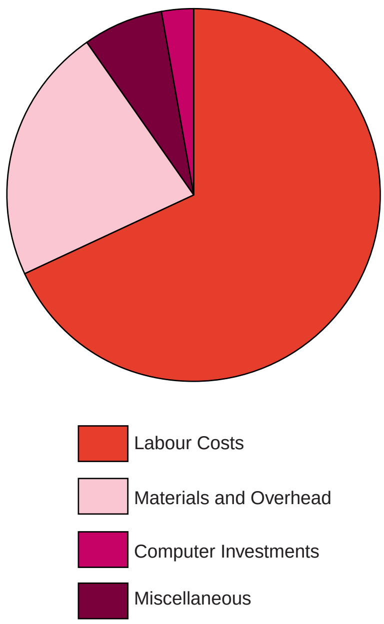

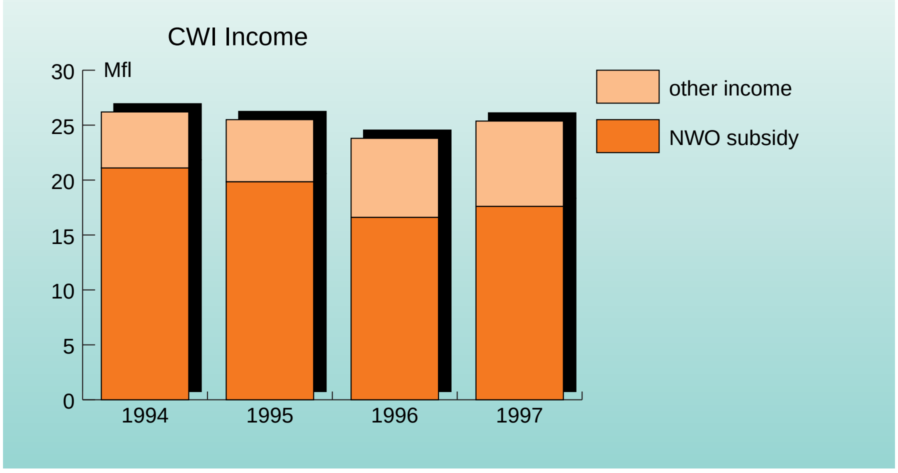
CWI Income
Mfl
30
25
20
15
10
5
0
1994
1995
1996
1997
other income
NWO subsidy

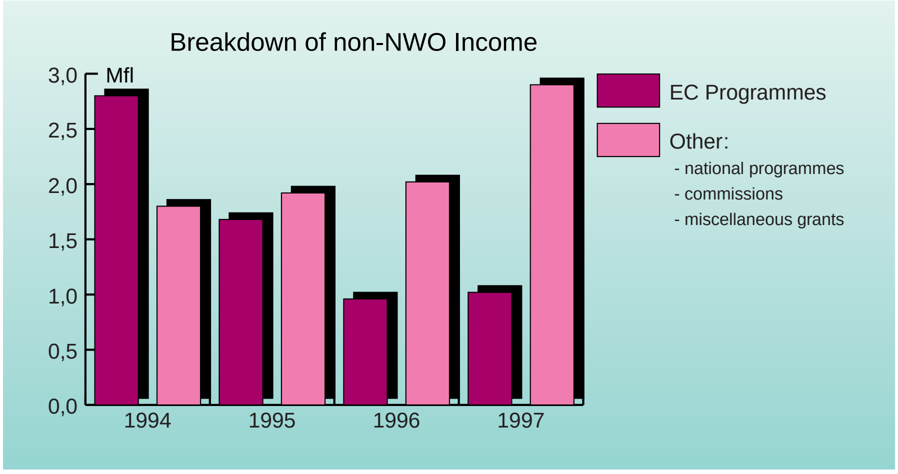
Breakdown of non-NWO Income
Mfl
3,0
2,5
2,0
1,5
1,0
0,5
0,0
1994
1995
1996
1997
EC Programmes
Other:
- national programmes
- commissions
- miscellaneous grants

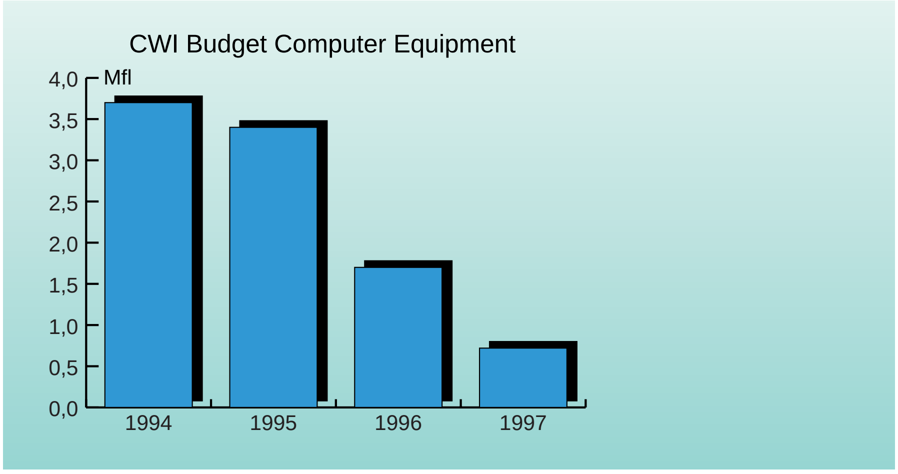
CWI Budget Computer Equipment
Mfl
4,0
3,5
3,0
2,5
2,0
1,5
1,0
0,5
0,0
1994
1995
1996
1997

CWI Ph.D. THESES

Author	Title	Thesis advisor(s)+)
M.A. Peletier	Problems in Degenerate Diffusion	C.J. van Duijn
P.M. de Zeeuw	Acceleration of Iterative Methods by Coarse Grid Corrections	P.W. Hemker, P. Wesseling (TUD)
R.J.F. Cramer	Modular Design of Secure Yet Practical Cryptographic Protocols	P.M.B. Vitányi
R.-M. Elkenbracht-Huizing	Factoring by the Number Field Sieve	R. Tijdeman (RUL)
W.A. van der Veen	Parallelism in the Numerical Solution of Ordinary and Implicit Differential Equations	P.J. van der Houwen
H. Boender	Factoring Large Integers with the Quadratic Sieve	R. Tijdeman (RUL)
D.J.B. Bosscher	Grammars Modulo Bisimulation	J.A. Bergstra (UvA)
J.E.A. van Hintum	Quality Constraints & Constrained Quality	D.K. Hammer (TUE), S.J. Mullender (UT)
A.V. Groenink	Surface without Structure – Word Order and Tractability Issues in Natural Language Analysis	D.J.N. van Eijck W.C. Rounds (UoM)
J. Pellenkoft	Probabilistic and Transformation based Query Optimization	M.L. Kersten
J.J.B. de Swart	Parallel Software for Implicit Differential Equations	P.J. van der Houwen
A.B. Ermakov	Percolation and Coalescing Particle Systems	M.S. Keane
M.I.J. van Dijke	Multi-phase Flow Modeling of Soil Contamination and Soil Remediation	C.J. van Duijn F.A.M. de Haan (LUW)

+) For external advisors the university's acronym is added:

TUD = Delft University of Technology
RUL = University of Leiden
UvA = University of Amsterdam
TUE = Eindhoven University of Technology
UT = Twente University
UoM = University of Michigan
LUW = Agricultural University Wageningen

COMPUTING EQUIPMENT RESOURCES

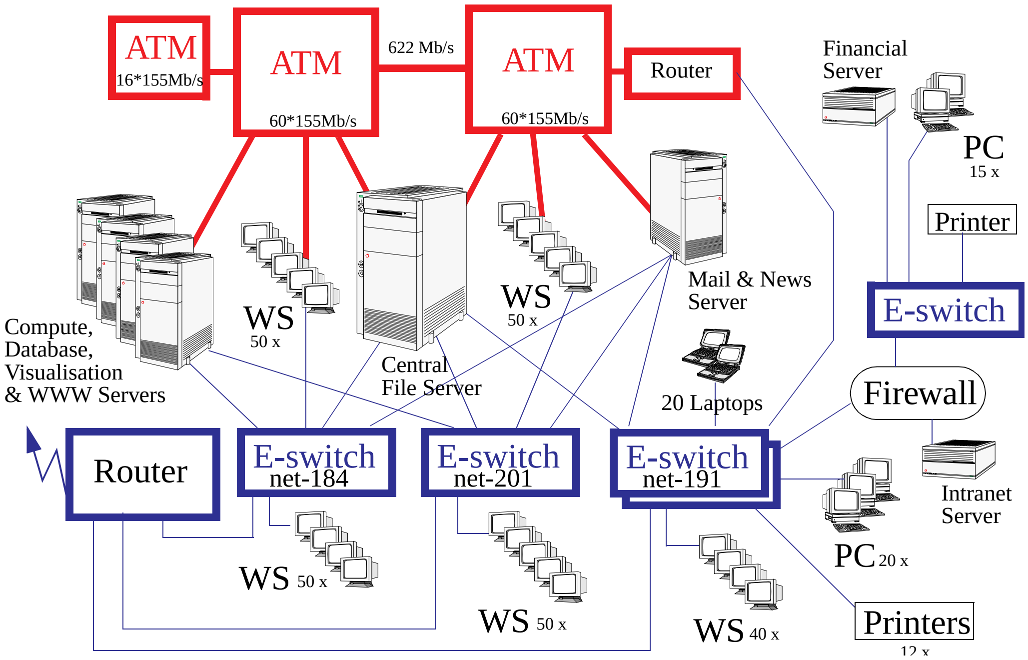

The colour red identifies our ATM-to-the desktop, 155 Mbit/sec, infrastructure. Blue refers to the 10Mbit/sec, coax Ethernet infrastructure. In 1997, the old coax backbone has been replaced by UTP from the distributed concentrators (not shown), to several Ethernet switches in the central computing room.

The above picture is only an abstraction of the total IT-environment in use. To give some more information of what constitutes the CWI IT-environment, the following numbers may give some idea: 175 research staff; 45 supporting staff; 350 user and pseudo accounts; ten central servers, ranging from a central file server as a spider in the network and the E-mail and News server, to smaller dedicated servers like name server, printer server, WWW server, SLIP/PPP server, etc.; 240 desktops and laptops; 400 network nodes; 1500 software packages, available to each (Unix) desktop; 50 PC-based desktops and laptops; three operating systems in nine flavours; three ATM switches; two routers; five 10Mbit/sec Ethernet switches; 24 Ethernet concentrators; 228 km fibre glass cable 'to the desktop'; supporting industry standards like: TCP/IP, SNMP, SMTP, SLIP & PPP, PostScript, the TeX family, PDF, AppleTalk, Novell, NFS, SMB, NIS, DNS, HTML, and just about anything else our researchers can come up with.

CWI RESEARCH PROGRAMMES

Probability, Networks & Algorithms

Cluster leader: O.J. Boxma

Networks and Logic – Optimization and Programming

Theme leader: A.H.M. Gerards

Networks and optimization

Design, analysis and implementation of optimization and approximation algorithms for combinatorial problems with the help of methods from graph theory, topology, discrete mathematics, geometry, and integer and linear programming, with special attention to network problems (flows, routing and VLSI-design) and scheduling and timetabling.

Constraint and integer programming

Study of the foundations and applications of constraint programming, in particular the design and implementation of an adequate programming environment for constraint programming, and the use of constraint programming for various optimization problems drawing on integer programming techniques.

Parallel declarative programming

Study of the relation between two declarative programming paradigms: logic and lazy functional programming. The envisaged translation should lead to an alternative, parallel implementation of a class of logic programs.

Traffic and Communication – Performance and Control

Theme leader: J.H. van Schuppen

Communication and computer networks

Modelling, performance analysis and control of traffic streams in communication and computer networks, in particular queueing-theoretic methods for congestion phenomena and network control by discrete event and stochastic methods.

Traffic networks

Performance analysis and control of traffic in urban, motorway, railway, air traffic, and other networks, in particular performance aspects of congestion, reliability and availability, and the development of control theory for discrete event and hybrid systems.

Control and system theory

Control and signal processing require mathematical models in the form of a dynamic system or a control system. Research is directed at realization and system identification problems.

Stochastics

Theme leader: M.S. Keane

Probability

Research in probability theory and its applications, with emphasis on combinatorial probability and large random systems.

Statistics

Theoretical and applied research in statistics, with emphasis on resampling methods and spatial statistics: statistics for inhomogeneous spatial Poisson processes, bootstrap resampling, nonparametric regression, inference on rare errors.

Stochastic analysis

Theoretical and applied research, in particular statistics for stochastic processes, with applications in the financial world, such as option pricing.

Ergodic theory

Theoretical and applied research in ergodic theory and stationary processes: applications of ergodic theory to number theory, convolutions on the motion group of the plane, simplification of proofs of ergodic theorems, nonexpansive mappings and their asymptotics, mathematical methods for stochastic discrete event systems.

Signals and Images (pilot theme)

Theme leader: H.J.A.M. Heijmans

Image modelling and coding

Statistical analysis of images by way of quadtrees, modelling of bit-error bursts in transmission lines, and fractal image coding.

Wavelets
Analysis of seismic data, in particular time-varying spectral estimation, and preprocessing with the wavelet X-ray transform.

Morphological image processing
Research on mathematical morphology and its applications to image and signal processing, in particular connected morphological operators, morphological pyramids and wavelets, and multi-scale image analysis and applications.

Stochastic geometry
Research on random sets (parameter estimation) and spatial statistics.

Image generation for virtual environments
Fundamental algorithms and software tools for viewer centered interaction paradigms and visual display techniques, in particular incremental visibility calculation and view volume culling, and the development of viewer centered virtual environments.

Software Engineering

Cluster leader: J.W. de Bakker

Interactive Software Development and Renovation

Theme leader: P. Klint

Software renovation
Development of new technology for the renovation and maintenance of legacy systems. Present focus is on the SOS Resolver project, addressing in particular the development of a language-independent, re-usable renovation architecture as well as actual renovation problems such as code restructuring and migration to object technology.

Optimization of scientific software
Design and implementation of language independent program analysis and transformation tools, restructuring of software for the numerical solution of PDE's using algebraic structures from global analysis, and investigation of the algebra of software module composition operations.

Interactive visualization environments
Management of scientific data, visual geometry specification, techniques for high level input, and distributed visualization environments, with applications to HPCN as well as traditional design and engineering.

ASF+SDF
Further development and reimplementation of the ASF+SDF Meta-Environment, in particular compilation, parser generation, unparser generation, and global architecture.

Specification and Analysis of Embedded Systems

Theme leader: J.F. Groote

Process specification and analysis
The study of specification and analysis techniques for process behaviour, in particular the search for a real time extension to the language μCRL. Development of test theory and, in cooperation with Philips NatLab, practical testing techniques, as well as means to analyse (real time) μCRL specifications, and implementation of these in the μCRL tool set.

Proof searching and proof checking
The study of proof search in simple logical systems, viz., propositional logic, in order to increase the efficiency of symbolic verification techniques applied to verify requirements on processes. Furthermore, the development of proof checking methods to establish the correctness of programmed systems 'beyond any reasonable doubt'.

Coordination Languages

Theme leader: J.J.M.M. Rutten

Formal methods for coordination languages
On the basis of transparant semantic models formal methods are developed for coordination languages, including the description of an operational model for Manifold.

Experimental testbed for control-oriented coordination
Design and implementation of an experimental testbed for practical control-oriented coordination programming on heterogeneous platforms and its programming support environment. This testbed is currently accessible through the coordination language Manifold, its utilities, and its visual programming interface, Visifold.

Coordination applications
Study and development of practically useful coordination patterns and protocols in various real-life applications, leading to program modules built on top

of the above experimental coordination testbed system.

Exploratory research: coalgebraic models of computation
Further development of coalgebra as a unifying mathematical framework for (transition and dynamical) systems, with emphasis on the coordination of coalgebras, as well as probabilistic transition systems.

Evolutionary Computation and Applied Algorithmics (pilot theme)
Theme leader: J.A. La Poutré

Evolutionary algorithms
Design of evolutionary algorithms for management-related problems like optimisation problems, strategic problems, and economic modeling and strategies.

Neural networks
Classification, visualisation and (strategic) prediction of data, including remote sensing, spatial stochastics, and artificial neural networks, as well as neural vision.

Discrete algorithms
Design of efficient algorithms for on-line optimisation problems underlying various management and design problems in computer systems and networks, as well as the use of quality of service for on-line scheduling for multimedia processes.

Modelling, Analysis and Simulation
Cluster leader: C.J. van Duijn

Environmental Modelling and Porous Media Research
Theme leader: J.G. Verwer

Numerical algorithms for air quality modelling
Numerical modelling of the long range transport and chemical exchange of atmospheric air pollutants, involving advection schemes, stiff chemistry ODE solvers, operator splitting, grid generation/adaptation and time stepping methods for PDEs, including implementations on supercomputers and parallel computers.

Numerical algorithms for surface water quality modelling
Design of parallel numerical methods for the simulation of water pollution (calamitous releases), the marine eco-system, dispersion of river water, sediment transport, etc.

Partial differential equations in porous media research
The modelling of transport processes in the subsurface, with emphasis on the analytical study of the governing partial differential equations.

Exploratory research: discretization of initial value problems
Fundamental research – though possibly application driven – into theoretical numerical questions and properties, currently focussing on special time integration methods for stiff, second-order initial-value problems, and on the Method of Lines for temporal discretisation.

Industrial Processes
Theme leader: P.W. Hemker

Computational fluid dynamics
Computation of flows in gases, liquids, plasticly deforming solids, or combinations of these (multiphase flows) for industrial applications. Current research includes advanced discretization methods for systems of nonlinear conservation laws, multigrid and sparse-grid solution methods, local grid adaptation and distributed computing, with emphasis on sparse-grid algorithms for 3D problems.

Parallel codes for process simulation
Design of parallel codes for the simulation of a wide variety of phenomena occurring in industrial processes like circuit analysis, control engineering, viscoelasticity, rheology, etc. The emphasis is on codes for initial-value problems for ordinary and implicit differential equations and Volterra integro-differential equations.

Plasma physics simulation
Development, implementation and analysis of new parallel iterative methods for the computation of eigenvalues and eigenvectors of generalized eigenvalue problems, where the matrices are very large, sparse and (complex) non-Hermitian. The algorithms are tested on problems coming from plasma physics and other application areas.

Discontinuous dynamical systems
Research into the interaction of ‘continuous’ and ‘discrete’ dynamics, which arises when regime switches, occurring for example in electrical networks, in robotics and in processes with a discrete controller, are modelled as discrete events.

Computational number theory and data security
The study of algorithms for factorization and primality testing, for computing discrete logarithms, and (derived from these) for the solution of large, sparse systems of linear equations over finite fields. In addition, low-discrepancy sequences for the evaluation of high-dimensional integrals are studied, in connection with the valuation of financial derivatives like options and mortgage backed securities.

Modelling of processes in biology and chemistry
Parameter identification in non-linear dynamical systems describing processes in biology and chemistry, modelling and numerical bifurcation analysis of structured populations, modelling of infectious diseases, and simulation and solution of evolutionary dynamics.

Exploratory research
The study of singular perturbation problems arising from the numerical computation of strongly convective flows, with emphasis on uniform convergence in 2D domains, in particular special non-uniform grid requirements and adaptive meshes. In addition, the derivation of uniform asymptotic expansions of integrals for large parameter values in non-standard cases, in order to obtain numerical approximations of special functions.

Mathematics of Finance (pilot theme)
Theme leader: J.M. Schumacher

In view of the pilot's current size no division in subthemes has been made so far. At present the following projects are underway:
- Statistics for random processes with applications to mathematical finance
- Error estimates in Quasi Monte Carlo computations
- Monte Carlo methods for derivatives with early exercise opportunities
- Construction of low-discrepancy sequences
- Event risk modelling.

Information Systems
Cluster leader: M.L. Kersten

Data Mining and Knowledge Discovery
Theme leader: A.P.J.M. Siebes

Data mining
For the data selection phase, research is focussed on structure in data, for example time-series, geographical data, multi-valued attributes, and non-universal relations, and on (non-)random samples. For data mining proper, the emphasis is on model-representation and search. An important aspect is the reformulation and generalization of well-known data mining algorithms in the KESO formalism.

Database architecture
Research on main memory and parallel DBMS architectures, emphasizing facilitation of data mining and efficiency; experiments with parallel data mining by porting MONET to systems like PARSYTEC and SP2; further development of performance monitoring and prediction software; (multi-)query optimization and support of special data types (e.g., time-series databases and picture databases), in connection with efficient data mining.

Multimedia databases
The objective is to achieve efficient storage and retrieval of multimedia data, such as pictures, video and audio, in particular by using feature detectors to simplify and speed-up multimedia data query.

Multimedia and Human-Computer Interaction
Theme leader: D.C.A. Bulterman

Multimedia authoring systems
Investigation of models and implementation environments for the development of complex multi-/hypermedia documents. Particular attention is paid to the author-controlled synchronization of independent data streams for presentation on heterogeneous (networked) user workstations.

Distributed multimedia applications
Presently the focus is on applications in distributed international electronic commerce, in particular the delivery of complex multimedia documents over high-speed (ATM) network infrastructures, and in remote access to publicly available information, as a core member of W^{3}C's Synchronized Multimedia working group.

Interactive structured documents
For authoring structured documents a good design of the supporting architectures is essential. To this end means are investigated to combine logical structure with constrained presentation of information in a variety of applications. Efficient cooperation of diverse software components in an interactive structured document requires a theory for the propagation of incremental changes across changes in data representation.

Interactive Information Engineering

Theme leader: P.J.W. ten Hagen

Information engineering framework
The focus is on the software tools and algorithms for interactive information environments. Results of the new multimedia programming ISO standard PREMO will be used in, e.g., facial animation, and graph visualization techniques are used for algorithms and tools for the display of and interaction with abstract information.

Digital libraries
Development of a thesaurus based search facility in the vast on-line collections of mathematical and computer science research papers, of tools for creating dictionaries of keyphrases, through which user-friendly access to the real material is achieved, and of a standard dictionary for mathematics on behalf of authors and publishers. Another goal is the interactive enrichment of on-line mathematical texts.

Applied logic and interactive books
Application of logic, viewed as a science of information, to key issues in the processing and engineering of information, and the dissemination of this view of logic. Production of interactive textbooks where logic is presented as a core part of an emerging science of information processing and information flow analysis.

Facial animation
Production of a prototype system capable of capturing facial emotional expressions as enacted by a speaking performer and of reproducing user-controlled transformations of those expressions as part of information presentations.

Quantum Computing and Advanced Systems Research

Theme leader: P.M.B. Vitányi

Quantum computing
Exploratory study of algorithms and systems based on quantum mechanical principles (in particular previous work on quantum cellular automata and quantum information retrieval algorithms). The aim is to enable the realization of practical quantum computation.

MDL Learning and evolutionary computing
Design, implementation, and comparative analysis of a series of practical applications of machine learning techniques. Applications include automatic grammar generation from large text corpora and comparative evaluation of predictive accuracy of MDL and new forms of stochastic complexity, and GP learning of neural network governed robot locomotion and general techniques improving speed and storage requirements of GP implementations.

Advanced algorithms and systems
Design and analysis of algorithms for distributed and parallel systems. Possibilities of future systems are identified by exploiting fundamental mathematical techniques of (Kolmogorov) complexity theory. A major item is descriptional complexity leading to the ‘incompressibility method’ and ‘learning by compression’. Also mobile and nomadic computing and communication are considered.

INTERNATIONAL AND NATIONAL PROGRAMMES

This chapter summarizes the major national and international projects in which CWI participates.

The following data are given for each project:
- title,
- period,
- cooperation with other institutes,
- special role of CWI (if any),
- CWI project leader(s).

European Programmes

ESPRIT

MERCURY (20089): Performance Management of Commercial Parallel Database Systems
January 1996 – December 1998
ICL, IFATEC, ING, Herriot-Watt Univ.
M.L. Kersten

KESO (20596): Knowledge Extraction for Statistical Offices
January 1996 – December 1998
National Statistical Offices of Finland, Greece (via FORTH), and The Netherlands, Infratest Burke (D), Data Distilleries (NL), GMD (D), Univ. Helsinki (FIN)
A.P.J.M. Siebes

CHAMELEON (20597): An Authoring Environment for Adaptive Multimedia Documents
November 1995 – November 1998
CLRC (UK), Epsilon SA (GR), Cartermill International (UK), Comunicacion Interactiva (SP), Egnatia Epirus Foundation (GR), Cycnos Systèmes Ouverts (F)
D.C.A. Bulterman

SAGA (21871): Scientific Computing and Algebraic Abstractions
September 1996 – September 1997
Univ. Bergen, Univ. Swansea
J. Heering

DELOS (21057): ERCIM Digital Library
March 1996 – March 1999
Elsevier, Univ. Michigan, all ERCIM Institutes
F.A. Roos

CONFER III (21836): Concurrency and Functions: Evaluation and Reduction
November 1996 – November 1999
INRIA, ENS, CNET, ICL, KTH, Universities of Bologna, Cambridge, Edinburgh, Pisa, Sussex and Warwick
J.W. Klop

COTIC (23677): Concurrent Constraint Programming for time-critical applications
1997 – 2000
Universities of Utrecht, Pisa, Lisbon and Kent, SICS, CR&T
K.R. Apt

COORDINA (24512): From Coordination Models to Applications
1997 – 2000
INRIA, Xerox, Univ. Leiden, 8 European Universities, Hollandse Signaalapparaten
J.J.M.M. Rutten

NeuroCOLT II (27150): Neural and computational learning
1998 – 2000
11 universities across Europe
P.M.B. Vitányi

COST

Verification and Validation Methods for Formal Descriptions (247)
1993 – 1997
Philips, Universities of Utrecht, Eindhoven, Amsterdam, Nijmegen, Groningen, Twente + institutions in 18 European countries
J.F. Groote

MAST Marine Science and Technology

MMARIE: Application of High Performance Computing Techniques for the Modelling of Marine Eco Systems
February 1995 – February 1998
Univ. Leuven, Delft Hydraulics, Univ. Southampton, IFREMER, CRS Cagliari, Univ. Hamburg, Univ. Liège, Univ. Delft, RIKZ, CETIIS, Univ. Bradford, Hydraulic Research Wallingford, Proudman Oceanographic Laboratory Bridston, UP de Catalunya.
P.J. van der Houwen

ACTS

SEMPER (AC026): Secure Electronic Marketplace for Europe
September 1995 – September 1998
Cryptomathic (DK), DigiCash (NL), Eurocom Expertise (GR), Europay International (B), FOGRA Forschungsgesellschaft Druck (D), GMD (D), IBM European Networking Center (D), Intracom (GR), KPN Research (NL), Otto-Versand (D), r3 security engineering (CH), SEPT (F), Univ. of Freiburg and Hildesheim (D)
D.C.A. Bulterman

TELEMATICS

DACCORD (TR1017): Development and Application of Coordinated Control of Corridors
January 1996 – January 1999
Hague Consulting, Univ. Delft, Univ. Lancaster, TNO, RWS, Univ. Naples, CSST, Autostrade Italia, INRTS, Ile de France, Ville de Paris, Univ. Crete, TCU
J.H. van Schuppen

TMR/HCM

EXPRESS: Expressivenes of languages for concurrency (CT93-0406)
1994 – 1997
Univ. Utrecht, SICS, Univ. Genova, Univ. Rome (La Sapienza), Univ. Hildesheim, Univ. Amsterdam, INRIA, GMD, Univ. Sussex, Univ. Nijmegen
Coordinators: J.F. Groote/J.W. Klop

DIA: Digital Identification and Authentication (CT94-0691)
1995 – 1997
ENS, Univ. Salerno, Univ. Saarland, Univ. Aarhus
Coordinator: R. Hirschfeld

ERCIM advanced databases technology network
1994 – 1997
M.L. Kersten

DONET: Discrete Optimization: Theory and Applications
April 1998 – March 2000
London School of Economics and Political Sciences, Univ. Pierre et Marie Curie (Paris), Rheinische Univ. Bonn (D), CNR (I), Univ. Lisbon (P), ALMA (F), DASH (UK), Ecole Polytechnique Fédérale de Lausanne (CH)
A. Schrijver, A.H.M. Gerards

ERNSI: Systems Identification
March 1998 – March 2002
Royal Tech. Univ. Stockholm, Tech. Univ. Wien, CNR, Univ. Leuven, INRIA, Univ. Rennes, Univ. Cambridge, Univ. Linköping
J.H. van Schuppen

Statistical inference for stochastic processes (CT92-0078)
1993 – 1997
Universities of Paris VI, Berlin, Aarhus and Freiburg, INRIA
K.O. Dzhaparidze

INTAS

Network Mathematical Methods for Stochastic Discrete Event Systems
1994 – 1997
O.J. Boxma

Mathematical methods for stochastic discrete event systems
1997 – 2001
Univ. Moscow, Univ. Novisibirsk, Univ. Cambridge, Univ. Braunschweig, INRIA
M.S. Keane

Symmetry and cohomology approach to equations of mechanics and mathematical physics
October 1997 – March 1999
Moscow Institute for Numerical Economy, Moscow State Univ., Univ. Twente, Univ. Salerno
M. Hazewinkel

ERETIMA – English-Russian enriched Thesaursus in Mathematics
November 1997 – April 1999
Univ. München, Yaroslav State Univ., Russian Academy of Sciences, Steklov Institute of Mathematics
M. Hazewinkel

National Programmes

SION *(Netherlands Computer Science Research Foundation)*

COCO – Computational Intelligence for Constraint Logic Programming
1997 – 2002
ERCIM partners in ERCIM WG on Constraints
K.R. Apt

ModAlg – Modular Algebra
September 1996 – August 1997
Univ. Amsterdam
J. Heering

A modular toolset for μCRL
1997 – 2000
Univ. Utrecht, Univ. Eindhoven, Univ. Amsterdam, Univ. Nijmegen, Univ. Groningen, Univ. Twente, Philips
J.F. Groote

Protocols, Reference models and Interaction schemes for Multimedia Environments
1997 – 2000
J.F. Groote

PROMACS – Probabilistic methods for the analysis of continuous systems
1998 – 1999
Univ. Eindhoven, Free Univ. Amsterdam
J.J.M.M. Rutten

COLA – Formal methods and refinement for coordination languages
1997 – 2002
Univ. Leiden, Hollandse Signaalapparaten
J.J.M.M. Rutten, F. Arbab

Dynamic Algorithms for On-line Optimization
1997 – 2000
Philips Research
J. La Poutré

Quality of Service for Multimedia Systems
1997 – 2001
Philips Research
J. La Poutré

AMIS – Advanced Multimedia Indexing and Searching
July 1997 – June 2001
Data Distilleries b.v., ICL, IFATEC, ING-Group, Tandem, Herriot-Watt Univ., Universities of Twente, Eindhoven, and Amsterdam
M.L. Kersten

Quantum Computing
1997 – 2001
Univ. Amsterdam
P.M.B. Vitányi

From ideas to reality – Implementing cryptography
1994 – 1998
L.G.L.T. Meertens

WINST – Themes for collaboration in mathematics and computer science
1994 – 1997
Universities of Nijmegen and Eindhoven
J.F. Groote, J.W. Klop

DEGAS – Design theory for autonomous databases
1993 – 1997
A.P.J.M. Siebes

Incremental parser generation and disambiguation in context
1993 – 1997
Univ. Amsterdam
D.J.N. van Eijck

MDL Neurocomputing
1994 – 1998
P.M.B. Vitányi

Equational term graph rewriting
1994 – 1998
J.W. Klop

Generic tools for program analysis and optimization
1994 – 1998
P. Klint

Checking verification of concurrent systems with type theory tools
1994 – 1998
Univ. Utrecht
J.F. Groote

Constraints in object-oriented interactive graphics
1994 – 1998
Univ. Eindhoven
P.J.W. ten Hagen

Constraint-based graphics
1994 – 1997
Univ. Eindhoven
F. Arbab

Cryptography, learning and randomness
1994 – 1997
Univ. Amsterdam
P.M.B. Vitányi

Classifying proof techniques for propositional logic
1994 – 1999
Univ. Delft
J.F. Groote

Scientific Visualization – from data visualization to interactive exploration
1995 – 1997
Univ. Delft, Free Univ. Amsterdam
J.J. van Wijk

Parallel declarative programming: transforming logic programs to lazy functional programs
1996 – 1998
K.R. Apt

Theorem proving and dynamic logic
1996 – 1998
ICCL, Univ. Stuttgart
D.J.N. van Eijck

SWON *(Netherlands Mathematics Research Foundation)*

Nonexpansive mapping and their asymptotics
March 1997 – February 2001
Univ. Utrecht, Univ. Delft, Univ. Leuven, Univ. Cambridge, Free Univ. Amsterdam
M.S. Keane

Sparse grid methods for time-dependent PDE problems
December 1997 – November 2001
UU/IMAU, RIVM, KNMI, TNO, Univ. Iowa
J.G. Verwer, B. Koren

Parallel computational magneto-fluid dynamics
February 1997 – January 2001
Univ. Utrecht, FOM Inst. Plasma Physics
H.J.J. te Riele

CAM – Computational number theory and data security
March 1997 – February 2001
Univ. Groningen, Univ. Leiden, Univ. San Rafael, Australian Nat. Univ. Canberra, Univ. Macquarie at Sydney, Univ. Bordeaux, Univ. Saarbrücken, Citibank
H.J.J. te Riele

Statistics for random processes with applications to mathematical finance
June 1998 – May 2002
Free Univ. Amsterdam
K.O. Dzhaparidze

Large random systems and combinatorial probability
September 1993 – December 1997
Cornell Univ., Math. Inst. Hungarian Ac. Sc., Univ. Utrecht, Delft, Leuven, Cambridge (UK), Chalmers (Gothenburg), Rome
J. van den Berg

WINST – Themes for collaboration in mathematics and computer science
1994 – 1997
Universities of Nijmegen and Eindhoven
J.F. Groote, J.W. Klop

Statistical properties of movements in the plane
February 1996 – January 2000
Math. Inst. Hungarian Ac. Sc., Universities of Utrecht, Delft, Leuven, Cambridge, Chalmers Gothenburg, and Rome
M.S. Keane

Integer polyhedra and binary spaces
June 1996 – May 2000
A.M.H. Gerards

Discontinuous dynamical systems
August 1996 – July 2000
Univ. Groningen, Twente, Brabant
J.M. Schumacher

STW *(Foundation for the Technical Sciences)*

Parameter identification and model analysis for non-linear dynamic systems
1993 – 1997
Univ. Delft, Heidelberg, Beer Sheva, Nova Scotia
P.W. Hemker

Parallel codes for circuit analysis and control engineering
1993 – 1997
Univ. Amsterdam, Philips, AKZO-Nobel
P.J. van der Houwen

FASE – Facial animation
1997 – 1999
Univ. Delft, Philips, NOB, Institute for the Deaf, KPN Research
P.J.W. ten Hagen

Wavelets: analysis of seismic signals
1996 – 1999
Universities of Delft, Eindhoven, and Groningen, Shell, KNMI, MARIN
N.M. Temme

Special NWO projects

CIMS – Parameter estimation for random sets
April 1997 – March 1999
Univ. Western Australia (Perth)
M.C. van Lieshout

Plasma simulation
1993 – 1998
Univ. Utrecht, FOM Inst. Plasma Physics
H.J.J. te Riele

Singular perturbation problems in strongly convective flows
1992 – 1997
Univ. Nijmegen, Inst. Math. & Mech. Ekaterinenburg, Univ. Novosibirsk
P.W. Hemker

Applied logic dissemination (Spinoza)
1997 – 2002
ILLC
D.J.N. van Eijck

NCF

CIRK – Mathematical modelling of global transport and chemistry of trace constituents in the atmosphere
1994 – 1997
Univ. Utrecht, RIVM, KNMI, TNO, Univ. Iowa
J.G. Verwer

Cray – Long Term Ozone Simulation
1994 – 1997
Delft Hydraulics, Min. Public Works-RIKZ, Univ. Delft, IFREMER, Univ. Leuven
P.J. van der Houwen

Virtual reality
1996 – 1999
ECN, Univ. Delft, Univ. Amsterdam
R. van Liere

ICES HPCN Programme

HPCN for Environmental Applications
1996 – 1999
Univ. Delft, Delft Hydraulics, TNO
J.G. Verwer

IMPACT – HPCN for Financial Services
1996 – 2000
ING, Univ. Amsterdam, Univ. Twente, Getronics, Univ. Delft, CAP Volmac, Data Distilleries, BIT by BIT
A.P.J.M. Siebes, M.L. Kersten

NICE – Computational Fluid Dynamics
1996 – 1999
Delft Hydraulics, CUNY Brooklyn
F. Arbab

High Performance Visualization
1996 – 1998
ACE, CAP Volmac, Arcobel, TNO
R. van Liere