



Centrum voor Wiskunde en Informatica

ANNUAL *REPORT*

'94



Centrum voor Wiskunde en Informatica

Two large, thick black arcs are positioned around the title. One arc starts on the left and curves upwards towards the right, while the other starts on the right and curves downwards towards the left, framing the text.

ANNUAL *REPORT*

'94

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 part of 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 member of ERCIM, the European Research Consortium for Informatics and Mathematics.

General Director

G. van Oortmerssen

ERCIM



Copyright © 1995 Stichting Mathematisch Centrum
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

CONTENTS

Introduction	4
Organization	10
Research Highlights	
Mathematical Epidemiology: Modelling the Force of Infection	12
Robust Control	15
Parameter Estimation in Dynamical Systems	18
Hybrid Systems	22
An Electronic Wallet for Digital Money	27
Computational Steering	31
ATM Networks: a new Infrastructure for Research Computing	36
Computing Equipment Resources	40
Financial and Other Data	42
CWI Research Programmes	45
International and National Programmes	49
Research Staff	55
Advisory Committees CWI	57

This Annual Report is complementary to the Jaarverslag SMC (in Dutch), which concentrates on SMC's National Activities in Mathematics. A complete overview of CWI's research activities, as well as SMC's Financial and Social Reports (in Dutch), are also available.

INTRODUCTION

The Research Scene

CWI operates as a leading research institute in a dynamic environment with abundant challenges. To achieve its mission within this environment, CWI must be alert for external developments on the domestic and international stages. A constant source of concern is the continuing *decline in spending* on R&D, whereby The Netherlands is in jeopardy of lagging behind in Europe. Meanwhile, the funding mechanism for fundamental research shifts from basic funding towards *project funding*, to be obtained in *competition* with others. In addition, there is considerable political pressure on research organizations to tune their research programmes more to the needs of *society*, and to increase the application of results by intensified interaction with *industry*. Lastly, the increasing complexity and 'cross-border' aspects of many problems demand a *multidisciplinary* approach and cooperation on an *international* scale.

MOBILE

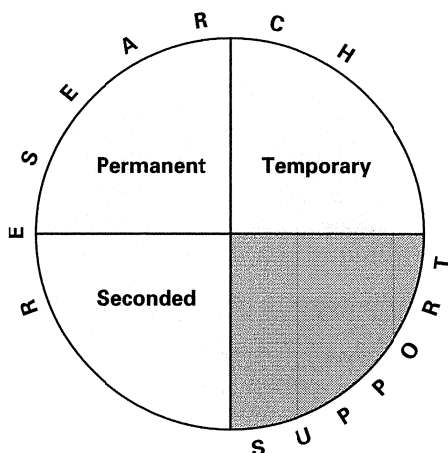
Over the past several years CWI has responded successfully to all these developments and factors. The action plans underway meld seamlessly with a mission statement virtually unchanged in almost fifty years - to conduct fundamental research complemented by the transfer of knowledge to the community at large. Indeed, CWI's

actual research topics are derived from issues in the public domain. In 1992 CWI's mission statement was updated and elaborated in *Towards New Equilibria: MOBILE*. The MOBILE document set out CWI's research strategy for the 1990s, the key words being innovation, quality and flexibility. These principles are in full accordance with the policy of CWI's major sponsor, NWO. During recent years, much effort was paid to reconfiguring CWI's research strategy and organisation in line with the principles set out in MOBILE. This resulted in more management responsibility for department heads, a lean and professional support staff, and improved management information systems. In 1994 CWI emerged as an energetic institute, with a high scientific output, open communication with the scientific world as well as with society, and in a healthy financial position.

During 1994 CWI was once again successful in securing research projects, at home and internationally. For example, CWI participated in 29 European Union projects and 19 projects funded by the NWO foundation SION (Netherlands Computer Science Research Foundation). For a full overview, please see the chapter on International and National Programmes (page 49).

CWI's research focuses on a limited number of themes, selected on the basis of scientific interest as well as societal needs. The transfer of the knowledge amassed is boosted through strategic alliances with institutes for applied research and research laboratories in industry. In the field of contract research for government and the private sector, CWI seeks long term joint-projects with the customer. A good example here is the research with Dutch Railways/NS into the 'rail timetable of the 21st century', which realised significant progress in 1994. This project resulted from initial contacts made back in 1992 at our annual 'CWI in the Market-place' event. This yearly event, held for the third time in 1994, has proved to be of great value in establishing a dialogue with potential users of scientific results and attracts attendance of many representatives of

CWI Personnel:
150 fte + 50 fte
seconded.



governmental, academic and business organisations.

Alliances

Much of CWI's research is multi-disciplinary in nature. Many projects here involve mathematicians and computer scientists working as a team, like in the SION-SMC project WINST (cooperative themes on mathematics and informatics) and in ACELA (Architecture of a Computer Environment for Lie Algebras); the latter project is funded by the NWO Foundation for the Technical Sciences (STW). Where other disciplines are involved, CWI works closely across a broad spectrum with other institutes. An excellent example is the *Mathematics & the Environment* programme which has been underway for some years, and in which CWI cooperates with several partners, including the National Institute of Public Health and Environmental Protection (RIVM), the National Veterinary Institute (ID-DLO), the Public Works department (RWS), the Royal Netherlands Meteorological Institute (KNMI) and Delft Hydraulics (WL). Another example is the project *Scientific Visualisation*, in which the Dutch national institute for energy research, ECN, is a partner.

Universities


CWI attaches much value to cooperation with research groups at universities. Collaboration is increasingly taking place in joint projects, often sponsored by NWO foundations such as SMC, SION and STW. In addition, ad-hoc partnerships are formed to perform contract research for the private sector. In 1994, two such projects were carried out with the University of Utrecht and Tilburg University.

Since academic research is structured more and more within Research Schools, the relations with these schools is of increasing interest. In order to foster cooperation, CWI signed agreements with the Research School in Logic, with the Thomas Stieltjes Institute for Mathematics, and with the Euler Institute for Discrete Mathematics and Applications. Similar agreements with the Dutch Institute for Systems and Control and the Institute for Programming Research and Algorithmics (IPA) are in preparation.

International Cooperation

CWI early recognized the importance of international cooperation. Evidence includes the set up

in the late 1980s of ERCIM, the European Research Consortium for Informatics and Mathematics; a twelfth member, Hungary, joined in 1994, and preparations were made for the entry of Switzerland. Furthermore, CWI can point to significant participation in European Union research programmes, notably ESPRIT and HCM (research networks). CWI coordinates five of the 29 European projects in which it is involved. Due to delays in the EU's Fourth Framework Programme (4FP), new projects can only be expected to start up in autumn 1995. The shift in emphasis away from long term research towards industrial applications and end-user involvement is a source of considerable concern. During the report year CWI started to prepare project pro-



CWI researcher S. Pemberton was appointed as editor-in-chief of the Bulletin of ACM's Special Interest Group on Computer-Human Interfaces.

posals for the Fourth Framework Programme in various areas, but due to the lower interest in long term research and increased competition, a lower yield is to be expected.

The *Research Institute for the Applications of Computer Algebra (RIACA)* also started up officially in 1994. RIACA is a working alliance between CWI, the Netherlands Computer Algebra Foundation (CAN) and the Kurt Gödel School of Linz, Austria (formerly Research Institute for Symbolic Computation RISC). This startup was possible thanks to funding for initial investments granted by the ministry of Education, Culture and Science.

HPCN

Research in the field of High Performance Computing & Networking (HPCN) requires cooperation of scientists from various fields, because of the multi-disciplinarity of problems involved, and

in order to share expensive advanced facilities. To this end, CWI and several universities - notably the University of Amsterdam - work together in IC³A (Interdisciplinary Centre for Complex Computer facilities Amsterdam). Under an agreement with both Amsterdam universities and IBM, CWI conducts experimental research on the SPI parallel computer at SARA (Academic Computer Centre Amsterdam); areas covered include numerical algorithms, graphic interaction and distributed databases.

The HPCN & the Environment consortium, already mentioned in passing, focuses on *High Performance Computing for Long-Range Transport Problems in the Atmosphere, Surface Water and Groundwater*. CWI works with six partners here: Delft University of Technology (TUD), the National Coastal and Marine Management Institute (RWS/RIKZ), Delft Hydraulics (WL), the National Public Health and Environmental Protection Institute (RIVM), Utrecht University Institute for Marine and Atmospheric Research (IMAU), and the Netherlands organization for applied scientific research (TNO). In PARES (Parallel Research Consortium) CWI's partners are the universities of Utrecht, Groningen and Leiden. Important research support in the area of HPCN comes from CRAY Research Grants, due to which CWI has access to an advanced parallel CRAY T3D computer in Switzerland. Over the

next several years, CRAY grants will enable the highly parallel and vectorized one-pollutant model developed by CWI in 1994, to be expanded into a multi-pollutant model, bringing the realistic modelling of certain environmental issues nearer.

During the course of the report year, CWI and a number of partners submitted proposals for the national HPCN programme financed by the economics ministry. Two CWI proposals in the field of numerical mathematics were approved under NWO's MPR (Massively Parallel Computing) priority programme; these were *Parallel computing in air pollution problems*, and *Parallel aspects of 3D multilevel and domain decomposition techniques for PDEs appearing in technical applications*.

Spin-offs

Spin-off effects have been an important aspect of CWI strategy since the late 1980s. Spin-off companies can be instrumental to the development and commercial exploitation of applications, activities which are outside the scope of CWI's mission. New start-up companies can create employment opportunities for (temporary) CWI staff, and can provide important feed-back from practice, thus giving new impulses to CWI's research programme. In 1989, CWI was behind establishment of the CAN Foundation, which has

CWI is involved in several research projects in the traffic area: (a) traffic lights replacement strategies, (b) railroad tables for the 21st century, (c) motorway traffic control.



the exploitation of computer algebra software as a major goal. In 1990, CWI's Cryptography research group spawned the DigiCash company to develop and license payment technology products - chipcard, software only, and hybrid - that both show the true capability of technology to protect the interests of all participants and are competitive in the market. During the report year, a number of ex-CWI staff formed the *General Design* company devoted to user-friendly user interfaces, with a special focus on Internet. These operations are clustered near CWI at the Watergraafsmeer Science Park. This would be an ideal site for a similar spin-off company planned in the Data Mining area. In all cases, there is an ongoing partnership between CWI and the spin-off companies.

Contract Research

In dealing with the private sector, CWI is aiming for long-term, sustainable cooperation, in which CWI's contribution consists of performing high-quality, innovative research. Significant contract research projects and clients in 1994 included:

- verification of audio-protocols in Philips stereo equipment (resulted in 'best paper award' at a recent leading congress);
- statistical consultation for the Public Works department and for Coopers & Lybrand;
- development of mathematical methods for the design of Dutch Railways timetables;
- traffic light replacement strategies and optimal use of bus stations, jointly with the University of Utrecht and Tilburg University, for Nederland Haarlem BV;
- theoretical research into traffic control algorithms for motorway networks within the DYNA project completed in 1994, as part of the European Union DRIVE programme.

As evidenced by the last three, CWI also has close involvement in traffic and transportation-related research.

Our annual 'CWI in the Market-place' is a proven formula for getting close to prospect in the public and private sectors. In 1994, the one-day event broke previous records, with 120 participants. A popular new item was the colloquium on *Questions from IT-practice*. In a series of five separate sessions, senior people representing software houses, banking and other sectors confronted researchers with their most urgent problems.

Electronic Superhighway

The report year also saw CWI - as a leading computer science research institute - getting closely involved in developments in-and-around the *electronic superhighway*. To synergise our considerable expertise in multimedia, networks and telematics, effective 1 January 1995, research in the Computer Systems & Telematics department was redeployed to the Algorithmics & Architecture department. As part of the *Interoperable Multimedia Systems* project, generic interoperability is studied in the context of World Wide Web, e.g. how present WWW-protocols need to be adjusted to enable distributed computation. CWI's organisational input included the appointment of Steven Pemberton as coordinator of the ERCIM World Wide Web Working Group (W4G) which was established in 1994.

Training

CWI's mission of knowledge transfer is partly achieved by training of young researchers. In so doing, an ongoing contribution is made to the renewal of expert human resources in The Netherlands. No less than 15 CWI researchers were awarded PhDs at Dutch universities during the report year. Three of them, A.J. Cabo, M.N.M. van Lieshout and P.F.M. Nacken, performed their research in the Image Analysis group, led by A.J. Baddeley, who took a professorship in Perth, Western Australia in early 1994. He was succeeded by M.S. Keane, joining from Delft University of Technology, while maintaining a

Digital two-dimensional lung-section.



part-time professorship there. New research in the image analysis field started up after Dr Keane's arrival includes modelling of lungs, a joint project with Rotterdam's Daniel den Hoed Clinic and Delft University of Technology.

ESPRIT

CWI has a prominent position in three ESPRIT projects, PYTHAGORAS, CAFE, and COMPULOG which we coordinate, and in MADE. PYTHAGORAS involves development of a Software Test Pilot to assess the quality of new database management systems. ING Bank has now joined in PYTHAGORAS as a partner, and the project has been granted a 12-month extension to research concrete applications. The objective of CAFE (Conditional Access for Europe) is to give Europe an open and secure electronic payment system with potential for expansion including the use of personal attributes like passports and house keys. In 1994, preliminary work was done towards a concrete practical test of the electronic wallet already developed in the project. This test will be carried out at several EU agencies in the first half of 1995. CWI's cryptography group will also be involved in activities of the National Chipcard Platform Foundation, via CAFE. The ESPRIT Basic Research Action COMPULOG II aims at studying extensions of logic programming to improve its knowledge representation and problem solving capabilities. Extensions consist of related paradigms from the areas of mathematical logic (automated theorem proving), deductive database systems, and logical aspects of artificial intelligence. Two large software projects are part of the COMPULOG effort: Goedel - a tool for declarative programming, and ConceptBase - a deductive object-oriented database system. Finally, MADE (Multimedia Application Development Environment) will provide the basis for multimedia object creation and authoring by introducing advanced, novel object-oriented techniques. CWI coordinates an effort to develop an international ISO standard for multimedia programming PREMO (Programming Environment for Multimedia Objects).

Multiple Computing Agents

CWI is tackling the research field of multiple computing agents in two ways. Firstly, the large body of expertise in machine learning amassed by CWI in recent years, provides the basis for research into neurocomputing (artificial neural networks, with applications in computer vision and speech recognition) and gencomputing (genetic algorithms which use natural selection mechanisms to solve optimization problems), with computational linguistics forming a concrete test environment. Secondly, at the end of 1994

joint research started with the University of Utrecht into methods and tools, based on genetic algorithms and neural networks, which make efficient use of parallel computers, with applications in image processing, planning and data analysis. This research focuses on hybrid systems, interactive planning tools and experiments on IC³A's massively parallel ParsyTec computer.


Research and Practice

Further evidence that good, fundamental science can be of direct practical interest is given by the following examples:

- The Ministry of Public Works commissioned CWI's expertise in statistical *bootstrap resampling* methods in locating and analyzing oil spillage in the North Sea. The same methodology was applied at the international accountancy firm, Coopers & Lybrand.
- Some years back, CWI's highly successful course on *wavelets* drew more than 100 participants. In 1994, CWI made a successful project application to the STW foundation (awarded early 1995). The research will be focused on signal processing problems in geophysics - a traditional area of concern for the oil industry.
- As part of the SION/SMC project *WINST* (Dutch acronym meaning 'profit', standing for Mathematics and Informatics Cooperation Themes) CWI is researching ways to convert its fundamental results realised in the past into practical applications, by developing interfaces between Term Rewriting Systems, Computer Algebra and Proof Checking.
- The *Dynamical Systems Laboratory (DSL)*, which is located at CWI, was officially opened in mid-1994. DSL's mission is to provide an interactive research environment for the study of dynamical systems by computers. DSL is supported by the NWO Priority Programme 'Mathematical Aspects of Non-linear Dynamical Systems' DSL maintains user software and performs research in symbolic computation and bifurcation analysis. Applications can be found in many fields. A new environment for numerical study of dynamical systems is presently under development. Workstations specially tailored for the study of dynamical systems were installed at CWI and several Dutch universities.
- In the *Data Mining* field, CWI was commissioned by Philips to examine reject patterns in video screen production.

Change in Management

The report year saw a change in the management configuration at CWI. The managing director, Dr. G. van Oortmerssen was appointed general director, effective 1 May 1994. The scientific director Prof. P.C. Baayen, stepped down at the end of the year. In December, his departure was marked by a scientific symposium in Amsterdam, with speakers including A. Schrijver (CWI/University of Amsterdam), J.F.A.K. van Benthem (University of Amsterdam), D. Scott (Carnegie Mellon, USA) and A. Bensoussan (INRIA, France). The proceedings were printed in a Liber Amicorum, with the appropriate title 'From Universal Morphisms to Megabytes: a Baayen Space Odyssey', edited by Krzysztof Apt, Lex Schrijver and Nico Temme. It includes some 40 contributions by researchers from within and beyond CWI. In his capacity as Vice President of ERCIM, Prof. Bensoussan took the opportunity to announce that Prof. Baayen had been appointed Honorary President of ERCIM. Prof. Baayen played a key role in the establishment and expansion of this consortium and was ERCIM's first president from 1991-1994. An annual 'Cor Baayen Award' was also instituted for promising young researchers in ERCIM. At the symposium, Prof. Baayen was also decorated with a Royal award (Officer in the



Professor Cor Baayen, Scientific Director of SMC 1980-1994, during his farewell speech (Amsterdam, 20 December 1994).

Order of Orange-Nassau). Recently, the Royal Netherlands Academy of Arts and Sciences awarded him its 1995 Medal for his considerable services to science in this country. Prof. Baayen's special qualities helped secure CWI's prominent position today. Building on these foundations we can look to the future with great confidence.

G. van Oortmerssen, General Director

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 SMC and 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 and industry in particular.

CWI's research is carried out in six scientific departments. There is considerable inter-departmental collaboration, for example in the ongoing multidisciplinary programmes *Mathematics & the Environment* and *Multimedia*. Researchers at CWI are supported by state-of-the-art computer facilities and a well equipped library of national importance and, hence, ideally prepared to handle the dynamic and interdisciplinary demands of present day research.

Besides being responsible for CWI, SMC also finances research projects in mathematics at Dutch Universities (National Activities in Mathematics). These activities comprise a total of almost sixty projects. CWI can obtain project


funding from this source as well and cooperates in various projects with university researchers.

SMC is administered by a Board of Trustees. The Dutch mathematics community is represented in the Board by three members, appointed on recommendation of the section Mathematics of the Royal Netherlands Academy of Arts and Sciences. Another three members are appointed in consultation with the Netherlands Computer Science Research Foundation (SION). This reflects the importance of the relationship between CWI and SION (CWI participates in many research projects sponsored by SION, and SION advises SMC about CWI's research programme in computer science).

Actual management of SMC and CWI is delegated to the General Director, who is supported by a scientific management committee, consisting of the heads of CWI's scientific departments.

A Science Council advises the Board of Trustees on matters of research policy and organization involving both the National Activities in Mathematics and CWI. The Science Council consists of five researchers from universities, including one from CWI. A number of Advisory Committees make recommendations to CWI scientific departments on implementing research plans.

ORGANIZATION



Organizational chart: the Stichting Mathematisch Centrum SMC and its research institute CWI.

RESEARCH HIGHLIGHTS

Mathematical Epidemiology: Modelling the Force of Infection

Research Programme	:	Modelling and Analysis
Researcher	:	O. Diekmann
E-mail	:	odo@cw.nl

Introduction

As the name suggests, it had always been assumed that malaria was due to 'bad air'. That changed towards the end of the 19th century, when Sir Ronald Ross discovered that it was actually an infectious parasitic disease with two hosts, man and the anopheles mosquito. The latter transfers infection when it bites to feed on blood. Thus one can build on a mechanistic basis, viz. a submodel for biting behaviour, when modelling the force of infection (i.e. the probability per unit of time, of a susceptible human or mosquito, becoming infected). In turn, the force of infection is crucial for the understanding and estimation of threshold parameters that determine the success or failure of an eradication campaign.

For many infectious diseases the modelling of the force of infection is still a problematic issue. Recent CWI research has concentrated on two specific cases: Aujeszky's Disease Virus (ADV) among pigs and Phocine Distemper Virus (PDV) among seals.

Some more history

Many aspects of epidemiology have a high appeal for a mathematical approach, and this was recognized early in the history of science. In 1766 D. Bernoulli compared the risk of smallpox vaccination with the risk of actual smallpox. In 1927, following in Ross's footsteps, Kermack and McKendrick derived a threshold population density for the spread of an infectious disease - somewhat on the lines of critical mass for an atom bomb. During the course of World War II, Kendall went on to analyze the travelling epidemic wave in a spatially distributed population,

but with an eye to implications for biochemical warfare, he delayed publication until 1965. May and Anderson drew attention to the role of parasites as regulators of natural populations and highlighted the evolutionary aspects of this relationship.

Epidemiology was a key topic at CWI during the early days (late sixties and early seventies) of the 'Werkgroep Biomathematica', with Lauwerier providing much of the impetus. Results on the speed of spatial propagation were made operational at the Institute of Theoretical Biology of the University of Leiden and then applied successfully at the Agricultural University in Wageningen to investigate the spread of fungus diseases in various crops.

For a while the attention for epidemiology in the mathematical community waned, both at CWI and elsewhere. However, the resurgence of malaria and the rise of HIV made it clear that infectious diseases still constitute a major public health problem. In 1989 CWI organised a colloquium on 'The spread of infectious diseases in structured populations' and in the period that followed most of the research concerned the definition and computation of the threshold condition in situations where individuals differ widely in characteristics relevant for disease transmission (such as sexual behaviour in the case of HIV). The colloquium also initiated contacts with the Central Veterinary Institute (now ID-DLO) in Lelystad and the National Institute of Public Health and Environmental Protection (RIVM) in Bilthoven.

The problem

Many viruses are transported from the mucus of




one host to that of another by aerosoles when hosts 'meet'. As a consequence one needs to model the contact process first, and then superimpose the transmission of the virus. The situation resembles that of chemical reactions, where molecules have to come close enough before they can react. Inspired by that similarity, standard deterministic epidemic models are in terms of densities and model the contact process by the law of mass action (which says, in this case, that the per individual per unit of time number of contacts with other individuals is proportional to the population density).

Meanwhile, in real life one often has to work with numbers, rather than densities. The sizes of a pig farm and a colony of seals are usually expressed as the number of individuals in them. If one population is twice as large as another, what difference does that make for the spread of the virus? And if PDV hits the seal colony, does the per individual per unit of time number of contacts with other individuals change? What influence does the subdivision of a farm in compartments (stalls), with reduced contacts between individuals in different compartments, have on the spread of the virus? Such questions require us to scrutinise the law of mass action and to indulge in the assumptions underlying a mathematical description of the contact process.

The results

In fact, it is not at all difficult to formulate a model in which density is a given constant, while population size is variable (over populations and, possibly, in the course of time). And this is what



photos of seals sunbathing on sand banks, as well as information about pig farm organisation, suggest as the real situation in these two cases.

This model was analysed, the threshold parameter determined and an equation for the final size of an epidemic derived.

The first test on data from a classic 1936 experiment by Greenwood with mice suffering from *Pasteurella muris* living in a network of cages which was enlarged when population size increased, was inconclusive: the model with the per capita number of contacts per unit of time independent of population size, and the model for which this number was proportional to population size, could both be made to fit the data with roughly the same accuracy. Incidentally, this work was performed at the Isaac Newton Institute in Cambridge as part of a special programme on mathematical epidemiology.

This finding prompted ID-DLO to perform experiments with ADV in two groups of pigs, a large group in a large stall and a small group in a correspondingly smaller stall. Fortunately the outcome here was very clear: the hypothesis of proportionality of contact rate and population size had to be rejected.

Another convincing argument was found in the data about the spread of PDV during the 1988 epidemic in the coastal waters of Northern-Europe: the final size appeared to be independent of colony size. This observation had puzzled researchers applying the 'standard' model in which contact rate is proportional to population size.

Mathematical models developed at CWI help to understand how farm or colony size affects the severity of an outbreak of a virus disease among pigs or seals. Photo's courtesy ID-DLO Lelystad (left) and IBN-DLO Texel (right).

Thus a combination of modelling considerations, mathematical analysis, experiments and observations helped to disentangle some aspects of the complicated relation between mechanisms at the individual level and phenomena at the population level.

Conclusion

When trying to bridge the gap between general

abstract mathematical theory and specific concrete real life situations or simulation models, one frequently encounters modelling problems which previously escaped notice. And analysis of these problems usually enriches general theory, while improving the simulation model. So the tension between abstract and concrete, though undeniably real, is not at all a negative force.

Robust Control

Research Programme	: System and Control Theory
Researchers	: J.M. Schumacher, J. de Does, M.K.K. Cevik
E-mail	: Hans.Schumacher@cwi.nl

Introduction

No model is perfect. A truism perhaps, but one that needs to be taken seriously into account by engineers and applied mathematicians alike. In particular this holds for the control of dynamic processes ('plants' in engineering jargon), in which models are used as the basis for the development of feedback rules aimed at achieving high performance standards. In standard control design methods one usually employs linearized, low-order models. A controller obtained in this way may perform nicely when coupled to the design model in a computer simulation, but what will happen when it is connected to the real plant? The theory of robust control attempts to answer this question. A controller is called robust if it keeps up performance not only for the nominal model, but also for other systems that deviate to a certain degree from that model. Although only partial results are available, these have already allowed engineers to assess model uncertainty more precisely, enabling improvement of the balance of robustness and performance.

Development of the field

Feedback control design on the basis of explicit mathematical models can be traced back to James Clerk Maxwell's paper 'On Governors' presented to the Royal Society in London in 1868. Although Maxwell's fame mainly derives from work in electromagnetic theory, his interests were much wider, notably including the theory of stability, for which he saw an interesting application in the explanation of the behaviour of steam engines controlled by Watt's governor or other devices. Indeed Maxwell's paper led to the formulation of an explicit test for the stability of

linear systems, now known as the Routh-Hurwitz criterion. However, the test was only of limited use to engineers, as the criterion (which calls for the construction of an array of numbers computed from the coefficients of the characteristic polynomial) gives little feel for what happens if there is a small change in the initial data. The American engineer Harry Nyquist came up with a graphical test for stability of feedback systems in the 1930s. Although no explicit uncertainty model was used, the graphical nature of the test made it easy for control designers to get at least a rough idea of the possible effects of unmodelled dynamics, and the Nyquist criterion became a standard tool in the control of single-input-single-output systems.

With the advent of dynamic programming and optimal control theory in the 1950s, followed by the development of the Kalman filter and linear-quadratic control theory in the early sixties, an extensive mathematical theory of control evolved. However, the theme of robustness remained relatively unexplored until the late seventies, when it was discovered that certain tools in the boundary area of complex function theory and linear operator theory could be profitably used to solve certain sensitivity minimization problems for linear systems. To the dismay of bibliographers the term that has become standard for this research area is ' H^∞ optimization', after the function space that plays a central role. This space is one of the family of H^p spaces, which are named after the British mathematician G.H. Hardy (1877-1947).

The new robustness theory is similar in spirit to the Nyquist criterion in that modeling errors are not necessarily viewed as perturbations in parameters. In this way it becomes possible to

Robustness, conservatism, complexity

Robust control is needed in several areas alongside the modelling of uncertainties. When a controller is applied in mass produced items like a beam tracker in compact disk players, one has to take into account the natural variability in individual products. One may try to squeeze out optimal performance from a nominal model only to find out that slight variations will cause severe deteriorations. An alternative strategy is to 'play it safe' and be content with modest performance that is relatively insensitive. If this is done on the basis of a rough assessment of possible variations, performance may be compromised more than is strictly necessary; in such a case the controller design is said to be 'conservative'. One can then attempt to model the uncertainty more precisely and optimize performance with respect to a refined uncertainty specification. However, this may lead to highly complex mathematical problems; hence the challenge of robustness theory is to find good compromises between tractability and accuracy in dealing with variability.

compare models with different numbers of state variables, which is important in engineering applications since models never include all dynamical effects in a plant. The comparison of models is based on a distance measure. The distance measure that has become most popular is the gap metric, which essentially looks at the largest angle between the linear spaces representing the input-output relations determined by different models. This metric has been known in mathematics since the late 1940s, when it was developed, notably in the Soviet Union, as a tool in operator theory. It turned out that the problem of optimizing the robustness of stability of the controlled system in terms of this metric allows an attractive solution which gives rise to a reliable design method ('loop shaping', developed mainly at the University of Cambridge in the group led by Keith Glover).

Alternative recent approaches to robustness include parametric robustness analysis and the so-called μ -analysis. In the parametric methods the set of all considered perturbations is described by a finite number of parameters. This


approach received a strong impetus from a remarkable result by the St. Petersburg mathematician V.L. Kharitonov, who showed that the stability of so-called interval polynomials (polynomials of which the coefficients are only known to lie in certain intervals) can be checked simply by looking at a finite number of extreme cases. However, like many other parametric approaches, the usefulness of this result is highly dependent on specific properties of particular applications. In μ -analysis one tries to obtain very precise uncertainty models, which in principle is important to avoid conservatism (see Box) in controller design. The approach tends to lead to difficult computational problems.

Measures of closeness

Enabling a mathematically precise definition of robustness depends on a precise specification of a 'small deviation'. In parametric robustness analysis, this problem is easily solved because one can use the differences between parameters as a measure of deviation - although even then issues of scaling must be carefully considered. However, the situation is entirely different if one wants a nonparametric theory of robustness, and reasonable measures of deviation can be proposed that differ from each other even more fundamentally than merely in terms of scale. Metrics that differ in this qualitative sense are said to generate different topologies. So there are really two issues: the choice of a topology, which is a qualitative decision, and the choice of a metric, which is a quantitative decision concerned with scaling issues.

Although it is possible to give some general guidelines, the selection of a suitable metric will always be partly application-dependent. However, the choice of a topology is a fundamental decision. One of the important achievements of post-1980 systems theory is the development of a topology for linear systems, which has become

Robust control forms an essential part of the laser beam tracker in compact disk players. Photo courtesy Philips Research Laboratories Eindhoven.



generally accepted as a basis for robustness theory. It is known as the gap topology. After initial formulation in 1980 by Zames and El-Sakary, the significance of the gap topology became much clearer when it was proved that it is actually the same as the graph topology defined by Vidyasagar in 1981. In due course several other characterizations of the gap topology were found which together helped to establish its present status. Efficient methods were also found to compute corresponding metrics.

CWI played a role at several stages in this development. The fact that gap topology and graph topology are the same was proved by Siquan Zhu. A PhD student at Eindhoven University of Technology, he learned about the graph topology at CWI's System and Control Theory seminar, where in the 1986 spring semester Vidyasagar's recently published book was studied. Later on, characterizations of the gap topology became an object of study for CWI staff member Hans Schumacher and his PhD student, Jesse de Does. They pursued a point of view in which the gap topology is looked at as a topology of uniform convergence for a certain function of a complex variable that can be associated to a linear system. This point of view is related to the classical Nyquist criterion, but it adds the explicit measure of deviation and is not limited to systems with just one input and one output. De Does was also able to obtain a new method of computing the standard gap metric which in many cases improves on the computational methods that were known before.

Following a step

As soon as one specifies a measure of deviation, the question arises how to optimize with respect to this measure. For instance, one seeks to design a controller that stabilizes the largest possible 'ball' of plants around a given nominal model. This problem may be posed for any choice of a deviation measure, but in general it is not so easily solved. Hence, it came as a surprise when

in 1988 researchers at Cambridge University showed that a particular metric for the gap topology allows an elegant solution to the problem of optimal robustness of stability; previously the problem had been believed to be of a more general H^∞ type which can only be solved by iterative procedures.

Robustness of stability is a good thing to have, but usually controllers are built with other additional purposes in mind. In particular controllers are designed to achieve 'regulation properties' - e.g. the ability to follow a step change in a reference input. So one may pose the problem of designing a controller that meets regulation requirements and is in addition optimally robust with respect to closed-loop stability. This problem was addressed at CWI by Hans Schumacher in collaboration with Klmiz Cevik of the Department of Electrical Engineering of Istanbul Technical University, who visited CWI for a nine months period in 1992/93. The problem can be viewed as an optimization problem with added side constraints. Adding extra conditions is often quite fatal to elegant solution methods, but it turned out that the regulation constraints can be dealt with in a nice way. The practicality of the design philosophy based on optimizing robustness of stability is thus enhanced considerably.

References

1. J.M. SCHUMACHER (1992). A pointwise criterion for controller robustness, *Systems & Control Letters*.
2. J. DE DOES, J.M. SCHUMACHER (1994). Interpretations of the gap topology: a survey, *Kybernetika*.
3. J. DE DOES, J.M. SCHUMACHER (1994). Continuity of singular perturbations in the graph topology, *Linear Algebra and Its Applications*.
4. J. DE DOES (1994). *The gap topology for linear systems, a geometric approach*, Ph.D. thesis defended at Tilburg University.

Parameter Estimation in Dynamical Systems

Research programme	: Boundary Value Problems, Multigrid and Defect Correction
Researchers	: P.W. Hemker, W.J.H. Stortelder
E-mail	: pieth@cw.nl

Introduction

Demand for mathematical support has burgeoned in many areas of experimental science, including chemistry, biology and electronic engineering. The shared factor is that these use mathematical models to describe and predict process behaviour and experiments are conducted to validate the models. Interest here has been boosted by the urgent need for energy-efficient chemical plants with minimal harmful by-products.

The models used describe time-dependent processes with numerous state variables and with many interactions and feed back loops between such variables. These models can be formulated mathematically by a set of differential algebraic equations (DAEs). It often happens that these model equations contain a number of unknown parameters, e.g. reaction constants (birth rates in population dynamics, etc.) which must be known for a complete process description. The project seeks to determine these unknown parameters on the basis of experimental data. This is accomplished by finding an optimal fit between the measurements from practice, or the laboratory, and the theoretical results obtained by solving the set of DAEs.

In simulation the model equations are assumed to be entirely known and the future behaviour of the process can be calculated. In our case we do have some additional information, i.e. the measurements, but the model contains some unknown quantities. The information from the measurements is used to retrieve these unknown quantities (the parameters). It is said that parameter estimation is the inverse problem of simulation.

Development in the field

The earliest work in the field of the parameter estimation was done by experimental scientists. There are numerous examples of fitting a curve to a set of experimental data in order to estimate unknown physical constants (a reaction rate, the gravity constant, the capacity of a condenser, etc.). As the equations are simple and few in number, in many cases these calculations can simply be done by hand or by simple recipes from statistics.

With the wider availability of computers and the higher levels achieved in numerical analysis, the early sixties marked the onset of a more thorough approach to models described by non-linear differential equations. Sophisticated solvers and good hardware are essential where model equations get more difficult by containing feed-back loops, being highly non-linear and stiff, and indeed, because of the burgeoning number of equations, measurements and unknown parameters. More recent research has resulted in some packages for parameter estimation problems in ordinary differential equations; one is already on the market.

Research at CWI

Researchers at CWI - or Mathematical Centre as it was then - were already looking at parameter estimation back in the early seventies. Alongside theoretical work, they had implemented a working code [1] in ALGOL 60, which met the standards of the day, and in 1974 this was included in the software library NUMAL [2]. However, it was 1991 before active research resumed. The impetus came from AKZO's research unit which

BOX 1

This example treats a simple process from population dynamics describing the population density in a two-species population with a predator-prey relation ('foxes' and 'rabbits'). The model is known as the Lotka-Volterra equations. We consider an area where the fox and rabbit population densities are the state variables of interest. The development of the densities depends on the initial state as well as on the birth-rate of the rabbits and the appetite and death-rate of the foxes. We denote the state vector of densities (rabbits and foxes) by (y_1, y_2) . The parameters for the birth-rate of the rabbits and the appetite and death-rate of the foxes are denoted by p_1, p_2 and p_3 respectively. The predator-prey model is now described by the following set of differential equations:

$$\frac{dy_1}{dt} = p_1 y_1 - p_2 y_1 y_2, \tag{1}$$

$$\frac{dy_2}{dt} = p_2 y_1 y_2 - p_3 y_2. \tag{2}$$

The initial condition of this problem is given by, e.g.:

$$\begin{pmatrix} y_1(0) \\ y_2(0) \end{pmatrix} = \begin{pmatrix} 1 \\ 0.3 \end{pmatrix}.$$

Each measurement is represented by the triple:

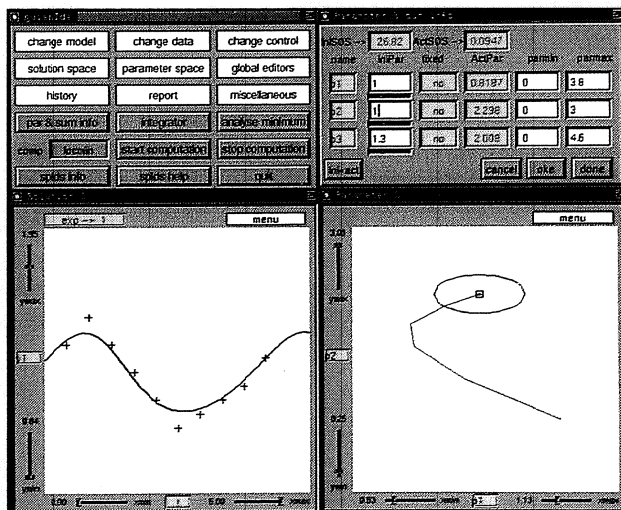
$$\hat{y}_i, c_i, t_i, \quad i \in \{1, \dots, N\}$$

with \hat{y}_i the measured value; c_i indicates the component in the vector of state variables (whether it concerns foxes or rabbits); t_i is the time of the measurement, and N the number of measurements.

Now we want to estimate the parameters p_1, p_2 and p_3 in such a way that the calculated values fit the measured values. In case of a least squares criterion we want to minimise

$$S(p) = \sum_{i=1}^N \left(y_{c_i}(t_i; p) - \hat{y}_i \right)^2, \tag{3}$$

by varying the three parameters, p_i . The minimisation is done by an iterative procedure requiring an initial estimate for the parameters, *IniPar* (see Figure below). The corresponding sum of squared discrepancies is denoted by *IniSoS*. The improved parameter values plus the corresponding squared sum after minimisation is given by *ActPar* and *ActSoS* respectively. The solution of the model equations with the optimal parameters is drawn in the *solution window*, the measurements are marked by: '+'. The successively improved parameter values are drawn in the *parameter window*. The ellipses in these windows show the confidence regions for the parameters.



sought insights into chemical engineering models with unknown parameters by the use of numerical techniques. Such insights can lead to answers on fundamental questions concerning consistency and uniqueness of the model, accuracy of the estimated parameters and strategy for possible additional experiments.

This collaboration with AKZO resulted in a rewritten and updated computer program, based on the techniques developed for the old ALGOL 60 code. The new code, when running on a state-of-the-art workstation, was 100 times faster than the old one. Alongside faster computer technology the use of the Maple computer algebra package contributed to this speed up factor.

Most recent research

Following successful completion of the AKZO project, research was continued with funding from the Dutch Technology Foundation (STW). This had the immediate impact of broadening scope and boosting links with several sectors of industry and research laboratories.

CWI's work now focuses on the design of an interactive tool which can be used by experimenters in model simulation and parameter estimation problems. This program package uses a graphical user interface (GUI) enabling the user to steer interactively through the wide range of possible calculations. The results of the calcula-


tions are shown by direct visualisation. See Box 1 for a simple, brief example of the mathematical formulation and the resulting response from the GUI.

The mathematics underlying the parameter estimation problem is based on results from many areas. The numerical solution of DAEs is a well-known subject in numerical analysis. Alongside the DAEs describing the model, we numerically solve the corresponding variational equations. These equations are derived automatically using computer algebra software by differentiation of the DAEs with respect to the unknown parameters. The DAEs and the variational equations are integrated simultaneously (in contrast to the frequently used finite differences method). The result of the variational equations is indispensable for the calculation of the gradient, which -in turn- is useful for the efficient minimisation of a fitness criterion. We had previously used a (weighted) least squares criterion combined with a local search method. Box 2 gives an impression of the landscape describing the function to be minimised over the parameter space. Research focuses on the use of more sophisticated fitness criteria and global minimisation techniques. Obviously, non-linear constraints for the unknown parameters can also be added.

The last main field of mathematics associated with parameter estimation concerns statistics.

Box 2

In order to give an indication of the irregularity of the object function ($S(p)$ from equation (3), Box 1), we visualised a two-dimensional manifold containing all parameter values leading to the same value of the object function (an iso-surface). The plots below immediately show that, despite the equation's simplicity, the object function is far from pleasant and its minimisation is not straightforward.



Having to deal with errors in the measurements is an inevitable consequence of working with experimental data. This leads to a probability region for the final estimate of the unknown parameters. Standard statistics may lead to poor results, because of the highly non-linear character of the problem. Additional problems arise if errors in the measurements are mutually dependent, or if parameters in the model cannot be estimated separately, but only in certain combinations (see Box 2).

Merging the solution techniques for the above-mentioned problems in an interactive computer-code leads to a helpful tool for experimenters, enabling insights which are hard or impossible to obtain via laboratory experiments. The solution of a parameter estimation problem is certainly not only a matter of computer science. Validation of mathematical models for dynamical systems (occurring in biochemistry, population dynamics, etc.) on the basis of experimental data, still has many practical and mathematical challenges as well, in particular if it is not clear that

the experimental data are sufficient to determine the unknown factors. This often appears to be the case indeed. Then it is important to decide what information *can* be derived and/or what additional experiments would be necessary to determine the undetermined factors.

References

1. P.W. HEMKER (1972). Numerical methods for differential equations in system simulation and in parameter estimation. H.C. HEMKER AND B. HESS (eds.). *Analysis and Simulation of biochemical systems*, Amsterdam, The Netherlands, North-Holland, 59–80.
2. P.W. HEMKER (1981). *NUMAL, Numerical Procedures in ALGOL 60*, volume 47.1–47.7, of *Mathematical Centre Syllabus*. Mathematical Centre, Amsterdam, 7 Vols.: 1. Elementary procedures, 2. Algebraic evaluations, 3. Linear algebra, 4. Analytic evaluations, 5. Analytic problems, 6. Special functions, 7. Interpolation and Approximation.

Hybrid Systems

Research Programme	: Concurrency and Real-time Systems
Researcher	: F.W. Vaandrager
E-mail	: fritsv@cw.nl

Introduction

Hybrid systems are mathematical models of applications consisting of a non-trivial mixture of discrete and continuous components, such as a digital controller controlling a continuous environment. Due to the rapid development of processor and circuit technology, we see more and more devices, ranging from aircraft and nuclear plants to consumer electronics, in which computers interact with the physical world. Given that many of these applications are safety critical, the specification, design and verification of hybrid systems has recently become an active area of research.

In this contribution, we will briefly sketch the development of the field, which is currently spread over distinct disciplines, most notably Computer Science and Control Theory. We will also discuss recent work on hybrid systems in CWI's Concurrency and Real-time Systems group, and in particular the correctness proof of a Philips audio control protocol.

Development of the Field

The first digital computers like the MARK-1 and ENIAC hardly interacted with their environments, behaving almost like autistic children, totally wrapped up in their own 'thoughts' and not reacting to stimuli from the environment. After initialization these machines computed continuously until termination of the program or a heavy-handed interrupt by the operator. However, designers soon became aware of massive potential in the control context for marrying these machines' digital computing powers with an ability to interact with the physical environment. The design of the next generation of computers


was adapted to enable this, and as early as 1953 MIT's WHIRLWIND computer was used for real-time process control in an aircraft tracking system. Since then, the use of computer technology in technical applications has grown exponentially until the present day, and examples of applications are now omnipresent in areas such as telecommunication, transportation, consumer electronics, process control, and medical systems.

Characteristics of these applications are that (1) the digital computer systems are *reactive*: they accept stimuli from the outside world at any time and react to those stimuli; (2) in many cases it is vital that reaction occurs in *real-time*, i.e., within a fixed amount of time; and (3) in order to reason about the correctness of the computer systems, one has to take into account the behavior of their *continuous environment*. For example, consider a computer system controlling a continuous plant. The control program regularly reads the plant's sensor data, like temperature and pressure. Based on these data, the computer will compute a new control law, and impose it on the plant by turning on a heating system, switching off a pump, etc. When a dangerous situation arises, for instance the pressure in a tank gets too high, the computer must quickly initiate appropriate action, like opening a valve, to avoid a catastrophe. For a formal proof of correctness properties like 'No explosion will ever occur', one needs knowledge on chemical reactions which take place, the differential equations determining the evolution of the pressure in the tank, etc.

Mathematical models of the type of phenomena described above, i.e., reactive, real-time applications comprising a non-trivial mixture of

discrete and continuous components, will be referred to as *hybrid systems*. The burgeoning information technology infrastructure and its increasingly powerful applications have caused a widespread use and rapid proliferation of hybrid applications in everyday life, and an increasing dependency of their correct operation. Clearly, design faults and bugs can have disastrous consequences for safety critical systems such as nuclear plants control software and digital flight control systems. The same applies to non-critical systems like consumer electronics where errors in software or digital hardware can have a major financial impact.

Although practical construction of hybrid systems started back in the early fifties, their mathematical study has only recently secured some of the attention it merits. Up to the late seventies, the main research thrust on program verification was directed at analysis of programs for 'autistic' batch computers. The eighties have witnessed a revolution in the formal methodology for the specification, verification, and development of reactive programs and - in general - reactive systems. Once a good understanding of the reactive system concept had been achieved, it was far easier to involve quantitative, real-time aspects. Although the development of reliable real-time systems will remain an important research topic for many years to come, several important theoretical results concerning real-time systems had been obtained by 1991. Around that time, the first two important workshops on hybrid systems organized in Ithaca, NY, USA and in Lyngby, Denmark, triggered an explosion of research in this new area. The history of the field from the perspective of Theoretical Computer Science is summarized as follows:



Timed and Hybrid Automata

The analysis and design of hybrid systems require a synthesis of ideas, concepts, mathematical theories and tools that are currently spread over distinct disciplines. The key role played by Control Theory alongside Computer Science, has had a significant impact at CWI. For example, hybrid systems have attracted interest and spurred synergistic benefits at and between the Concurrency and Real-time Systems group (computer science), and the System & Control Theory group (mathematics). Reference 1 provides an introduction to the area from the perspective of CWI's System & Control Theory group.

The notion of a *transition system (TS)* plays a central role in the theory of reactive systems. A TS consists of a set of *states*, a subset of *initial states*, a set of (*discrete*) *actions*, and a set of (*discrete*) *transitions*, which are triples

$$s \xrightarrow{a} s'$$

specifying that from state *s* the system can evolve to state *s'* by the instantaneous occurrence of the action *a*. A run of a TS starts in an initial state. The system jumps from state to state via instantaneous transitions, and in between these transitions, it can remain in any state for an arbitrary period.

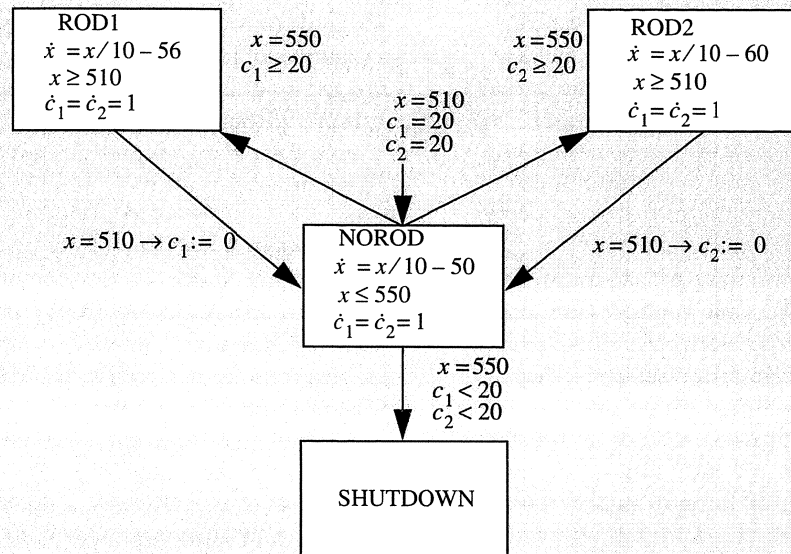
At the lowest level, timed and hybrid systems can be described by TSs with as an additional component a collection of *time transitions*, which are triples

$$s \xrightarrow{d} s'$$

specifying that from state *s* the system can evolve in a positive, real-valued amount of time *d* to state *s'*. In the model of *timed transition systems (TTS)* proposed by N.A. Lynch and F.W.

Vaandrager, two axioms are imposed on time transitions. The first axiom states that if there are time transitions $s \xrightarrow{d} s'$ and $s' \xrightarrow{d'} s''$, there exists a time transition $s \xrightarrow{d+d'} s''$. Stating the second axiom is rather more complex; it postulates that for each time transition $s \xrightarrow{d} s'$ there exists a *trajectory*, a function *w* that specifies an intermediate state for each intermediate point in time, such that $w(0) = s$, $w(d) = s'$ and for all $t, t' \in [0, d]$ with $t < t'$,

$$w(t) \xrightarrow{t'-t} w(t') .$$



A hybrid automaton for a simple temperature control system. The system controls the coolant temperature in a reactor tank by moving two independent control rods. The temperature of the coolant is represented by the variable x . Initially x is 510 degrees, both rods are outside the reactor core and the system is in state NOROD. In this state the temperature rises according to the differential equation $\dot{x} = x/10-50$. If the temperature reaches 550 degrees, three things can happen: either rod 1 is put into the reactor core and the automaton moves to state ROD1, or rod 2 is put into the reactor core and the automaton moves to state ROD2, or a complete shutdown occurs and the automaton moves to state SHUTDOWN. Mechanical factors make that a rod can only be placed in the core if it has not been there for at least 20 seconds. Shutdown will only occur if no rod is available. In the automaton, variables c_1 and c_2 are used to measure the times elapsed since the last use of rod 1 and rod 2, respectively. The initial values of both c_1 and c_2 are set to 20 seconds, so that initially both rods are available. Since they represent perfect logical clocks, the first derivatives of c_1 and c_2 with respect to time are always equal to 1. Control rod 1 refrigerates the coolant according to the differential equation $\dot{x} = x/10-56$; control rod 2 has a stronger effect and refrigerates the coolant according to the differential equation $\dot{x} = x/10-60$. If the temperature has decreased to 510 degrees, the system moves back from state ROD1 or ROD2 to state NOROD, and variable c_1 or c_2 , respectively, is reset to 0. The correctness property for this system is not difficult to prove and says that the SHUTDOWN state cannot be reached. (This example is from Leveson et al.)



Core of High Flux Reactor in Petten.

Thus a trajectory describes *how* the system evolves from s to s' . A run of a TTS consists of a sequence of two-phase steps. The first phase of a step corresponds to a continuous state transformation described by a trajectory. In the second phase the state is submitted to a discrete change taking zero time, described by a discrete transition.

We can add more structure to timed transition systems by defining states to be pairs (\vec{x}, \vec{y}) of a vector \vec{x} of *discrete variables* and a vector \vec{y} of *continuous variables*. Both the discrete and continuous variables can be changed in a discrete transition. However, only the continuous variables may change in a time transition. In the Timed Automata of Alur and Dill, which is a popular model of real-time systems, all time transitions are of the form

$$(\vec{x}, \vec{y}) \xrightarrow{d} (\vec{x}, \vec{y} + d)$$

where $\vec{y} + d$ is the vector obtained by adding d to each of the variables in vector \vec{y} . In this model, the continuous variables behave as perfect logical *clocks*, whose values increase apace with time. In the timed transition systems associated with the more general Hybrid Automata models of Nicolin, Sifakis and Yovine, and Maler, Manna and Pnueli, the way in which the continuous variables change can be specified via differential equations. On the opposite page an example is presented of a


Hybrid Automaton modelling the temperature control system of a nuclear reactor.

Work at CWI — The Audio Control Protocol

Central to the approach of the Concurrency and Real-time Systems group at CWI is the I/O automata model of N.A. Lynch and M. Tuttle. N.A. Lynch and F.W. Vaandrager jointly extended this model for reactive systems to the setting of real-time and hybrid systems. In particular this involved the development of simulation techniques, which are used to prove that one hybrid system is an implementation of another hybrid system, or an implementation of a reactive or real-time system. Jointly with I. Polak of Amsterdam Univer-



Figure 1:
Philips 900 audio system.
Photo courtesy
n.v. Philips Industrial Activities
Leuven.



The Philips protocol uses the well-known Manchester encoding of bit strings. In this encoding, the time axis is divided into *bit slots* of equal length, one bit slot for each bit in the string. In the block-shaped signal that is sent over the bus, a bit '0' is represented by a downgoing edge in the middle of a bit slot, and a bit '1' by an upgoing edge in the middle of a bit slot. If the same bit is sent twice in a row then an additional edge is required, which is placed exactly in between the corresponding two bit slots. The diagram displays the Manchester encoding of the string '1001100'.

sity, D.J.B. Bosscher and F.W. Vaandrager developed a language for the specification of so-called *linear hybrid systems*, a class of extended I/O automata that is closely related to the hybrid automata of Nicollin et al. The theoretical work on hybrid systems has been applied in a project for Philips; this is briefly discussed below.

Fully-fledged computer networks are standard features in today's consumer electronics, like the Philips 900 audio system shown in Figure 1. These networks enable the different devices to talk to each other, and to offer a series of new, useful services to the consumer. For example, a consumer can wake up the whole system at the touch of a single button: there is no need to switch on the tuner first, followed by the CD player, then the amplifier, etc. Instead the system will perform this task by broadcasting a 'wake up' message over the network. The main technical difficulty in building the network for the Philips 900 audio system was that it had to be cheap: consumers are only willing to pay fractionally more for the additional services provided by the network. In fact, the only additional hardware that Philips needs to implement in the network comprises a few transistors, resistors, etc., for the bus interface. The software runs on microprocessors that have to be present anyway. Because the clocks of these microprocessors drift, and because the programs dealing with the network sometimes have to wait for other programs that run on the same microprocessors, the network protocol has to deal with a significant uncertainty in the timing of events. In fact, Philips allows for a tolerance of 1/20 on all the timing.

Despite this very large timing uncertainty D.J.B. Bosscher, I. Polak and F.W. Vaandrager have proved the correctness of part of the Easy-link real-time protocol used by Philips to achieve

reliable communication between the devices. The protocol is modeled as a linear hybrid system, with continuous variables to represent the drifting logical clocks of the sender and receiver in the protocol. In formal terms, the drifting is expressed by the requirement that the first derivative of the clock variables is in the interval $[1 - T, 1 + T]$, where T is the tolerance on the timing. Correctness of the protocol has been proved if the tolerance is less than 1/17. This value is larger than the tolerance of 1/20 allowed by Philips. A counter-example shows that the protocol fails for tolerances greater or equal to 1/17.

Mechanical support is vital in managing the complexity of real world applications. With this in mind, much of the research effort on hybrid systems currently focuses on development of mechanical tools to support specification and verification. At CWI, W.O.D. Griffioen has succeeded in mechanically checking the complete verification of the audio control protocol using the LP general purpose theorem proving tool. An impressive complementary result has recently been obtained by P.-H. Ho and H. Wong-Toi of Cornell University. Based on the CWI/UoA modelling of the Philips protocol, they fully automatically verified an instance using the HYTECH symbolic model checker. They also automatically synthesized the maximum clock drift of 1/17.

Reference

- A. OVERKAMP, J.J. VAN SCHUPPEN (1994). Control of discrete event systems – research at the interface of control theory and computer science. K.R. APT, A. SCHRIJVER, N.M. TEMME (eds.). *From Universal Morphisms to Megabytes – a Baayen Space Odyssey*, CWI, Amsterdam, 453–467.

An Electronic Wallet for Digital Money

Research Programme	: Cryptography
Research	: R. Hirschfeld
E-mail	: ray@cw.nl

Introduction

The last quarter century has witnessed explosive growth in the technology for automated handling of information. This has resulted in many conveniences, but has also introduced new dangers, such as increasing opportunities for unauthorized access to sensitive or personal data, and for tampering with such data. With the advent of electronic commerce has come the need for electronic money, i.e., a digital representation of cash. Electronic money introduces additional associated security problems, such as forgery of money, respending the same money, and invasion of privacy of people's spending habits. The field of cryptography has addressed the problems of authentication and data security in general and of the security of electronic money in particular.

The project CAFE, which is carried out by a consortium of thirteen European institutions and is funded by the European Commission's ESPRIT program, has applied modern cryptographic techniques to produce a highly secure but also open and flexible system for consumer payments using electronic money. As a leader in theoretical research on electronic money and as coordinator of the CAFE project, CWI has played a major role in the development of the protocols used by the system.

Background


Europe leads the world in the introduction of smart cards, wallet-size cards with embedded computer chips. French bankcards and the telephone cards of several countries are chip-based. These early chip cards do not really merit the name 'smart' - they contain memory chips and work in essentially the same way as magnetic

stripe cards except that they are more difficult for a counterfeiter to overwrite. More recent systems incorporate microprocessor chips, so that the card can participate actively in transaction protocols. This allows the implementation of digital signature techniques, upon which electronic money is based.

An application such as a telephone card, in which the card issuer is the same as the service provider (so no clearing is necessary), and all points of payment are online, does not really require a sophisticated microprocessor on the card. But the availability of smarter cards has led to the introduction of prepaid electronic purse systems, in which value is stored on a card and can be used for payment at a variety of shops and other service providers. Because these transactions can be completed offline, they are suitable for small value payments for which cash is traditionally used; credit and debit cards require costly online authorization, which is infeasible for low-value payments. Electronic purse systems are undergoing trials in several European countries;

*The CAFE infrared wallet.
Courtesy CAFE
Project.*





The CAFE card.
Courtesy CAFE
Project.

the Dutch banks plan to introduce one in the Netherlands next year based on a system developed and currently under trial in Belgium.

The security of electronic purse systems is based on digital signatures, a cryptographic method of certifying the origin of a digital message. Messages (which can represent banknotes or card balances) are signed by an algorithm that uses a secret key provided by the signer, and authentication is performed by another algorithm

that also uses a key. In most of the electronic purse systems underway, the signing key and the authenticating key are the same. Because the authenticating key present at the point of sale could also be used to sign messages (i.e., create money), it is protected against discovery and possible misuse by storing it in a tamper-resistant hardware module. This reduces the flexibility of the system; it is difficult to combine multiple issuers of electronic money into the same system, and security modules cannot be given out indiscriminately, but only to service providers who can be trusted not to try to compromise them.

An alternative to symmetric systems that use the same key to create and to verify a signature is public-key digital signatures. In a public-key system, the signing and authenticating keys are different, and no knowledge of the signing key can be obtained from the authenticating key. Such an asymmetric system is ideal for an open and interoperable electronic purse environment, because the signing key need be known only to the issuer, and the authentication key can be made public and need not be protected in any way.

Despite their advantages, public-key signatures are not yet widely used in commercial sys-

Partner	Description	Country
CWI	National research institute for mathematics and computer science	The Netherlands
CardWare	consultant for the financial industry on electronic payment technologies	United Kingdom
IFS	social research institute	Germany
Gemplus	smart card manufacturer	France
DigiCash	software and hardware designer	The Netherlands
Ingenico	point-of-sale terminal manufacturer	France
Siemens	industrial electronics manufacturer	Germany
SEPT	national telecommunications and postal research institute	France
Catholic University of Leuven	university	Belgium
Royal PTT Nederland, PTT Research	national telecommunications company research laboratory	The Netherlands
SINTEF DELAB	university research institute	Norway
Aarhus University, Mathematics Institute	university research institute	Denmark
University of Hildesheim, Informatics Institute	university research institute	Germany

The CAFE Consortium.

One of the fundamental notions of modern cryptography is that of a one-way function - a function that is efficient to compute but impractical to invert. For example, it is easy to multiply large integers, but is thought to be difficult to factor a large integer product into its constituent prime factors. Similarly, modular exponentiation is relatively efficient, but its inverse, the discrete logarithm, remains intractable. Although no proof of the difficulty of either of these number-theoretic problems is known, they have thus far resisted all attempts at efficient solutions. Many cryptographic protocols are based upon assumptions of their intractability.

A related notion is that of a one-way trapdoor function, which is normally difficult to invert, but which becomes easy to invert if some additional information is known (this is the trap door). This additional information can form the secret key (or part of the key) in a cryptosystem. An example is the well-known RSA cryptosystem (named after its inventors: Rivest, Shamir, and Adleman), which is based on the observation that if $n = pq$, the product of two large primes, then computing powers mod n is easy but computing roots mod n seems difficult unless the factors p and q are known, in which case it is easy.

Cryptosystems were originally developed for encryption rather than signing of messages. In general, a public-key cryptosystem consists of a public encryption algorithm E and a secret decryption algorithm D with the property that for a message m , $D(E(m)) = m$. Each person has his own pair of encryption and decryption algorithms. To send a secret message, the sender encrypts it using the recipient's E (which is publicly available), and then only the recipient can decrypt it, using his own secret E .

Digital signatures turn this situation around. The secret algorithm D is used to sign messages, and the public algorithm E is used to authenticate them. The property required is that $E(D(m)) = m$. To sign a message, the signer applies her own secret algorithm D , and then anybody can verify it using her publicly available E . The RSA cryptosystem can be used for both encryption and for digital signatures because it has the property that $D(E(m)) = m = E(D(m))$, since encryption and decryption are the inverse operations of raising to a power and extracting a root.

Commercially, the most widely-used cryptosystem is the Data Encryption Standard (DES). This is a symmetric algorithm and is unsuited for public-key use, but has the advantage that it does not rely on time-consuming modular arithmetic. As hardware support for these operations becomes more widely available and inexpensive, however, this is becoming less of a consideration.

tems. This is because they require more elaborate computation, and until recently performing this computation on a smart card chip has been too slow or too expensive. But with the development of specialized cryptographic smart card chips, which implement complex operations needed for public-key signatures in hardware, public-key electronic purses are now feasible. The CAFE project has developed a public-key system for electronic purses, and, more generally, for electronic wallets, which combine an electronic purse with other applications, such as digital passports, driver's licenses, house keys, etc.

In addition to a smart card, CAFE has developed a hand-held wallet that communicates with payment terminal via infrared (much like a television remote control, except bidirectional) and allows consumers to confirm payments with their own device and to complete the payment without the wallet ever leaving their hands.

CWI's role

CWI is the coordinating partner of the CAFE project, and as such is responsible for the overall management of the project. In addition, CWI

plays a major role in the design of the protocols used in the system, drawing on its many years of active research experience in digital signatures and electronic money.

In collaboration with the other protocol partners, CWI worked on all aspects of the cryptographic protocols developed for CAFE, including not only the fundamental protocols for secure transactions (withdrawal, payment, deposit, etc.), but also currency exchange, tolerance of loss and faults, key management, and other related items.

CWI has also applied its special expertise on privacy protection and user-moderated transactions in the design of the wallet protocols and the provision of untraceability as a system option.

Progress in 1994

CAFE is a three-year project that is roughly divided into a year of design, a year of implementation, and a year of trial. In 1994 the project was in its second year, and so the emphasis was on implementation. In the first half of the year, specifications were finalized and the construction of prototypes of the devices begun. CWI was primarily involved in the specification of the basic pro-

ocols to be implemented in the trial.

In the second half of the year, work began in full on the development of the various devices, and the firmware for the microprocessor chip was completed and fabrication of the chip begun. For CWI and the other protocol partners, emphasis shifted to the design of extended protocols, which would not be implemented in the trial but which form an important part of the system architecture. These include enhanced electronic purse protocols, as well as additional electronic wallet applications, such as electronic credentials and loyalty schemes.

Plans for 1995

In 1995, the CAFE system will undergo a trial on

the premises of the European Commission in Brussels. If successful, this will expand to include other EU institutions in other cities and perhaps the surrounding communities. Although the currency exchange mechanism developed by the project is very general, the trial will focus on the introduction of an electronic ECU. Users will load their cards with any combination of their home currency and ECU, and (in the trial) they will be able to spend their home currency only in their home country, but the electronic ECU at a CAFE terminal in any country.

The results anticipated include an evaluation of the technology, surveys of user opinions, and a cost assessment and business case analysis of the system.

Computational Steering

Research Programme	: Computer Graphics and Visualization
Researchers	: R. van Liere, J.J. van Wijk
E-mail	: robertl@cw.nl

Introduction

The standard cycle in simulation is to prepare input, execute a simulation, and visualize the results. Performing these activities simultaneously realises greater insights and higher productivity. This is the underlying idea of *Computational Steering*: researchers change parameters of their simulation on the fly and immediately receive feedback on the effect. However, the development of a dedicated user-interface for a simulation is a time-consuming process, which requires the expertise of specialists on user-interfacing and computer graphics. Hence CWI's Computational Steering project aims to develop an environment in which researchers themselves can develop and use interfaces to their simulations.

Background

Scientific Visualization has been a separate research area since publication of an influential report by the US National Science Foundation in 1987. Many new methods, techniques, and packages have been developed over the past several years; but most of these are confined to post-processing of data-sets. The usual assumption is that all data is generated first, after which the researcher iterates through the remaining steps of the visualization pipeline (selection, filtering, mapping, and rendering) to achieve insight in the generated data. Hence, there is only limited interaction with the simulation.

Tracking is the first step to increasing interaction with the simulation. After each time-step of the simulation the resulting data for that time-step is sent into the visualization pipeline and can be inspected. If the researcher considers the results

invalid, the simulation can be stopped at an early stage, and restarted with a different set of input parameters. The next step, *Computational Steering*, goes a lot further, and can be considered as the ultimate goal of interactive computing. Computational steering enables the researcher to change parameters of the simulation while the simulation is progressing.

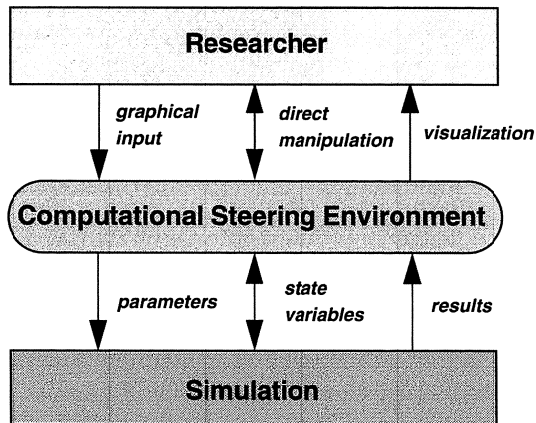
As an example, R.E. Marshall of the Ohio Supercomputer Center has applied computational steering to the study of a 3D turbulence model of Lake Erie. His conclusions were that: 'Interaction with the computational model and the resulting graphics display is fundamental in scientific visualization. Steering enhances productivity by greatly reducing the time between changes to model parameters and the viewing of the results.'

Requirements

Although computational steering is an attractive concept, implementation is cumbersome. A researcher will need to ask a specialist in user-interfaces and visualization to develop a suitable tool; and in the weeks or months this takes, the chances are high that the researcher's interests have shifted. Moreover, the use of computational steering will introduce new research questions, which prompt modifications of the tool. Hence, researcher and specialist will need to collaborate for an extended period.

The Computational Steering project started up in 1992 as a joint project between CWI and the Netherlands Energy Research Foundation (ECN). The aim is to develop insights, methods, techniques, and tools enabling researchers to apply computational steering. The current focus is on development of a Computational Steering Envi-

ronment (CSE) that encourages exploratory investigation by the researcher of his simulation.



What are the requirements for such an environment? The first set of requirements follows directly from the definition of computational steering. Researchers must be enabled to change parameters and to visualize results. The simulation must be enabled to read these parameters and to write the results. Certain variables fall both in the input and output categories. For instance, in a time-dependent simulation the state variables are calculated via integration, but occasionally the user might want to change the current value (position, velocity, etc.). This implies that the CSE must support direct manipulation: the user can interact with graphics objects that are also updated by the simulation.


Second, the environment should enable researchers to use computational steering without help from visualization experts. In the process of gaining insight via computational steering, a researcher typically wants to look at and to control other, possibly new, variables, and to visualize them in various ways. This implies that it should be easy to define and refine an interface, as well as to connect a simulation with the environment.

Finally, the environment must support multiple processes running simultaneously. This enables the integration of existing special purpose packages for input and visualization, and makes it possible to run complex simulations consisting of several separate processes.

Architecture

Given the set of requirements, how can we realize a solution? We have chosen for an architecture

concentrated around data. The main process in the CSE is the Data Manager, which manages a database of variables. Attributes (type, dimensions, etc.) and its current value are stored for each variable. Other processes (called satellites) can connect to and communicate with the Data Manager. Satellites can create, read and write variables. Furthermore, satellites can subscribe to events, such as notification of mutations of a particular variable. The Data Manager takes care of event notification. Thus, the Data Manager enables satellites to use the same data and to communicate with other satellites.



A simple library of subroutines is available for researchers, and they can connect their simulation with the Data Manager by declaring the relevant variables. Next, a single call in the inner loop of the simulation suffices to exchange data between the simulation and the Data Manager.

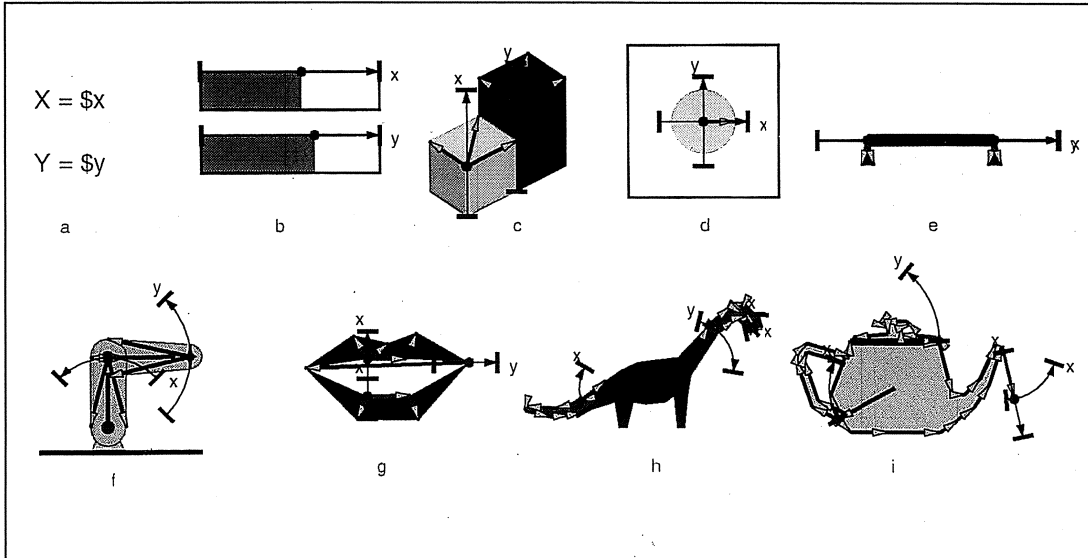
We have developed a variety of general purpose satellites for standard tasks. Data can be logged, sliced, transformed, and calculated. However, the most important satellite is a general purpose satellite for input and visualization of data.

Parametrized Graphics Objects


Many ingenious, but unrelated solutions are available for the input of data and the visualization of data. We have chosen a different solution: look for the greatest common divisor of these aspects, and provide an homogeneous solution. The greatest common divisor of user input widgets and visualization tools simply happens to be graphics. Buttons, sliders, graphs, histograms all boil down to collections of graphics objects. Therefore, we use Parametrized Graphics Objects (PGOs) as the

BOX



The researcher is free to design representations conveying the semantics of the underlying variables: Nine different ways to visualize two scalar variables x and y are shown, both in edit-mode and in run-mode. The variables can be presented via standard UI-widgets (a and b), or in a business graphics style (c). The parameters can also refer to a position (d), a range (e), or have a mechanical (f) or sensual (g) interpretation. Typical computer graphics interpretations are given in (h) and (i).



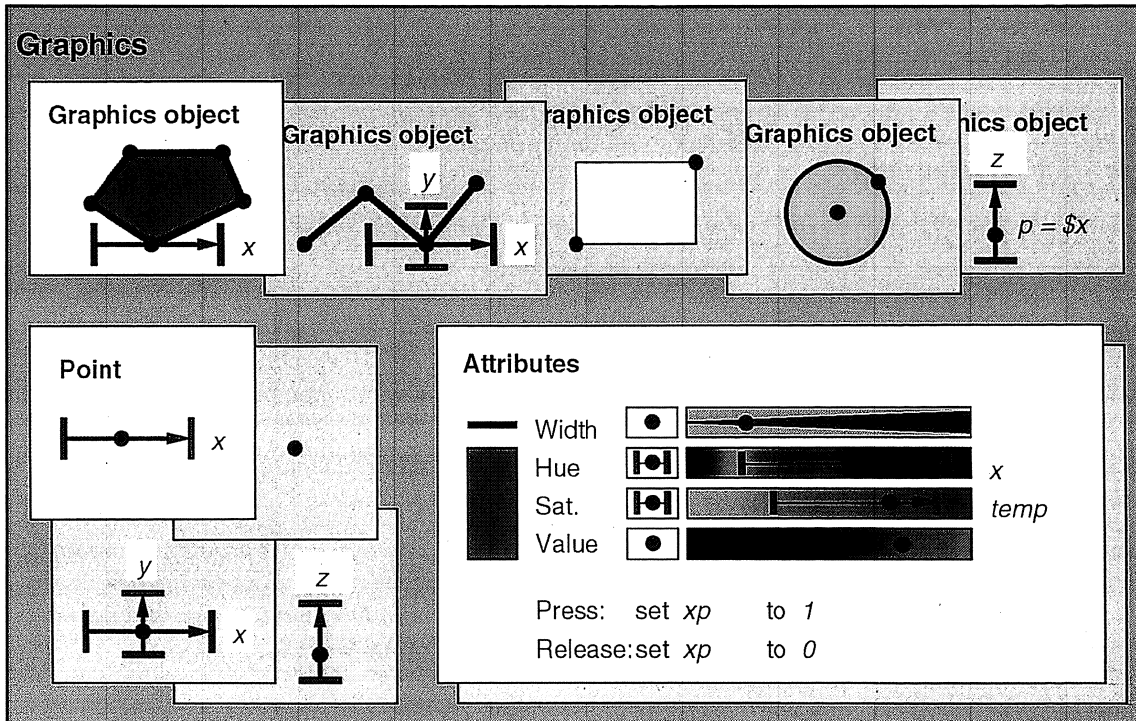
edit-mode



run-mode



Ceramic Foam Burner: computationally via CSE (above) and experimentally (below).



*Parametrized
Graphics
Objects.*

main concept for the general purpose interaction satellite, and the graphics editor as a metaphor for the design of the interface.

The graphics editor has two modes: specification and application, or shorter, *edit* and *run*. In edit-mode, the researcher can create and edit graphics objects much like in MacDraw-like drawing editors. The geometry of the objects is defined by points. To each point degrees of freedom can be assigned, which can be parametrized to values of variables in the Data Manager. The colour and linewidth of objects can also be parametrized.

In run-mode, a two-way communication is established between the researcher and the simulation. Data is retrieved from the Data Manager and mapped onto the properties of the graphics objects. Also, the researcher can enter text, drag and pick objects, which is translated into changes of the values of variables in the Data Manager. Hence there is automatic support for the usually time-consuming and error-prone direct manipulation of graphics objects.

Application

Ceramic foam burners are studied intensively at ECN, both experimentally and computationally. The aim is to realise energy savings and to reduce NOx emissions. A one-dimensional mathematical model for ceramic foam burners has been developed that incorporates a one-step chemical reaction, heat transfer between gas and ceramic, radiation from the burner surface to the surroundings, and radiative heat transfer within the ceramic foam. The CSE has been used to steer those simulations. The concentrations of oxygen and gas, and the temperatures of the ceramic and the gas are shown along the burner. The radiation and resulting NOx concentration are also displayed. The user can control the surplus of air in the mixture and the desired power. These parameters control the air and gas flow. In the setting shown here, the burner is in a radiant mode: the mixture burns at the downstream edge of the ceramic foam. As a result, the foam heats up, and radiates its energy to the surroundings. This particular interface was developed in half a day, demonstrating that our approach effectively enables researchers to apply computational steering.

ATM Networks: A New Infrastructure for Research Computing

Research Programme : Multimedia Kernel Systems
Researcher : D.C.A. Bulterman
E-mail : dcab@cwil.nl

Introduction

It has been clear for over a decade that a reliable and (relatively) fast communications network is an essential—perhaps *the* essential—component of a research computing infrastructure. In 1992, CWI received a large grant from the Dutch government to upgrade its network infrastructure from standard 10 mega-bit/second (Mbps) Ethernet to new, multi-hundred Mbps networks. This grant, which provides funds to upgrade facilities ranging from network cables to new workstations, will allow CWI to be among the first European research institutes to migrate nearly its entire research staff to these new technologies when the project completes in early 1996. While new applications areas such as high-performance computing, multimedia and scientific visualization provide a compelling incentive to expand data transfer bandwidth, the heavy use of more


mundane' services such as network file systems, file transfer protocols and access to other shared resources provide an equally important motivation to expand networks.

New Technologies and New Needs

Unlike past user requirements, in which simple connectivity was the central demand, current and future users need networks that support two basic characteristics: *low latency* and *high throughput*. Low latency means that the time between transmission and receipt of information is short—short enough, for example, to support data types that have strict relative arrival time constraints (such as audio and video information). High throughput connections allow large amounts of data to be sent quickly, in a sustained manner. While all computer networks would seem to provide such services, the reality of current-generation technology such as Ethernet is that there are no guarantees for data throughput and latency, meaning that they fail to meet the needs of all but the simplest time-constrained applications.

Several new and emerging technologies exist that can provide the basis of new infrastructures supporting high-throughput, low-latency transfers. Among these technologies are: Fiber-optic Digital Data Interface (FDDI) networks, a mature token-ring 100 Mbps technology; Asynchronous Transfer Mode (ATM) networks, an increasingly available mixed-medium circuit-switched 155 Mbps technology; and 100 Base-T networks—or, more simply, 100 Mbps Ethernet, a new version of Ethernet that promises a relatively inexpensive speed upgrade of existing Ethernets. Of these choices, ATM provides the greatest potential for a

Figure 1: ATM Architecture. A workstation and server exchange information as streams of cells, with each cell holding 48 bytes. The switch internally routes data between ports (f) and (m). Some switches allow broadcasting or multicasting of cells, allowing one sender to address multiple receivers. At present, typical port-to-port data rates are 155 Mbps. Top rates are currently in excess of 600 Mbps.



new central infrastructure: it is not bound to a 100 Mbps maximum (with existing switches already able to support more than 600 Mbps) and it provides its data transfer rates on a *per-user* basis, rather than as a shared resource. The architecture of a simple ATM infrastructure is illustrated in Figure 1.

Initial ATM Implementation Results

Early in 1994, we began our first experiment with ATM networking. As this technology is still under development — especially at the higher, more user-oriented layers — we felt it was important to get practical experience with initial implementations of ATM hardware, even if this meant that the several generations of facilities would need to be used before a final infrastructure would be put into place. In order to evaluate the technology, a small test ATM network was installed connecting 16 file/compute servers and workstations. The ATM infrastructure was provided by Fore Systems, and included an ASX-100 ATM switch and interfaces for Sun and Silicon Graphics workstations and servers.

To the surprise of those conducting the experiments, the installation of the ATM network interfaces and software occurred without major problems. The real surprise came, however, when we conducted a series of file transfer experiments using a mix of machines and ATM interfaces. Results varied between a ten-fold improvement in performance and a nearly thirty-fold *decrease* in transfer performance compared with standard Ethernet.

Table 1 summarizes three sets of tests conducted using the Fore ATM interfaces and switch. (In all cases, Silicon Graphics workstations were

used to minimize operating systems differences.) As can be seen, performance using ‘out of the box’ settings range from dismal to impressive, depending on the architectures involved. In tests of a light-weight version of TCP, a performance improvement can be achieved only after careful system tuning; unfortunately, tuning parameters were not constant across architecture types. In the final test, we modified the IP MTU sizes and measured FTP performance. Here again, different (architecture-dependent) settings were required to obtain satisfactory performance.

It is clear that a combination of workstation processing speed and parameter settings strongly influence the benefits of ATM. In addition, the architecture of the ATM interface at the workstation is critical in off-loading the CPU’s processing of in-coming ATM cells. (In our tests with SGI GIO-100 interfaces, up to 80% of total CPU time was spent servicing cell interrupts—essentially crippling the user’s workstation; use of intelligent interfaces from the same manufacturer provided a substantial improvement.)

The real challenge: applications

Our initial tests provided insights into *wires-and-pliers* level ATM testing. At this level, the placement of an ATM network remains an out-of-the-box experience, especially if all components are acquired from a single supplier. Connecting ATM switches from different suppliers is much less straightforward, and getting user workstations to function correctly using a mix of equipment vendors — a normal occurrence for Ethernet networking — is still an exciting and unpredictable activity with ATM.

Still, all of these technical issues can expected

Table 1: ATM performance tests from an SGI Challenge server to SGI workstations

Test #	Settings	Destination Workstations (all figures in Mbit/sec)		
		SGI R3K Indigo/ GIO-100	SGI R4K Indigo/ GIO-100	SGI Onyx/ VME-200
1: Basic TCP/IP	‘Out of box’	0,3	0,9	66,0
2: tcp IB=input blk size OB=output blk size	OB=IB=8192	0,4	1,0	85,0
	OB=8192, IB=256	0,4	12,0	32,0
	OB=256, IB=8192	10,0	14,0	17,0
	OB=IB=256	8,0	14,0	14,0
3: IP MTU size test (avg FTP speed)	MTU = 9188	—	—	50,0
	MTU = 1064	13,4	21,5	—

to be solved within a generation or two of ATM hardware. The commercial pressures to come to a variety of ATM standards is so great that one can expect quick—but not immediate—harmonization of network protocols. A bigger challenge concerns the effective *use* of the possibilities that multi-megabit networks offer.

If we look backward to early communication strategies, we can see that the availability of a base technology has nearly always driven the applications that follow. In the early days of computer-based telecommunications, developments such as full-screen text editors and transaction processing were made possible by a stable (if slow) underlying architecture.

More recently, the use of workstations coupled with client/server networks led to the development of window managers and point-and-click computing, where information could be accessed transparently from servers that a particular user never knew existed. A glimpse of the future of networked computing is currently available via the World-Wide Web, which provides a multi-media information service across the Internet. Here, users can surf the network, pulling complex data pages over for presentation on their local hosts. (See Figure 2 for an example.) At present, most Web users access the Internet using slow-speed modems or — at best — network connections of a few mega-bits. Imagine the use of network connections that provide each user with over 100 Mbps of data at a


sustained rate: this makes high-quality audio and (compressed) broadcast video from multiple sources available to even relative simple user presentation systems.

Beyond the movement of data, ATM also holds other prospects. One exciting idea for most researchers who are forced to share their offices with one or more high-speed workstations is that ATM networks can restore peace and quiet to the workplace. Since the network connection is already faster than most hardware disk drives and I/O busses, it should be possible to move the disks and even the processors into separate rooms, leaving the user with a monitor, keyboard, mouse and loudspeakers on the desktop.

Future plans for CWI

The point of our early activity in ATM is not only to confirm that protocols developed elsewhere actually work; they usually do, although often with a creative reading of the documentation supplied by the manufacturer. Instead, experiments are intended to be a catalyst for involving researchers in the development of ATM applications. By selecting user groups to participate in experimentation, and by anticipating their needs, we can continue to assist in the development of new applications and new solutions that will determine the next generation of information technology. The facility we provide will enable existing projects to examine aspects of multi-protocol use of the underlying network to solve existing pro-

Figure 2:
Amsterdam (left)
and The Netherlands (right)
on World Wide Web.



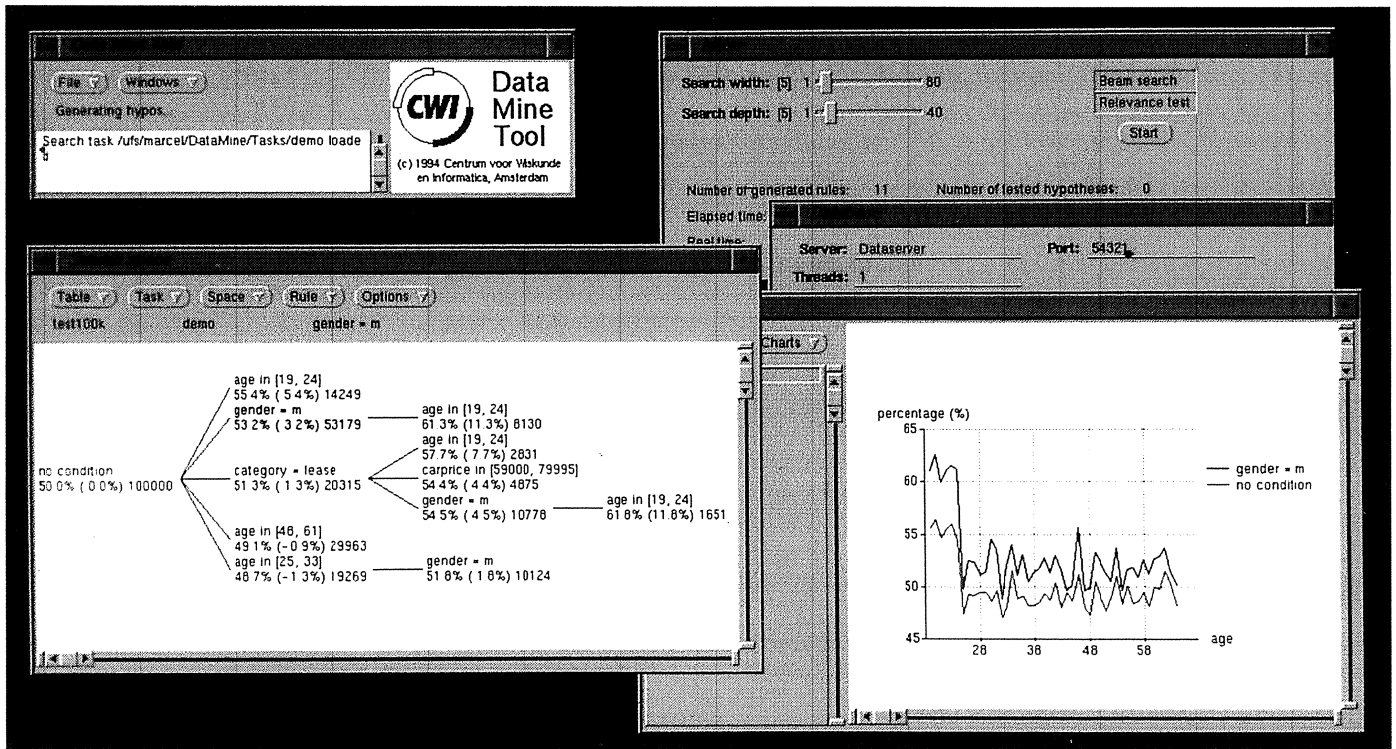
RESEARCH HIGHLIGHTS

blems. We also expect that the problems that can be studied using such a network will encourage researchers in both mathematics and computer science to work together to integrate theoretical, heuristic and systems solutions to solving common problems. Finally, we expect that the presence of a distributed, high-bandwidth network will make CWI an attractive partner in national and international research projects.

At present, CWI is participating in several wide-area networking experiments, on both the national and international levels. Unfortunately, the marketing policies of national PTT's severely limit the nature of wide-area connections available: at best, one can expect to get a 2- or 4-Mbps connection, even though the internal infrastruc-

ture can support connects in excess of 150 Mbps. This puts Dutch (and European) researchers at a distinct disadvantage in developing local approaches to solving wide-area information distribution problems. As a result, we can expect our colleagues across the Atlantic Ocean to maintain their lead in developing ATM protocols and application solutions. In the local-area network world, the situation for CWI is much better. Even with the modest infrastructure installed to date, new approaches to problems in scientific visualization and multimedia computing have been implemented. We expect these to be the first in a series of new network applications that will extend into the first decade of the new century.


COMPUTING EQUIPMENT RESOURCES



In general, data mining can be defined as the discovery of strategic information in large databases. In research at CWI this has been concretised as the search for 'interesting subsets' in a database. For an insurance company, such an interesting subset could consist of a group of clients with a far higher chance of causing an accident than the average insurant. As many topics in computer science, research in data mining has both theoretical and experimental components. In particular, the development of the data mining tool called Data Surveyor, is regularly assessed in pilot studies for industry. One of the pilot studies in 1994 was for an insurance company, with the goals given above. In the picture above, Data Surveyor is seen in action on a demo version derived from this case study (the real data and the results reached on that data are, of course, confidential). In particular, the tree depicted in the search space window shows which subsets have been

considered as interesting. To discover these potentially interesting subsets, Data Surveyor searches heuristically through the space of all possible subsets. This process encompasses the evaluation of a large collection of subsets. This implies that the underlying database management system has to evaluate tens of thousands of queries. To speed up this process (data mining has to be interactive) Data Surveyor uses the Monet database system, a main memory database system developed jointly by the University of Amsterdam and CWI. Monet ensures that the database hotspot is in main memory. In this way, Data Surveyor is hampered by I/O as little as possible. The result is that Data Surveyor lies large claims both on the available main memory and on the processing power of a computer. In pilot studies such as the insurance case, the capacity of high performance file servers such as the Zeus is easily stretched to their limits.

COMPUTING EQUIPMENT RESOURCES





CWI Computing Equipment Resources & Network Infrastructure as of December 31, 1994
 ktm=1000 triangular meshes
 Sm92=SPECmark92 integer&floating point performance

FINANCIAL AND OTHER DATA

FINANCES 1994



In 1994, SMC spent Dfl. 28.45 million, of which about Dfl. 2.96 million was allocated to university based research and Dfl. 25.49 million to CWI. The expenses were covered by a subsidy from NWO (Dfl. 24.30 million), other subsidies and grants (Dfl. 0.41 million), and from the international programmes (mainly EC programmes, e.g. BRITE, ESPRIT, SCIENCE and RACE) (Dfl. 2.79 million). Finally, an amount of Dfl. 1.99 million was obtained as revenues out of third-party-services and other sources. During 1994 CWI also hosted 96 researchers in externally financed positions. These are not included in the adjacent financial summary.

Income CWI

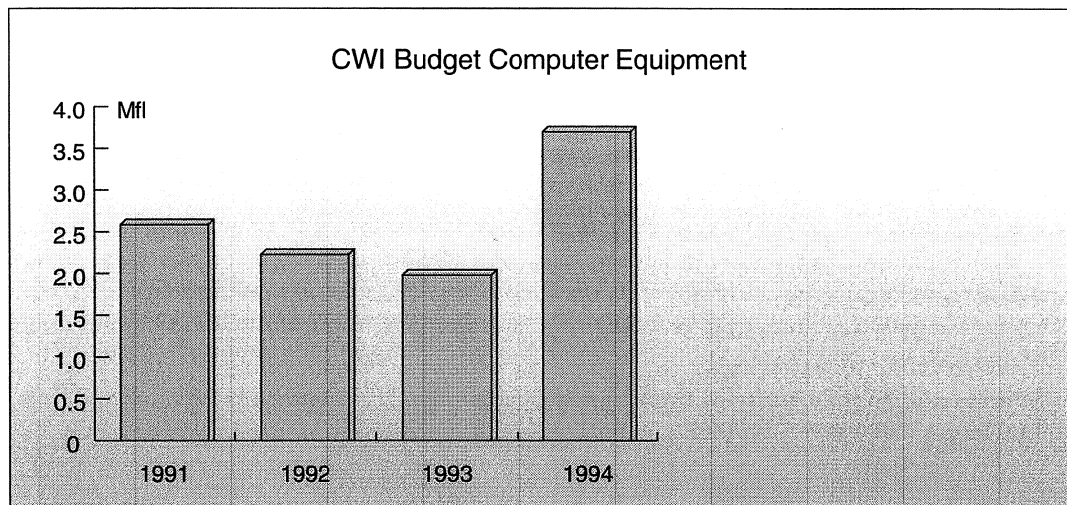
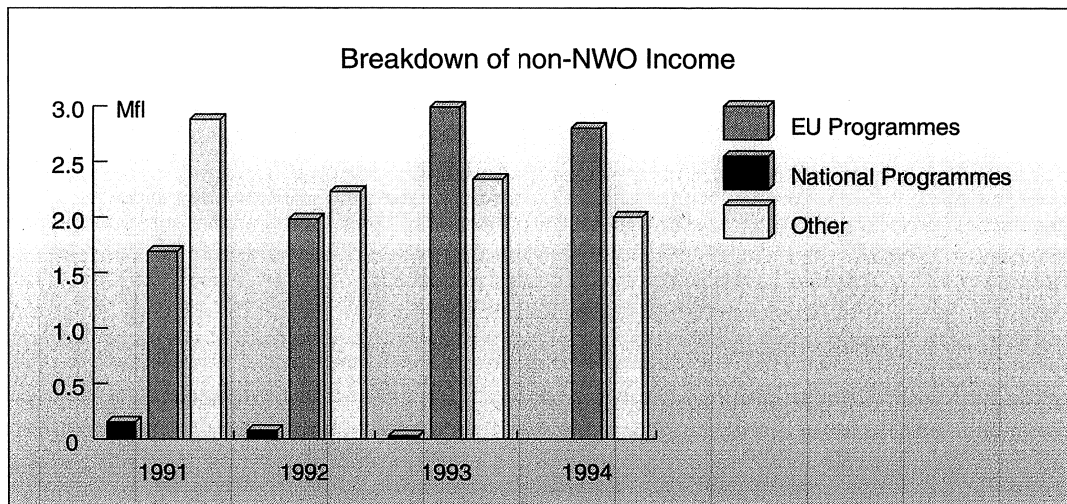



	<i>university based</i>	<i>CWI</i>	<i>SMC</i>
	* Dfl. 1000		
INCOME			
subsidy and grants			
- NWO	3170	21132	24302
- other	138	273	411
national programmes	-	-	-
international programmes	-	2791	2791
other revenues	-3	1993	1990
total income	3305	26189	29494
EXPENSES			
labour costs	2226	16581	18807
materials and overhead	254	2556	2810
investments	-	4785	4785
transfer to third parties	472	1684	2156
miscellaneous	-	-110	-110
total expenses	2952	25496	28448

Expenses CWI



FINANCES 1991 - 1994



CWI Ph.D. THESES

Author	Title	Thesis advisor(s) ⁺⁾
E.A. van der Meulen	Incremental Rewriting	P. Klint
R.A. Trompert	Local Uniform Grid Refinement for Time-dependent Partial Differential Equations	P.J. van der Houwen
M.S. Dijkhuizen	Compact Quantum Groups and Quantum Homogeneous Spaces	T.H. Koornwinder (UvA)
C.A. van den Berg	Dynamic Query Processing in a Parallel Object-oriented Database System	P.M.G. Apers (UT)
Nguyen huu Cong	Parallel Runge-Kutta-Nyström Methods	P.J. van der Houwen
A. van Deursen	Executable Language Definitions -- small Case Studies and Origin Tracking Techniques	P. Klint
M.H. van der Voort	A Design Theory for Database Triggers	M.L. Kersten
S.C. Borst	Polling Systems	O.J. Boxma
A.J. Cabo	Set Functionals in Stochastic Geometries	P. Groeneboom (TUD) A.J. Baddeley (Perth)
M.N. M. van Lieshout	Stochastic Geometry Models in Image Analysis and Spatial Statistics	A.J. Baddeley (Perth)
P.F.M. Nacken	Image Analysis Methods Based on Hierarchies of Graphs and Multi-scale Mathematical Morphology	F.C.A. Groen (UvA)
H. Korver	Protocol Verification in μ -CRL	J.C.M. Baeten (TUE) J.A. Bergstra (UvA)
F. van Breugel	Topological Models in Comparative Semantics	J.W. de Bakker
W.J. Fokkink	Clocks, Trees and Stars in Process Theory	J.C.M. Baeten (TUE) J.A. Bergstra (UvA)
J. de Does	The Gap Topology for Linear Systems	J.M. Schumacher

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

UvA = University of Amsterdam
 UT = University of Twente, Enschede
 TUD = Delft University of Technology
 Perth = University of Western Australia, Perth
 TUE = Eindhoven University of Technology

CWI RESEARCH PROGRAMMES

Algebra, Analysis & Geometry

Algebra, discrete mathematics and computer algebra

Algorithms and calculation techniques in selected areas of algebra and combinatorics, notably the narrowly intertwined cluster of objects: Lie algebras, Hopf algebras and Coxeter groups and the creation of computer based calculation and knowledge tools for mathematical research.

Projects:

- Algorithmic algebra and discrete mathematics
- Computer assisted mathematics

Group leader: M. Hazewinkel

Modelling and analysis

The mathematics of biology, dynamical systems and asymptotic techniques. In mathematical biology the emphasis is on dynamical biological phenomena, especially the dynamics of structured populations.

Projects:

- Population dynamics and epidemiology
- Dynamical systems
- Asymptotics

Group leader: O. Diekmann

Operations Research, Statistics & System Theory

Combinatorial optimization and algorithmics

Fundamental and applied research, with an orientation towards mathematics (discrete mathematics, geometry, number theory, probability theory), operations research (linear and integer programming, optimization, sequencing, scheduling), computer science (complexity theory, computational geometry) and applications (VLSI-layout, robotics, pattern recognition, railway and airplane routing and scheduling).

Projects:

- Design and analysis of algorithms

- Polyhedral methods and polynomial-time algorithms
- Multicommodity flows and VLSI-layout
- Applications

Group leader: A. Schrijver

Analysis and control of information flows in networks

Fundamental and application-oriented research concerning the behaviour of complex stochastic systems: mathematical analysis of queuing models; performance analysis of computer and communication networks; integration of queuing and reliability theory in order to assess the behaviour of systems subject to breakdown, replacement and repair; stochastic phenomena in lattice-type networks, including applications in mathematical physics and communications.

Projects:

- Analysis of mathematical queuing models
- Stochastic processes on networks
- Reliability and availability of networks
- Performance analysis and control of computer and communication networks

Group leader: O.J. Boxma

System and control theory

Formulation of system theoretic concepts, development of realization theory, synthesis and analysis of system identification algorithms, robust control, stochastic control and filtering, and supervisory control of discrete event systems and hybrid systems.

Projects:

- Stochastic system theory
- Deterministic system theory
- Control of discrete event systems
- System identification for compartmental models
- Control of distributed computer systems

Group leader: J.H. van Schuppen

Image analysis and spatial stochastics

Theory and applications of mathematical and sta-

tistical problems arising in the analysis of digital images and spatial stochastic processes, with emphasis on stochastic and geometric methods for stochastic models and dynamical systems.

Projects:

- Stochastic geometry
- Bootstrap resampling
- Random fields
- Mathematical morphology and discrete image transforms
- Ergodic theory of spatial processes

Group leader: M.S. Keane

Numerical Mathematics

Discretization of evolution problems

Fundamental and applied research into numerical methods and software for evolutionary differential equations. Both ordinary and partial differential equations are treated. Attention is given to (i) theoretical stability and convergence analysis, (ii) method of lines, (iii) large-scale applications, currently from circuit analysis, control engineering, atmospheric air pollution and shallow water pollution, and (iv) high performance computing, use of vector/parallel and massively parallel computers.

Projects:

- Equations of fluid mechanics and related topics
- Three-dimensional flux modelling in shallow water
- Parallel initial-value-problem algorithms
- Algorithms for atmospheric flow problems

Group leaders: P.J. van der Houwen, J.G. Verwer

Boundary-value problems, multigrid and defect correction

The numerical solution of boundary-value problems by multigrid and the method of defect correction, primarily special methods for 3D-problems, with adaptive grids and the fast solution of the corresponding systems; special attention is given to sparse grids. The major application is the compressible Euler- and Navier-Stokes equations, in particular for aerodynamics.

Projects:

- The analysis of defect correction and adaptive techniques
- Application of multigrid techniques to fluid dynamics problems
- Parameter identification in ordinary differential

equations

Group leader: P.W. Hemker

Large-scale computing

High performance scientific computing, focused on the optimization and comparison of mathematical and numerical algorithms on massively parallel processors, on parallel vector processors and on clusters of workstations; the development and use of tools for enhancing portability of parallel software, and for performance evaluation of parallel hardware.

Projects:

- Parallel numerical algorithms
- Computational number theory

Group leader: H.J.J. te Riele

Software Technology

Computational models

Mathematical modelling of a wide range of programming notions from contemporary computer languages, together with foundational studies of the applied techniques, and the study of biologically inspired computation models such as neural networks and genetic algorithms.

Projects:

- Semantics
- Planning based on neural networks and genetic algorithms

Group leader: J.W. de Bakker

Concurrency and real-time systems

Development and application of formal methods and tools for the specification and verification of timed distributed systems, both to control the complexity of such systems and to achieve the desired degree of correctness. Technically, the emphasis is on assertional verification techniques (the I/O automaton model) in combination with the use of general purpose theorem provers.

Projects:

- Structural operational semantics
- Testing and verification of timed systems
- Checking verifications of concurrent systems with type theory tools
- Transfer

Group leader: F.W. Vaandrager

Extensible programming environments

Generation of incremental programming environments from algebraic language definitions. More

specifically, the group is extending, maintaining, and promoting the ASF+SDF system for the interactive development of programming and application languages.

Projects:

- Generation of interactive programming environments
- Generic tools for program analysis and optimization
- Parallel rewriting on SP1

Group leader: P. Klint

Algebraic and syntactic methods

Foundational research in term rewriting systems, with an emphasis on higher-order term rewriting and term graph rewriting (including infinitary term rewriting). Applied research aims at making bridges between facilities for term rewriting, computer algebra, and proof checking.

Projects:

- Higher-order rewriting
- Term graph rewriting and infinitary rewriting
- WINST (Mathematics and Informatics Cooperation Themes)

Group leader: J.W. Klop

Logic and language

The use of mathematical logic in modeling concepts from artificial intelligence and computational linguistics, in connection with issues in knowledge representation, common sense reasoning and semantics of various linguistic constructs. Secondly, the use of computational models developed by computer scientists in connection with logic programming, formal aspects of Prolog, program verification, non-monotonic reasoning, natural language processing and the parallels between natural and programming languages.

Projects:

- Logic programming and non-monotonic reasoning
- Verification of Prolog, logic and constraint logic programs
- Structural and semantic parallels in natural languages and programming languages
- Incremental parser generation and disambiguation in context
- A framework for computational semantics

Group leaders: K.R. Apt, D.J.N. van Eijck

Algorithmics & Architecture

Algorithms and complexity

Algorithmic aspects of non-conventional computer networks, machine learning and distributed information systems, covering the design, construction and use of hardware and software, as well as applications. In particular, realistic models for multi-computers, design and analysis of algorithms suitable for distributed computations, theory and development of systems of multiple computing agents (neurocomputing and genetic programming), and novel modes of computation (quantum computing).

Projects:

- Parallel and distributed algorithms
- Multiple computing agents

Group leader: P.M.B. Vitányi

Cryptography

The research concerns all aspects of cryptology related to information security. This involves the construction and analysis of cryptographic protocols and their underlying cryptographic algorithms, and mathematical proofs of their soundness and reliability. Emphasis is placed on the security and privacy of individuals in protocols for the transmission of messages, payment systems, and handling of personal data by organizations.

Projects:

- Public-key cryptography
- Specification of cryptographic protocols
- CAFE (Conditional Access For Europe)

Group leader: R. Hirschfeld

Interoperable multimedia systems

Development of formalisms and tools for composing software systems from various independently developed multimedia components, in particular: the specification of a transportable, multi-machine environment in which both passive and active multimedia components can be integrated into a 'multimedia document'; protocol rules for adaptive communication of multimedia data; and distributed resource allocation algorithms.

Projects:

- Wide-area distributed interoperability
- Interactive books

- Constructive algorithms
 - Multimedia document specification and construction
 - OS-based constraint resolution for multimedia applications
 - CMIFed native C++ implementation
- Group leaders: D.C.A. Bulterman, L.G.L.T. Meertens

Databases

Research on database design theory and architectures for advanced database management systems. In particular, experimentation with novel architectures to exploit the potential parallelism of database management in shared-store and distributed-store processor systems. In addition: the study of design and architecture of active databases; data mining; and parallel database management architectures.

Projects:

- Parallel DBMS architectures
- Query optimization
- Active databases
- Data mining
- Performance analysis

Group leader: M.L. Kersten

Interactive Systems

Computer graphics and visualization

Research on computer graphics, visualization and image processing with a strong commitment to approaching these areas from a Human-Computer Interaction point of view. The research covers application-driven exploitation of fundamental techniques. In particular: adaptive image synthe-

sis; wavelet based multi-resolution image coding and decoding techniques; and interactive control of numerical simulations.

Projects:

- Computer graphics
- Image coding and decoding
- Computational steering

Group leader: A.A.M. Kuijk

Interaction and parallelism

Development of conceptual models and practical languages for coordination of interactions among a potentially large number of co-operative concurrent processes that make up a single application.

Projects:

- MANIFOLD language and system
- Visual parallel programming and constraint solving
- Meta-computing and parallel applications

Group leader: F. Arbab

Interaction and multimedia

The creation of basic facilities for media integration, in particular the design of an application-specific concrete set of multimedia products which represent information perceived as units. Work in the ESPRIT project MADE will provide the basis for multimedia object creation and authoring, and introduce advanced, novel object-oriented techniques for this purpose.

Projects:

- Multimedia fundamentals
- Multimedia systems
- Multimedia applications

Group leader: P.J.W. ten Hagen

INTERNATIONAL AND NATIONAL PROGRAMMES

This chapter summarizes the major national and international projects in which CWI participates. While participation in European research projects (e.g. ESPRIT) remained on the same level in 1993, participation in European research networks and national projects increased substantially.

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

COMPARE (5399): Compiler Generation for Parallel Machines
January 1991 - April 1995
Ace BV, STERIA, GMD, INRIA, Harlequin Ltd, Univ. Saarland
P. Klint

MADE (6307): Multimedia Application Development Environment
May 1992 - December 1995
Bull SA, SNI, Iselqui, British Aerospace, INESC, Gipsi SA, ESI, Barclays Bank, NR, FhG-IAO, INRIA
P.J.W. ten Hagen

PEPS (6942): Performance Evaluation of Parallel Systems
November 1992 - April 1994
Thomson Sintra, Intecs Sistemi, Simulog, Univ. of Warwick, NPL (UK), AFNOR (France), CNR (Italy), PTB (Germany)
J. Kok

CAFE (7023): Conditional Access for Europe
December 1992 - December 1995

Digicash, PTT, Cardware, Gemplus, SEPT, Ingenico, SINTEF-Delab, Institut für Sozialforschung Frankfurt, Institut für Informatik Hildesheim, Siemens, Universities of Leuven and Aarhus
Coordinator
R. Hirschfeld

PYTHAGORAS (7091): Performance Quality Assessment of Advanced Database Systems
May 1992 - November 1995
ICL, Bull SA, Heriot-Watt Univ., CCIP, Infosys, IFATEC
Coordinator
M.L. Kersten

ESPRIT Basic Research

SEMAGRAPH II (6345)
October 1992 - October 1995
Univ. East Anglia, ECRC GmbH, Univ. Rennes, Univ. Nijmegen, Imperial College
J.W. Klop

CONFER (6454): Concurrency and Functions: Evaluation and Reduction
September 1992 - September 1995
INRIA Rocquencourt, ECRC GmbH, Univ. Edinburgh, CNRS-ENS, Imperial College, INRIA Sophia Antipolis, Univ. Pisa, SICS
J.W. Klop

COMPULOG II (6810): Formal Aspects of Prolog and Logic Programming
August 1992 - August 1995
Univ. Leuven, ECRC GmbH, RWTH Aachen, Univ. Saarland, Univ. Pisa, Univ. Rome (La Sapienza), Univ. Rome (Tor Vergata), UNINOVA Lisbon, Univ. Uppsala, Imperial College, Universities of Bristol, Edinburgh and Aix-Marseille II
Coordinator
K.R. Apt

CONCUR 2 (7166): Calculi and Algebras of Concurrency: Extensions, Tools and Applications

September 1992 - September 1995
Universities of Eindhoven, Aalborg, Edinburgh,
Sussex and Oxford, INRIA, SICS, INPG, Sharp,
Chalmers Univ., ECRC
F.W. Vaandrager

QMIPS (7269): Quantitative Modelling In Parallel Systems
October 1992 - October 1995
Univ. René Descartes LAA, Univ. Erlangen-Nürnberg, Univ. Torino, Imperial College, Univ. Newcastle, INRIA Sophia Antipolis
O.J. Boxma

NeuroCOLT (8556): Neural and Computational Learning
January 1994 - January 1997
Royal Holloway and Bedford New College, Univ. Mons, Rheinisch-Westfälische Tech. Hochschule, Univ. Pompeu Fabra, Techn. Univ. Graz, London School of Economics, Helsingin Yuopisto, Lab. de l'Informatique du Parallelisme, Univ. Milan P.M.B. Vitányi

BRITE/EURAM

AERO II (AER2-CT92-0040): Solution adaptive Navier-Stokes solvers using multidimensional upwind schemes and multigrid acceleration
January 1993 - January 1996
Von Karman Institute for Fluid Dynamics, Free Univ. Brussels, Politecnico di Bari, Technical Univ. Denmark, Royal Institute of Technology, Dornier Deutsche Aerospace, Fokker Aircraft B.V., Aerospaziale, British Aerospace, Dassault Aviation
P.W. Hemker

DRIVE

DYNA (V2036): A Dynamic traffic model for real-time applications
January 1992 - January 1995
Hague Consulting Group, CSST, Univ. Naples, Elasis, RWS, Univ. Lancaster, Univ. Libre Bruxelles, Univ. Delft
J.H. van Schuppen

MAST Marine Science and Technology

NOWESP: North-West European Shelf Programme
September 1993 - September 1996
RWS, Institut für Meereskunde, Univ. Leuven, NIOZ, Proudman Oceanographic Laboratory Bridston, Sir Allister Hardy Foundation for Ocean Science, Institute of Marine Research, Inst. für Ostseeforschung, Delft Hydraulics, BSH, IfBM, IFREMER, MUMM, Univ. Delft, Trinity College, Universities of Bordeaux and Liverpool
P.J. van der Houwen

RACE

BOOST (2076): Broadband Object-Oriented Service Technology
January 1992 - January 1995
MARI Computer Systems Ltd, IPSYS Software Plc, Bull S.A., Société Française de Génie Logiciel S.A., GIE Emeraude, Detecon Technisches Zentrum, Intrasoft S.A., Telefonica, Intecs Sistemi Spa, Standard Elektrik Lorenz AG, Alcatel SEL, Centro de Estudos de Telecomunicações, Univ. College of Wales, Universities of Athens and Aveiro
F.W. Vaandrager

LRE

FRACAS: A Framework for Computational Semantics
January 1994 - April 1996
Univ. Edinburgh, Univ. Saarland, Univ. Stuttgart
J. van Eijck

Libraries Programme

RIDDLE(1038): Rapid Information Display and Dissemination in a Library Environment
February 1993 - February 1995
Longman Cartermill Ltd, Rutherford Appleton Laboratory
F.A. Roos

VALUE

POWER: Performance Oriented Workbench Experiment on Real Information Systems in the Energy Field

1994-1995

IFATEC

M. Kersten

Promotional Activities related to the ESPRIT

7023 project CAFE

1994-1995

Cardware

R. Hirschfeld

*SCIENCE***MASK:** Mathematical Structures in Semantics for Concurrency (CT92-0776)

September 1, 1992 - September 1, 1995

Univ. Pisa, CNRS/INRIA, Universities of Udine, Mannheim and Koblenz

Coordinator

J.J.M.M. Rutten/J.W. de Bakker

System Identification: Modeling, Realization and Parameter Estimation for Problems of Engineering, Economics and Environmental Science
July 1992 - June 1995

Univ. Groningen, Technical Univ. Wien, Univ. Leuven, INRIA, Univ. Rennes I, Univ. Cambridge, LADSEB-CNR, Linköping Univ.

CWI participates through the Systems & Control Theory Network of Univ. Groningen, seat of the coordinator

J.H. van Schuppen

*Human Capital and Mobility Networks***EXPRESS:** Expressiveness of languages for concurrency (CT93-0406)

1994-1997

Free Univ. Amsterdam, SICS, Univ. Genova, Univ. Rome (La Sapienza), Univ. Hildesheim, Univ. Amsterdam, INRIA, GMD, Univ. Sussex

Coordinator

F.W. Vaandrager

Statistical inference for stochastic processes (CT92-0078)

1993 - 1996

Universities of Paris VI, Berlin, Aarhus and Freiburg, INRIA

K.O. Dzhaparidze

EUROFOCS: European institute in the logical

foundations of computer science (CT93-0081)

1994-1996

Univ. Edinburgh, INRIA, Universities of Pisa and Cambridge, ENS

J.W. de Bakker

The equations of fluid mechanics and related topics (CT93-0407)

1994-1996

CMAP, Univ. Paris VI, IX, XIII, Univ. Pisa, Univ. Ferrara, Univ. Nantes, IST (Lisbon), Universities of Trento, Pavia, Grenoble, Coimbra, London and Valladolid

J.G. Verwer

Algebraic combinatorics (CT93-0400)

1993-1996

Univ. Magdeburg, KTH Stockholm, Univ. Perugia, Univ. Cagliari, Univ. Bielefeld, Univ. Strasbourg, Univ. Bayreuth, Univ. Vienna, Univ. Paris VI, Univ. College of Wales, Universities of Copenhagen, Erlangen and Bordeaux I, Konrad Zuse Inst.

M.A.A. van Leeuwen

DONET: Discrete optimization and applications (CT93-0090)

1993-1996

Univ. Joseph Fourier, ZOR Bonn, Univ. Oxford

A. Schrijver

DIMANET: Discrete Mathematics Network (CT94-0429)

1994-1996

Universities of Bielefeld, Bologna, Cambridge, Montan, Lisbon, Madeira, Milan,

Oxford, Paris 1, and Umea, Konrad Zuse Zentrum, Danmarks Tekniske Univ., Techn.

Hochschule Darmstadt, Ecole Polytechnique Lausanne, Queen Mary and Westfield

College, ENS Lyon, CNRS, Kungliga Tekniska Hogskolan

A. Schrijver

ERCIM computer graphics network (CT93-0085)

1993-1996

P.J.W. ten Hagen

ERCIM advanced databases technology network

1994-1997

M.L. Kersten

INTAS

ERCIM-FSU Cooperative Network in Informatics and Applied Mathematics

1994-1995

M. Hazewinkel

Network Mathematical Methods for Stochastic Discrete Event Systems

1994-1995

O.J. Boxma

European Science Foundation Networks

Dynamics of complex systems in biosciences

O. Diekmann

Highly structured stochastic systems

A.J. Baddeley

National Programmes

SION (Netherlands Foundation for Computer Science)

Incremental program generators

1990-1994

P. Klint

Mathematical morphology in hierarchical graph representations of images

1990-1994

Inst. voor Zintuigfysiologie TNO, Univ. Amsterdam

H.J.A.M. Heijmans

Design implementation and application of a transparent distributed computing system

1989-1994

Univ. Twente

M.L. Kersten

Nonwellfounded sets and semantics of programming languages

1991-1995

J.J.M.M. Rutten

Extensions of orthogonal rewrite systems - syntactic properties

1992-1995

J.W. Klop

Computational Learning Theory

1992-1996

P.M.B. Vitányi

Declarative and procedural aspects of non-standard logics

1992-1996

K. Apt

MathViews - Functional and architectural aspects of mathematical objects in an Integrated System

1992-1996

A.M. Cohen

From ideas to reality - Implementing cryptography

1994-1998

R. Hirschfeld

WINST: Themes for collaboration in mathematics and computer science

1994-1998

Universities of Nijmegen and Eindhoven

H.P. Barendregt, J.W. Klop, M. Hazewinkel,

A.M. Cohen

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

INTERNATIONAL AND NATIONAL PROGRAMMES

- F.W. Vaandrager
Universities of Eindhoven, Leiden, Limburg and Twente
- Constraints in object-oriented interactive graphics
1994-1998
Univ. Eindhoven
P.J.W. ten Hagen
L.G.L.T. Meertens
- Higher-order and object-oriented processes (HOOP)
1994-1999
Universities of Eindhoven and Leiden
J.W. de Bakker
ALADDIN - Algorithmic Aspects of Parallel and Distributed Computing
1992-1996
Univ. Utrecht
P.M.B. Vitányi
- MAGNUM, Database technology for multimedia information systems
1994-1998
Universities of Twente and Amsterdam
M.L. Kersten
Systematic design of user interfaces
1990-1994
Free Univ. Amsterdam, Universities of Delft, Brabant and Twente
L.G.L.T. Meertens
- Constraint-based graphics
1994-1996
Univ. Eindhoven
F. Arbab
Distributed Algorithms
1993-1995
P.M.B. Vitányi
- Cryptography, learning and randomness
1994-1996
Univ. Amsterdam
P. Vitányi, R. Hirschfeld
Special NWO projects
- NFI (National Facility Computer Science)*
- Performance analysis and control of distributed computer systems
1990-1995
O.J. Boxma/J.H. van Schuppen
AIDA: Algorithms in algebra
1993-1996
Universities of Eindhoven, Groningen and Twente
A.M. Cohen
- Structural and semantic parallels in natural languages and programming languages
1991-1995
Univ. Amsterdam, OTS, Univ. Utrecht
D.J.N. van Eijck
Nonlinear systems
1993-1996
Universities of Groningen, Delft, Utrecht, Wageningen and Leiden, KSLA
O. Diekmann
- Intelligent CAD systems
1986-1994
TNO/IBBC, Univ. Amsterdam
P.J.W. ten Hagen
Computationally intensive methods in stochastics
1993-1996
Universities of Leiden, Amsterdam, Rotterdam, Utrecht and Nijmegen
M.S. Keane
- Formal methods for the description of information systems and the analysis of information systems descriptions (ISDF)
1989-1994
Application of numerical methods for singularly perturbed PDE's to the mathematical analysis and simulation of environmental problems in the Ural
1992-1994
Ural branch of the Russ. Acad. of Sciences, Ural State Univ., Nuclear Safety Institute Moscow, Univ. College Cork, Inst. for Numerical Computation and Analysis Dublin
P.W. Hemker

STW (Foundation for the Technical Sciences)

Parameter identification and model analysis for
non-linear dynamic systems

1993-1997

P.W. Hemker

Parallel codes for circuit analysis and control
engineering

1993-1997

Univ. Amsterdam

P.J. van der Houwen

ACELA - Architecture of a Computer Environ-
ment for Lie Algebras

1993-1996

Univ. Eindhoven

A.M. Cohen, L.G.L.T. Meertens

NCF

CIRK: Mathematical modelling of global trans-
port and chemistry of trace constituents in the
atmosphere

1994

Univ. Utrecht, KNMI

J. Verwer

Three dimensional transport of pollutants in shal-
low seas

1994

J. Verwer

RESEARCH STAFF

Analysis, Algebra and Geometry

M. Hazewinkel (head of department)

M. Biemond

A.E. Brouwer

A.M. Cohen (advisor)

O. Diekmann

N. Elhoussif

F.C.A. Groen (advisor)

T.W. Hantke

N. van den Hijligenberg

R.A. Hirschfeld

R. Hoksbergen

A.A. de Koeijer

Yu. A. Kuznetsov

A.M.A. van Leeuwen

M.A.A. van Leeuwen

J.A.J. Metz (advisor)

G.H.M. Roelofs

A.M. de Roos

J.A. Sanders

N.M. Temme

S. Verduyn Lunel

J. de Vries

O.J.M. Weber

Programmer:

B. Lisser

Operations Research, Statistics and System Theory

O.J. Boxma (head of department)

N. Bayer

J. van den Berg

S.C. Borst

R.J. Boucherie

J. Coelho de Pina

J.W. Cohen (advisor)

M.B. Combé

F.A. van der Duyn Schouten

K.O. Dzhaparidze

A. Ermakov

A.M.H. Gerards

H.J.A.M. Heijmans

R. Helmers

J.M. van den Hof

H. van der Holst

R.H.P. Janssen

M.S. Keane

M. Laurent

J.K. Lenstra (advisor)

A.A.F. Overkamp

F.K. Potjer

J. Rosenthal

A.J. van der Schaft

A. Schrijver

J.M. Schumacher

J.H. van Schuppen

A.A. Stoorvogel

S.J. van Strien

P.R. de Waal

Programmers:

A.G. Steenbeek

R. van der Horst

Trainee:

S. Vos de Wael

Numerical Mathematics

P.J. van der Houwen (head of department)

J.G. Blom

H. Boender

J.G.L. Booten

P.W. Hemker

R.-M. Huizing

W.H. Hundsdorfer

J. Kok

B. Koren

M. van Loon

H.J.J. te Riele

B.P. Sommeijer

E.J. Spee

W.J.H. Stortelder

J.J.B. de Swart

W.A. van der Veen

J.G. Verwer

H.A. van der Vorst (advisor)

P. Wesseling (advisor)

Programmers:

C.T.H. Everaars

W.M. Lioen

M. Nool

D.T. Winter

P.M. de Zeeuw

Software Technology

J.W. de Bakker (head of department)

K.R. Apt	M. Gabrielli	A.S. Klusener
H.P. Barendregt	J.J. Ganzevoort	E. Marchiori
J.A. Bergstra (advisor)	W.O.D. Griffioen	F. van Raamsdonk
I.M. Bethke	A.V. Groenink	J.J.M.M. Rutten
D.J.B. Bosscher	J. Heering	M. de Rijke
J. Brunekreef	B.P.F. Jacobs	F. Teusink
T.B. Dinesh	J. Jaspars	F. Tip
H. Elbers	J.F.Th. Kamperman	D. Turi
S. Etalle	P. Klint	F.W. Vaandrager
D.J.N. van Eijck	J.W. Klop	M. van Wezel
W.J. Fokkink	H.P. Korver	H.R. Wiklicky

Algorithmics and Architecture

M.L. Kersten (head of department)

J.F.P. van den Akker	J.-H. Hoepman	A.P.J.M. Siebes
A. Berthiaume	M. Holsheimer	M. Theodoridou
A.M. Bleeker	J. Keller	C.J.E. Thieme
H.M. Buhrman	F. Kwakkel	P.M.B. Vitányi
D.C.A. Bulterman	H.A.N. van Maanen	O.J.M. Weber
R.J.F. Cramer	L.G.L.T. Meertens	
H.H. Ehrenburg	S.J. Mullender (advisor)	<i>Programmers:</i>
C. Galinda-Legario	J. Pellenkofft	F. van Dijk
P. Gruenwald	S. Pemberton	A.J. Jansen
L. Hardman	G. van Rossum	K.S. Mullender
R. Hirschfeld	L.A.M. Schoenmakers	

Interactive Systems

P.J.W. ten Hagen (head of department)

F. Arbab	R.H.M.C. Kelleners	<i>Programmers:</i>
P.J. Bouvry	A.A.M. Kuijk	C.L. Blom
P.A. Griffin	R. van Liere	F.J. Burger
M. Guravage	J.D. Mulder	C.T.H. Everaars
M. Haindl	G.J. Reynolds	H. Noot
F.C. Heeman	T. van Rij	M.M. de Ruiter
I. Herman	R.C. Veltkamp	
J.E.A. van Hintum	J.J. van Wijk	

ADVISORY COMMITTEES CWI

Analysis, Algebra and Geometry

G. van Dijk (RUL)
M.A. Kaashoek (VUA)
E.J.N. Looijenga (UvA)
L.A. Peletier (RUL)
M. van der Put (RUG)
E.G.F. Thomas (RUG)

Operations Research, Statistics and System Theory

R.D. Gill (RUU)
P. Groeneboom (TUD)
H.C. Tijms (VUA)
H.L. Trentelman (RUG)
W.H.M. Zijm (UT)

Numerical Mathematics

A.O.H. Axelsson (KUN)
M.N. Spijker (RUL)
H.A. van der Vorst (RUU)
G.K. Verboom (WL, Delft)
T.M.M. Verheggen (KSLA)

Software Technology

H. Brinksma (UT)
C. Hemerik (TUE)
H.J. van den Herik (RL)
R.L.C. Koymans (PLN)
J. Landsbergen (IPO)
J.-J.Ch. Meyer (UU)
C.A. Middelburg (PTT Research)
A. Ollongren (RUL)
W.P. Weijland (PTT, Telecom)

Algorithmics and Architecture

H. Brinksma (UT)
H.H. Eggenhuisen (Philips Natlab)
S.D. Swierstra (RUU)
L. Torenvliet (UvA)

Interactive Systems

F.W. Jansen (TUD)
G.R. Joubert (TU, Clausthal, Germany)
F. Klok (Philips)
W. Loeve (NLR)
C.W.A.M. van Overveld (TUE)
H.J. Sips (TUD)

