ERCIM NEWS

European Research Consortium for Informatics and Mathematics
www.ercim.org

Number 46

July 2001

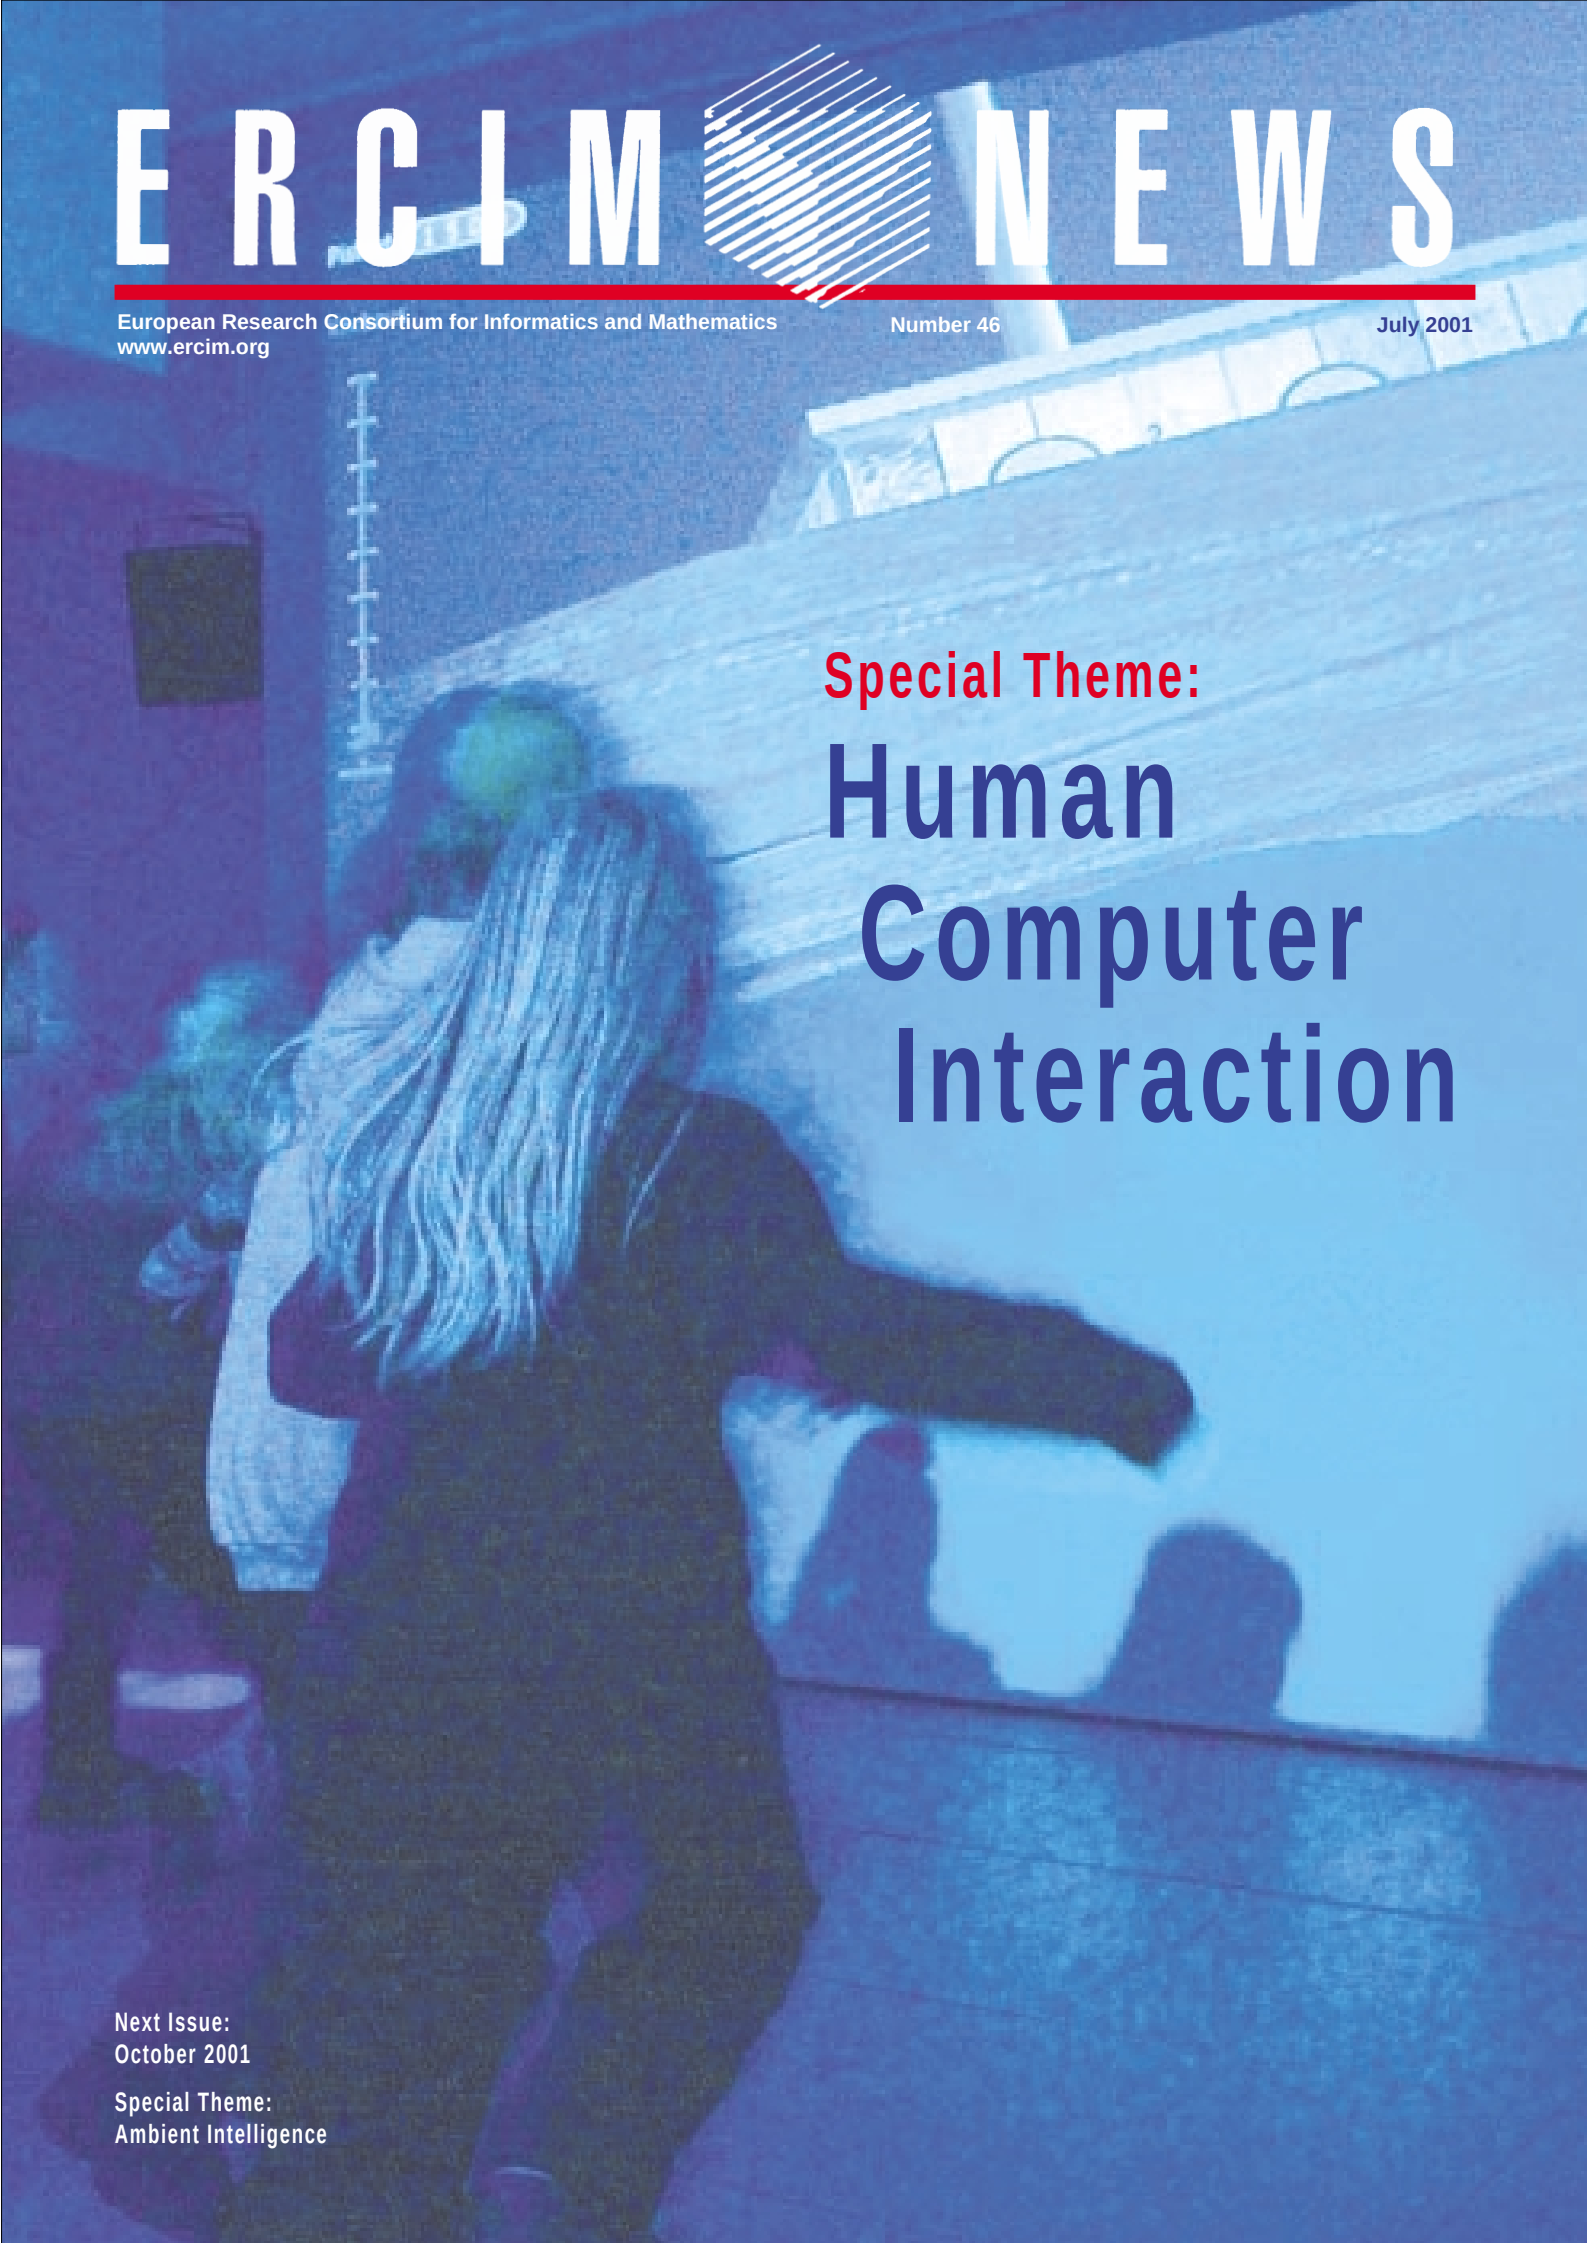

Special Theme:

Human Computer Interaction

Next Issue:
October 2001

Special Theme:
Ambient Intelligence

CONTENTS

Cover image: Novel User Interfaces for the Future — A group of school children steering the diving bell Nautilus. See article on page 24.

Gavriel Salvendy is Professor at Purdue University and the author of 400 publications, including 23 books, and is a member of the National Academy of Engineering.
E-mail: salvendy@ecn.purdue.edu

Human-computer Interaction (HCI) entails the study of physical, social, cognitive and engineering aspects of designing information technology for ease of use. The field has experienced tremendous growth over the past two decades, by making use of, and extending, basic science, especially engineering and the social sciences. Today, there are over 40,000 HCI professionals worldwide, two-thirds of whom are computer scientists and one-third behavioral scientists. Their work is being published in over 17 journals, presented in over 20 annual conferences, and published in over 25 books annually.

In the 1980s and '90s HCI was the prime productivity enhancing force which led the penetration of the computer towards a wide range of applications areas such as the business environment, health care, and education, thus considerably advancing the quality and standards of life. This momentum is continuously growing in the Information Society era, setting new targets and broadening the type, nature and scope of computer mediated human activities.

The European Commission (EC) plays a pivotal role in making information technology accessible by all individuals at all times. This is made possible by the visionary Information Society technologies programme of the European Commission, and in particular by the commitment towards a user-friendly Information Society.

A significant number of results from EC-funded projects will be presented at HCI International 2001, held jointly with the Symposium on Human Interface (Japan) 2001, the 4th International Conference on Engineering Psychology and Cognitive Ergonomics, and the 1st International Conference on Universal Access in Human-Computer Interaction, in New Orleans, USA, 5-10 August 2001 (http://hcii2001.engr.wisc.edu).

The past twenty years of HCI research and development have been pushed by advances of technology. It is foreseen that the next 20 years of HCI will be characterized by HCI research and development requirements pushing technology. The ultimate goal of HCI should be to establish an independent international laboratory to evaluate and test the universal accessibility and usability, and joy of use of each information technology product and service before it is released to the market, and assign a usability index (from 0 to 100) for each information technology product or service with which individuals need to interact. This will provide valuable information to the consumers, and thus enable them to make intelligent informed decisions regarding the selection of information technology for their needs. It is envisioned that this evaluation will be performed automatically by embedded software in the test-computer, thus also providing intelligent output as to where, how and what needs to be changed in order to increase the index of accessibility, usability and joy in using the information technology. This information would prove most helpful to developers and manufacturers of information technology, and push HCI to new scientific and professional heights to the benefit of mankind.

Gavriel Salvendy

This is the first of a series of articles which from now on will appear in ERCIM News, reporting on the past and upcoming activities of the DELOS Network of Excellence on Digital Libraries covering a three-month period. Further information on the objectives, activities and membership of DELOS can be accessed from the DELOS home page: http://www.ercim.org/delos/

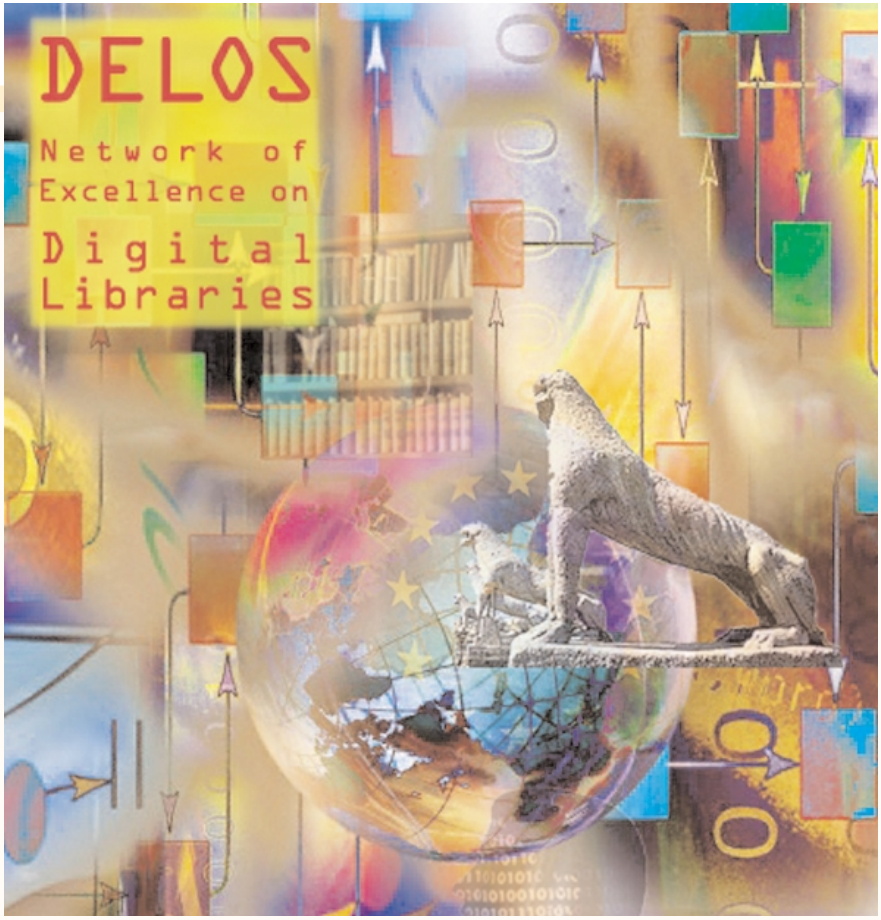

Digital Libraries National Initiatives and EU Actions: The Need for Harmonization

An EU-DELOS Workshop on Harmonization of European and National Initiatives in Digital Libraries took place on 11 May 2001 in Brussels.

Currently, the situation in Europe in the field of digital libraries varies widely from country to country. A number of national initiatives are either under way or under definition. The EU 5th Framework Programme has also funded a large number of DL projects. However, so far, all these initiatives have proceeded independently without interactions between individual projects or with the EU programmes. The goal of the meeting was therefore the identification of means which could make the harmonization of national and EU DL initiatives feasible. Other important goals of the meeting were to provide valuable input for the definition of appropriate actions in the DL field under the 6th Framework Programme and the establishment of bilateral relationships between those initiatives which share specific goals. To this end high-level representatives of national funding agencies and/or DL national initiatives, and representatives of EU programmes were invited to meet and exchange their views. Organizations from the following countries were represented: UK, Germany, Italy, Denmark, Portugal, Finland, Norway, France and Greece. The European Space Agency was also represented.

The meeting was organized into two sessions. In the first session each national initiative gave a presentation on its current achievements and/or planned programs. In the second session the crucial question of how to proceed was addressed. The central point, shared by all participants, was the identification of joint actions which are important to begin now, but which should be fully deployed under the forthcoming 6th Framework Programme.

Digital Libraries: Future Research Directions for a European Research Programme

A DELOS brainstorming meeting took place on 13-15 June 2001 in San Cassiano (Dolomites), Italy.

Digital libraries is an important research field with many technologies contributing to its progress (databases, information retrieval, networks, multimedia, user interfaces, artificial intelligence, visual computing, multilinguality, etc.). The field is of particular importance for Europe given the enormous richness of its cultural heritage and scientific content. This has been recognized by the 5th Framework Programme which has funded a number of projects in this area. With respect to digital libraries, the challenge of the 6th Framework Programme is to define a new vision in which the convergence of new information infrastructures and advances in digitalization and other enabling technologies can provide new models for making available, using

and enjoying them. The objective of the meeting was therefore to outline the main research directions of a future European research programme in the field of digital libraries. Prominent members of the European research community were thus invited to participate in and contribute to the definition of this new vision of digital libraries, and to translate it into concrete research actions which will become an integral part of the next 6th FP. Some prominent members from the US research community were also invited, in collaboration with the National Science Foundation. Some 20 EU and US high-level researchers attended.

DELOS-Russia Workshop

The first DELOS-Russia collaboration workshop took place in Moscow, Russia on 7-8 June 2001.

The goal of the workshop was to evaluate possibilities for establishing collaborations between EC-funded groups and Russian partners in the field of digital libraries and to identify key issues for such collaborations. In addition, the UNESCO Institute for Information Technologies in Education in Moscow expressed its interest in supporting this workshop, and therefore a session was held on 9 June dealing mainly with Digital Libraries in Education, with additional participants invited by UNESCO.

Collaboration with China

Costantino Thanos, Director of the DELOS NoE, was invited to present the main DELOS activities at the 12th International Conference on New Information Technology (NIT2001) held on 29-31 May 2001 in Beijing, China. Future possibilities for further collaboration were explored.

International Summer School on Digital Library Technologies

The first DELOS International Summer School on Digital Library Technologies will take place in Pisa, Italy on 9-13 July 2001.

The main purpose of the school is to foster research in and understanding of the fundamental technologies underlying the digital libraries field. The school is directed towards members of the research community in the wide sense, that is, graduate students, young researchers and professionals involved in R&D in DL-related areas, possibly representing both the information technology scientist, the industrial communities (electronic publishing, broadcasting, software industry, etc.) and the user communities interested in digital libraries technologies (libraries, archives, museums). The school will consist of a one-week intensive course of nine half-day lectures and one half-day dedicated to BOF (Birds of a Feather) sessions, ie discussions in smaller groups on specific topics of common interest, with the participation of the lecturers.

Workshop on Personalization and Recommender Systems in Digital Libraries

The second workshop on Personalization and Recommender Systems in Digital Libraries took place in Dublin, Ireland on 18-20 June 2001.

One of the important ways for users to feel comfortable with and become productive using information technology is the ability to personalise and tailor systems to individuals or groups of users. This covers both explicit personalisation directly by the user, and implicit tailoring by systems which track users usage patterns and preferences and adapt systems and interfaces accordingly. The concept of personalisation thus is about making systems different for individual people, but the concept of personalisation itself can mean different things. One type of personalisation that is growing in use is recommender systems. Such systems take input directly or indirectly from users and based on user needs, preferences and usage patterns, recommender systems will then make personalised recommendations of products or services. These vary from recommending books to buy or TV programs to watch, to suggesting web pages to visit. The ultimate goal of such recommender systems would be to reduce the amount of explicit user input and to operate, effectively, based on usage patterns alone, thus giving users what they want without them having to ask. The workshop focus will be on techniques and technologies that are now applied to personalisation and recommender systems in digital libraries, or that have a reasonable potential of applying within a few years.

Tutorial on Multilingual Access for Information Systems

A DELOS tutorial on Multilingual Access for Information Systems will be held at IFLA 2001 (International Federation of Library Associations and Institutions) taking place in Boston, USA on 16-20 August 2001.

Much attention has been given over the past few years to the study and development of tools and technologies for multilingual information access (MLIA) and cross-language information retrieval (CLIR). This is a complex, multidisciplinary area in which methodologies and tools developed in the fields of information retrieval and natural language processing converge. Two main sectors are involved: multiple language recognition, manipulation and display; multilingual or cross-language search and retrieval.

The tutorial will provide participants with an overview of the main issues of

interest in both these sectors. Topics covered will include: the specific requirements of particular languages and scripts, processing of multilingual document collections, multilingual metadata, techniques for cross-language retrieval, presentation issues, system evaluation. The focus of the tutorial will be on the needs of the user—whether information seeker, provider, or administrator—and on how today's multilingual retrieval systems can impact on each of these user types.

The third DELOS workshop on Interoperability in Digital Libraries will be held on 8-9 September 2001 in Darmstadt, Germany. It will be held in conjunction with the fifth European Conference on Digital Libraries (ECDL2001) taking place in Darmstadt on 4-8 September 2001 (see announcements on page 75).

Workshop on Audio-Visual Metadata Modelling

A DELOS workshop on Audio-Visual Metadata Modelling will be held on 23 September 2001 in London, UK in conjunction with the FIAT/IFTA Annual Conference.

The workshop will be organized into two sessions. The first session will be devoted to the presentation of research and standardization results in the audio-visual metadata modelling domain. The second session will be dedicated to the presentation of concrete experiences with metadata models and editors. Each presentation will last 20 minutes plus 5 minutes for questions-answers. All the speakers (about 10-12) will be invited speakers.

The DELOS Network of Excellence on Digital Libraries, funded by the 5th Framework Programme of the European Commission, provides a framework to foster national and international research activities in the Digital Libraries domain and to disseminate the research results in the relevant user communities.

The activities of the DELOS Network of Excellence are organized into five Forums: the Research Forum, the Evaluation Forum, the Standardization Forum, the Technology Transfer Forum and the International Cooperation Forum.

The activities of DELOS Network of Excellence are open to all those working or interested in DL-related research or applications.

Delos website:
http://www.ercim.org/delos

Please contact:
Costantino Thanos — IEI-CNR
Network Director
Tel: +39 050 315 2910/2912
E-mail: thanos@iei.pi.cnr.it

Order Form

If you wish to subscribe to ERCIM News

free of charge

or if you know of a colleague who would like to receive regular copies of ERCIM News, please fill in this form and we will add you/them to the mailing list.

send, fax or email this form to:

ERCIM NEWS
Domaine de Voluceau
Rocquencourt
BP 105
F-78153 Le Chesnay Cedex
Fax: +33 1 3963 5052
E-mail: office@ercim.org

Data from this form will be held on a computer database. By giving your email address, you allow ERCIM to send you email

Name: .....

Organisation/Company: .....

Address: .....

.....

Post Code: .....

City: .....

Country .....

E-mail: .....

You can also subscribe to ERCIM News and order back copies by filling out the form at the ERCIM web site at http://www.ercim.org/publication/Ercim_News/

Considerable Increase in R&D State Funding in Hungary

Recognizing the outstanding importance of research and development, the Hungarian Government initiated a Research, Development and Innovation Programme within the frame of an economic plan called 'Széchenyi Plan'. With this programme, Hungary will strengthen state involvement in R&D, and at the same time, motivate the engagement of enterprises in closer cooperation with the public research and development sphere. National Research and Development Programmes are important parts of the programme.

In 1999 the Hungarian Government established the Science and Technology Policy Council (STPC) and its Science Advisory Board (SAB) and empowered SAB to work out the principles for the Hungarian science and technology policy, to make an assessment of the research activities in Hungary, and to define thematic priorities for research. Independent studies and discussions resulted in the document 'Science and Technology Policy 2000' outlining a long- term development programme for Hungarian science, technology and innovation. The document contains, among other things, the following main recommendations on which the Research, Development and Innovation Programme of the Széchenyi Plan is based:

- working out an attractive, performance-based faculty and researcher career model
- strengthening the existing scientific research units and increasing their absorbing capabilities for competition-based funding
- strengthening the competition-based research funding schemes as the Hungarian Scientific Research Fund (OTKA) and the Technological Development Fund
- launching the new National Research and Development Programmes.

To achieve the above goals, the central budget of 2001 and 2002 allocate substantial state funds. In 2001 the R&D's budget support will increase at least by 17.5 billion HUF (approx. 68 million Euro), and in 2002 by an additional 19 billion (approx. 73 million Euro). This will mean a considerable nominal increase of 61% in two years. The total research and development expenditure is planned to reach 1.5 percent of the GDP by 2002, a significant increase in 4 years (starting from about 0.8 percent in 1998).

National Research and Development Programmes

As a result of an approved government resolution, the Ministry of Education of the Republic of Hungary launched a call for application in the following fields:

- Improving the quality of life
- Information and communication technologies
- Environmental and material research
- Research on agribusiness and biotechnology
- Research on the national heritage and contemporary social challenge.

The final call for proposals was announced on the 15th of January, 2001. The support for these programmes will be 5.75 billion HUF (approx. 22,25 million Euro) in 2001 and 10 billion HUF (approx. 39 million Euro) in 2002. In Programmes 1-4 the minimum amount for the whole project cycle is 100 million HUF (approx. 0,39 million Euro). In the review process the project proposals submitted by consortia of institutions enjoyed preferences. The programme and their projects are application-oriented and cooperation- targeted, maintaining an excellence in original research. The decision about the successful applications has been announced recently.

Programme 2 — Information and Communication Technologies

The programme focuses on the following main fields, not excluding other areas:

- Integrated intelligent sensors
- Development of devices and methods for human language technologies
- Mobile and integrated telecommunication networks and services
- Application of combined analogue and logic computing technologies in robotics and other fields, as well as telepresense.

SZTAKI's project proposals were well-accepted. The Institute is project coordinator in five complex projects and project partner in two other ones, out of the 14 successful applications (altogether there were about 110 submitted proposals in Programme 2).

The projects coordinated by SZTAKI are:

- Commercial vehicle fleet management system (project leader: József Bokor)
- Digital enterprises, production networks (project leader: László Monostori)
- Digital restoration system of motion pictures for film archives (project leader: György Kovács)
- Knowledge-intensive information technology for the safe and optimal operation of complex industrial systems (project leader: Péter Inzelt)
- Sensing computers and telepresence (project leader: Csaba Rekeczky).

Among the many distinguished industrial project partners we find Nuclear Power Plant, Paks, Hungary, General Electric Hungary, MOL, MATÁV, KNORR BREMSE GmbH, GE -Medicor, etc. The projects will be reported on in ERCIM NEWS separately.

Link:
Ministry of Education, Hungary
http://www.om.hu/

Please contact:
Erzsébet Csuhaj-Varjú — SZTAKI
Tel: +36 1 466 5644
E-mail: csuhaj@sztaki.hu

Human Computer Interaction

by Constantine Stephanidis

Human-Computer Interaction (HCI) is concerned with the design, implementation and evaluation of interactive computer-based systems, as well as with the multi-disciplinary study of various issues affecting this interaction. The aim of HCI is to ensure the safety, utility, effectiveness, efficiency, accessibility and usability of such systems.

In recent years, HCI has attracted considerable attention by the academic and research communities, as well as by the Information Society Technologies industry. The on-going paradigm shift towards a knowledge-intensive Information Society has brought about radical changes in the way people work and interact with each other and with information. Computer-mediated human activities undergo fundamental changes and new ones appear continuously, as new, intelligent, distributed, and highly interactive technological environments emerge, making available concurrent access to heterogeneous information sources and interpersonal communication. The progressive fusion of existing and emerging technologies is transforming the computer from a specialist's device into an information appliance. This dynamic evolution is characterised by several dimensions of diversity that are intrinsic to the Information Society. These become evident when considering the broad range of user characteristics, the changing nature of human activities, the variety of contexts of use, the increasing availability and diversification of information, knowledge sources and services, the proliferation of diverse technological platforms, etc.

HCI plays a critical role in the context of the emerging Information Society, as citizens experience technology through their contact with the user interfaces of interactive products, applications and services. Therefore, it is important to ensure that user interfaces provide access and quality in use to all potential users, in all possible of contexts of use, and through a variety of technological platforms. The field of HCI is now experiencing new challenges. New perspectives, trends and insights enrich the design, implementation and evaluation of interactive software, necessitating new multidisciplinary and collaborative efforts.

Europe is playing a protagonist role in this direction, by conducting collaboratively innovative scientific and technological work that has led to the establishment of a framework of reference for a user-friendly Information Society. The Information Society Technologies (IST) Programme of the European Commission has placed considerable emphasis on the human- and social-oriented aspects of technology. This commitment is expected to be continued and enhanced in the forthcoming Sixth Framework Programme for the period 2002-2006, currently being elaborated on the basis of the conclusions of the Lisbon European Council and the objectives of the eEurope initiative.

On the European scene, ERCIM appears as a key actor in the field of HCI, with many collaborative activities being conducted in several member institutions. The ERCIM Working Group 'User Interfaces for All' (http://ui4all.ics.forth.gr/)

The ERCIM Working Group 'User Interfaces for All'

Since its establishment in 1995, the ERCIM Working Group 'User Interfaces for All' (UI4ALL) has organised annual workshops, that have brought together researchers and teams working in the different ERCIM organisations (but also organisations beyond ERCIM or the European boundaries), who share common interests and aspirations, and contribute to the endeavours towards making the emerging Information Society equally accessible to all. The prolific and dedicated work of this growing community has led to the establishment of the new International Conference on Universal Access in Human-Computer Interaction (UAHCI). The first such event, UAHCI 2001 (http://uahci.ics. forth.gr/), will take place in New Orleans, Louisiana, USA, 5-10 August 2001, in cooperation with HCI International 2001, with 228 papers accepted for presentation covering a wide variety of topics in this area. The seventh ERCIM UI4ALL Workshop will be organised in October 2002, and henceforth workshops will take place every second year, alternating with the UAHCI Conference.

Furthermore, the work of the ERCIM Working Group UI4ALL has contributed to the establishment of the International Scientific Forum 'Towards an Information Society for All' ISF-IS4ALL (1997-2000), an international ad hoc group of experts which first recognised the need for a global approach towards an Information Society for All (http://ui4all.ics.forth.gr/isf_is4all). Two White Papers have been published and submitted to the European Commission, reporting on an evolving international R&D agenda in the field of HCI. Since then, the vision of an Information Society for All and the necessity for universal access to Information Society Technologies have acquired widespread acceptance and importance not only at a scientific and technological, but also at a European policy level, as demonstrated by the eEurope - eAccessibility initiative of the European Commission (http://europa.eu.int/information_society/eeurope/index_en.htm).

The activities initiated by the International Scientific Forum are now being continued in the framework of the Thematic Network (Working Group) 'Information Society for All' IS4ALL (http://is4all.ics.forth.gr). For more details see article on page 11.

The work of the ERCIM UI4ALL Working Group, recipient of the ERCIM WG Award for the year 2000, has also contributed to the establishment of an international, interdisciplinary refereed journal 'Universal Access in the Information Society' (UAIS), published by Springer (http://link.springer.de/journals/uais/) that solicits original research contributions addressing the accessibility, usability, and, ultimately, acceptability of Information Society Technologies by anyone, anywhere, at anytime, and through any media and device. 'User Interfaces for All' is also the title of an edited book dedicated to the issues of Universal Design and Universal Access in HCI, published by Lawrence Erlbaum Associates (ISBN 0-8058-2967-9, 760 pages).

has systematically promoted the proactive realisation of the Design for All principles in HCI. The vision of User Interfaces for All involves the development of user interfaces to interactive applications and telematic services, which provide universal access and usability to all potential users, including people with different cultural, educational, training and employment background, novice and experienced computer users, the very young and the elderly, and people with different types of disabilities, interacting with different technological platforms in different contexts of use.

The special theme of this ERCIM News reflects the increasing importance of the HCI field in the wider context of research and development activities for Information Society Technologies, and sheds light on the multidisciplinary and evolving nature of work in this field, as well as on current trends and initiatives. The call for contributions has met overwhelming interest, and an unprecedented number of contributions are included, reflecting a wide variety of issues, such as novel interface development approaches, interaction techniques and devices for multimodal, continuous and ubiquitous computing, access to information, and applications in different domains.

Please contact:
Constantine Stephanidis — ICS-FORTH
Tel: +30 81 391741
E-mail: cs@ics.forth.gr

User-Centred Methods for New Application Domains, New Users, New Interaction Media

by Dominique Scapin and Noëlle Carbonell

In order to face the challenges of a widespread usage of information technology, ergonomics and HCI research is needed. This is the scope of the MErLIn research group at INRIA.

Thanks to recent research advances, most hardware and software obstacles to the advent of a world wide information society are being overcome. However, unless universal software accessibility is achieved, this major social evolution may increase social exclusion. Appropriate ergonomic design guidelines, recommendations and standards are needed, particularly for meeting the requirements and expectations of users:

- in all standard user categories (eg, novices, occasional and expert users)
- in all contexts of use, especially new ones, such as home automation, online services for the general public (eg, e-commerce), mobile computing, etc
- providing users with special needs (eg, disabled and elderly citizens) with easy computer access.

These issues are at the centre of the research activities of the MErLIn project (INRIA, LORIA, University Paris V). MErLIn's research is focused on the study of human-computer interaction, using a pluridisciplinary user-centred approach, and capitalising on the methods and techniques developed in cognitive ergonomics, artificial intelligence and user interface engineering. The main objective is to contribute to the optimisation of the ergonomic quality of interactive software.

On the one hand, the project aims at integrating ergonomic results into the software design process by designing user centred methods and support software for the ergonomic evaluation and design of interactive software. On the other hand, it aims at acquiring the ergonomic knowledge required for improving the utility and usability of novel computer applications, that is software relating to new application domains, promoting new contexts of use, and/or intended for new categories of users.

In the area of user centred design methods and support, research topics include:

- Adaptation of standard ergonomic criteria and definition of new criteria for Web applications; design of support tools for the assessment of the ergonomic quality of Web applications. Current studies concern both task performance and user preferences (EvalWeb and ErgoCOIN projects, respectively with LIHS and University of Santa Catarina; ESPRIT project 'UsabilityNet'). In addition, 3D ergonomic guidelines are considered in the Eureka project COMEDIA.
- Integration of task models into the design process; in particular, modelling and predictive evaluation of cognitive workload (PREDIT project with SNCF and RATP, supported by the French Ministry of Transportation).
- Definition, implementation and ergonomic assessment of new contextual online help strategies for users of current software intended for the general public; the experimental evaluation of the prototype involves potential novice users (research project with IMASSA, supported by the French Ministry of Defence and the CNRS).

In the area of new multimodal interaction languages, ie, for new application domains, new users, and new interaction media, the research aim is to identify and tackle new ergonomic issues stemming from technological innovation, using specific ergonomic methods and techniques, such as observation, experimentation, user testing, evaluation of mock-ups and prototypes. Ongoing research projects include:

- The definition and evaluation of multimodal command languages which could prove adequate substitutes for direct manipulation in interaction contexts where the use of mouse and keyboard is awkward or impossible; the experimental evaluation of the defined languages, which combine speech and designation gestures on a touch screen, is focused both on their naturalness, usability, and their efficiency, that is the simplicity and reliability of their automatic interpretation (recent research project with IMASSA, supported by the French Ministry of Defence and the CNRS).
- The design, based on the expertise gained thanks to the previous project, of efficient and usable multimodal languages for interacting with virtual realities, especially for manipulating virtual 3D objects in the context of computer aided design applications, eg, for the design of ready-to-wear or the customisation of products on e-commerce Web sites (Eureka project COMEDIA).
- The design of adaptable user interfaces for ill-sighted and blind users (research project supported by the French MIRE and CNAV organisms).
- The study and participatory design of augmented reality laboratory notebooks (with Institut Pasteur), as well as the study of distant communities (ESPRIT project 'InterLiving', with LRI and University of Aarhus).
- The study of ergonomic issues related to the implementation of the universal accessibility principle and its application to the health sector (IST Thematic Network IS4ALL).

Links:
http://www.inria.fr
http://www.loria.fr

Please contact:
Dominique Scapin — INRIA
Tel: +33 1 39 63 55 07
E-mail: Dominique.Scapin@inria.fr

Noëlle Carbonell — LORIA
Tel: +33 3 83 59 20 32
E-mail: Noelle.Carbonell@loria.fr

IS4ALL: An Analytical Insight into Universal Access

by Constantine Stephanidis and Demosthenes Akoumianakis

IS4ALL — Information Society for All — is a European project that provides a methodological frame of reference which introduces universal access principles into the design of Health Telematics applications and services. The characteristics of Health Telematics, such as the variety of end users involved, the changing Healthcare contexts of use and the penetration of new computer-mediated activities, make it a rich field of inquiry and re-shape the way in which healthcare practices are structured and organized.

IS4ALL (IST-1999-14101) seeks to establish a wide, interdisciplinary and closely collaborating 'network of experts' (Working Group) to provide the European Health Telematics industry with a comprehensive information package detailing how to appropriate the benefits of Universal Design. While the specific technological / scientific objectives to be attained by IS4ALL have been described in ERCIM News 43, this article will elaborate on recent technical developments concerning the advancement of methodological support for introducing universal access principles into the design of Health Telematics applications and services. The project departs from the traditional perspectives of universal access as a set of high-level principles or guidelines, which can be interpreted to provide the required design insights, and instead, it emphasizes an analytical and process-oriented approach, which blends scenario-based perspectives on systems development and structured argumentation to re-shape and filter scenarios.

Scenarios are perceived as narrative descriptions of computer-mediated human activities in a Health Telematics environment. The social setting of a Health Telematics environment may be bound to a clinic within the hospital, a ward within a clinic or even to the end users business or residential environment. The scope of such scenarios is intended to be narrow and focused on a very specific issue. As an example, a possible scenario would be focused on patient's access to the electronic healthcare record from home, or while on the move, using a portable device, which runs on a particular platform.

Generating the Scenario

Scenarios can be developed at different levels. They can range from vision-oriented statements of intended actions to concrete experiences with an implemented artefact. Since scenarios may or may not have an explicit focus on universal access, a key question is to devise suitable mechanism or techniques for working with these scenarios to derive useful insights for the purposes of the project (see figure).

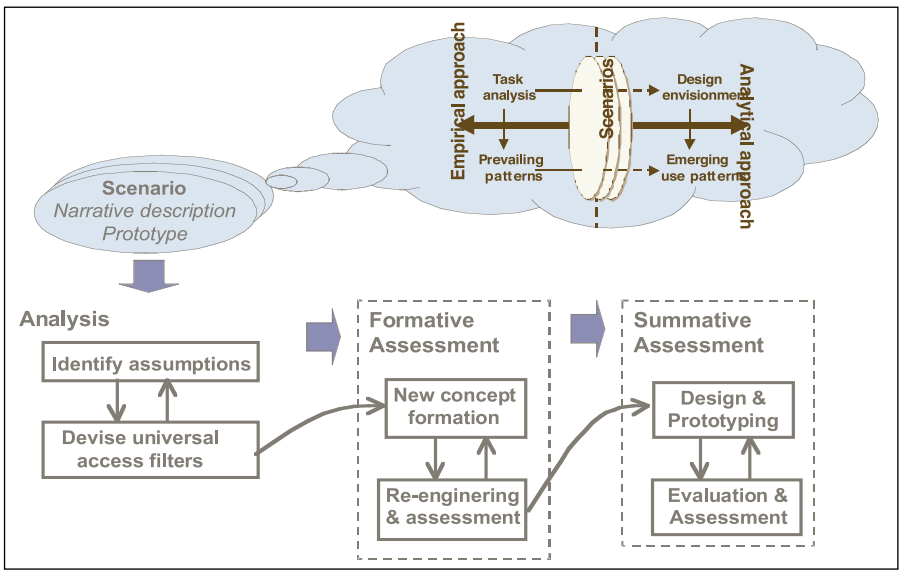

Working with scenarios — steps and phases in IS4ALL.

Working with the Scenario

Working with the scenario entails a process of structured argumentation, which results in unfolding, recording and reflecting upon implicit/explicit claims embodied in the original scenario. Since we are concerned with designing for diversity, an appropriate argumentation framework should accommodate arguments in support or against universal access. To this effect, rather than starting from high-level principles and guidelines, we are in favour of a more direct and structured approach leading to a structured argumentation technique based on the notion of accessibility filters. Accessibility filters may take several forms. The most popular is that of questions, which seek to provide insights into how a particular product is to be used (by specific user groups, novel contexts of use, or different platforms).

Scenario Screening: Setting Accessibility Filters and Relaxing some of the Implicit Assumptions

Accessibility filters are specific issues depicting emerging patterns of use necessitated either by an intention to accommodate diversity in the design of the task or by the need to re-engineer an artefact. For instance plausible accessibility filters could be the following:

- How can a user with gross-temporal control familiar with switch-based interaction perform the task?
- How can the task be carried out with an alternate pointing device (eg, a stylus of a palmtop computer)?
- How can the task be performed in a public kiosk?

Each one of these filters may be used to re-engineer the original scenario (or the

artefacts implied) so as to provide instances, which satisfy the filter.

Identifying Breakdowns and Envisioning New Design Alternatives
This involves a conscious effort to encounter potential problems, envision solutions and unfold alternatives. Thus for instance, developing further the first accessibility filter may lead to changes in the interaction with the artefact or perhaps to the introduction of new artefacts, such as visual keyboards for text entries or alternative window management facilities.

What is important to note about accessibility filters is that they play a dual role in design. Firstly, they are useful when there is a tentative design to be assessed. In such cases, they help designers identify potential break down or shortcomings in the use of the product. Moreover, they can also be used as a basis for argumentation during design, thus fostering an analytical insight to accessible design. Nevertheless, accessibility filters do not provide answers. Rather, they motivate designers to think about certain aspects of a problem.

Acknowledgements
IS4ALL comprises one co-ordinating partner, namely ICS-FORTH, and several member organisations, including Microsoft Healthcare Users Group Europe (MS-HUGe), European Health Telematics Association (EHTEL), CNR-IROE, GMD, INRIA and Fraunhofer-Gesellschaft zur Foerderung der angewandten Forschung e.V.—Institut fur Arbeitswirtschaft und Organisation (FhG-IAO).

Please contact:
Constantine Stephanidis — ICS-FORTH
Tel: +30 81 391741
E-mail: cs@ics.forth.gr

Supporting Collective Design Activities: A Cognitive Task Analysis Method

by Françoise Darses, Françoise Détienne, Pierre Falzon and Willemien Visser

The EIFFEL research group at INRIA carries out research activities in the domains of Cognitive Ergonomics and Cognitive Psychology. Their aim is to model both individual and collective design activities and to assess and specify tools and methodologies supporting design.

Studies on reasoning in design have usually been carried out on individual problem solving activities. In response to the increasing need to assist collective work in an industrial context, more recent studies have shifted their focus. A major concern in industrial modernisation is the creation of new organisations which support collective work, greater interaction between designers and manufacturers, as well as capitalisation and reuse of design knowledge.

Our working hypotheses are:
- Design problem is a particular kind of problem, often qualified as ill-defined
- There are cognitive invariants in the design activity whatever the application domain: eg reuse of past designs
- Research work on transversal themes such as collective work or organisational memory should advance significantly by focusing on particular task domains, design tasks being one of them
- The nature of the approaches for supporting design should be not only technical. They are most often techno-organisational.

Our methodological approach is as following. We conduct empirical studies, either field studies or laboratory experiments. Our empirical studies examine two kinds of design situations:
- face-to-face design situations (more generally, non-technology-mediated design situations)
- technology-mediated design situations.

The former studies concern upstream research on particular activities which need to be supported. They allow us to construct a reference model of the activity. Based on this kind of model, special support needs are identified. The latter studies allow us to assess different kinds of technological systems or methodologies.

COMET: A Cognitive Task Analysis Method for Collective Design
One example of the research contributions of the Eiffel group is a Cognitive Task Analysis (CTA) method for analysing and modelling collective design: COMET. The ensuing objective is to build cognitive models useful for intelligent assistance systems for designers.

The design of products is mainly achieved by teams, in which designers must cooperate to develop the solution. As any other professionals, designers need assistance tools, which can be either computational (CAD or CAM for instance) or methodological (project management tools, decision making methods, etc.). To be efficient, these tools must be able to support the cognitive processes that underlie the designers' activities. Thus, collective design cognitive processes must be identified and modelled. But it appears that more often than not, the cognitive dimensions of designers' work are directly extrapolated from observations of their activity or from interviews, without going through a phase of cognitive modelling. This approximate nature of the analyses makes cognitive system design more like the result of inspired innovation rather than the result of well-grounded methods. CTA methods are thus needed to understand design activities and to develop tools for designers.

COMET is based on the analysis of the dialogues exchanged between designers in co-design situations. COMET gives some methodological guidelines related to a cognitive ergonomics approach of the activity. These guidelines have been elaborated and used in two case studies of co-design situations: computer network co-design and software review meetings. COMET has also been applied and assessed in other design studies.

COMET takes place when the socio-technical macro-analysis of the activity has finished: the ergonomic problem to solve has been formulated, the concerned people have been identified, the prescribed task has been understood through interviews, the organisational demands have been highlighted. COMET stems from observation of work in situ, and especially from observations of design meetings. It provides a frame for analysing design dialogues recorded in such situations. The frame is made up of two distinct and hierarchical levels of coding: basic level (where individual UNITS are identified) and composite level (where the 'Co-operation Moves' are identified).

At the basic level, utterance turns are cut up into one or more individual UNITS according to a coding scheme developed on a Predicate — Argument basis. Predicates correspond to actions (ACT) implemented by participants; arguments (OBJ) correspond to objects related to the action. Each unit is modulated (MOD), according to the form of the predicate (Assertion or Request). Thus, each UNIT is coded as MOD[ACT/OBJ].

At the composite level, frequent and consistent UNITS are grouped into sequences. Such sequences can be formed in various ways, on a qualitative or on a quantitative basis, ie on a concept- or a data-driven basis. Sequences correspond to basic co-operative interactions called 'Co-operation Moves'. COMET was developed for analysing co-design processes occurring in the same place, where the designers are present in the same room. It can also be applied for design work mediated by communication tools (such as web-based tools), where the designers do not design at the same time. The growing importance of such cooperative design situations makes COMET a good candidate as a CTA method usable for such cooperative design environments.

Link:
http://www.inria.fr/Equipes/EIFFEL-eng.htm

Please contact:
Francoise.Detienne — INRIA
Tel: +33 1 3963 5522
E-mail: Francoise.Detienne@inria.fr

Towards a Generic Framework for Building Intelligent Collaborative Interface Agents

by George Vouros and Vangelis Kourakos-Mavromichalis

The project Intelligent Collaborative Interface Agents (ICIA) aims at the implementation of collaborative interface agents that act in conjunction with their human partners, increasing their competency in using a software application.

The ICIA project has been initiated by the Intelligent Collaborative Systems (InCoSys) Group at the Department of Information and Communication Systems of the University of the Aegean, Samos, Greece. We are currently developing a collaborative anthropomorphic interface agent that acts as an active partner for children from 10 to 14 years, in the context of an educational software application.

Although user interfaces designed with well-known interface design strategies, such as the information visualization strategy [Ahlberg C., Shneiderman B.: 'Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays'. In Human Factors in Computing Systems: Proceedings of the CHI '94 Conference. New York: ACM, 1994], provide facilities for the users to learn progressively how to perform on them—by experimenting and observing the effects of their actions—it is the case that novice users, and particularly small children, need the help of an expert user in order to learn basic functionalities. In this case, users need the help of an instructor that would teach them to perform the needed tasks. To instruct users effectively, an instructor needs to understand users' goals, ascribe to users mental states describing their knowledge, beliefs, capabilities, intentions and desires, and commit that he helpful behaviour will be exhibited towards achieving users' goals when the need arises. Understanding user goals and committing to exhibiting helpful behaviour towards achieving these goals assures that the instructor shall not perform actions that are in conflict with users' intended actions and desires.

In addition to interactive interfaces that provide facilities for users to learn progressively how to perform on them, this research proposes the use of intelligent collaborative agents. These agents aim to provide help to users taking into account their goals, intentions, information needs, preferences and characteristics. Such agents must be designed and implemented as active partners that look over the shoulder of users and collaborate with them towards achieving a common goal.

An approach to intelligent user interfaces that is close to the ICIA approach is that of COLLAGEN [Rich C., Sidner C., Lesh N.: 'COLLAGEN: Applying Collaborative

Discourse Theory to Human-Computer Interaction'. In AI Magazine, Special Issue on Intelligent User Interfaces, 2001 (to appear)]. COLLAGEN is complementary to the interface agents proposed in this paper since it can play the role of the middleware between an application, the user and a software agent that is actually collaborating with the user. Furthermore, COLLAGEN implies that the agent will always perform the next action that contributes to a goal, which is not appropriate for instructing the user on how to perform a task.

Key issues towards our goal are the following:

- The agent must form advices and information presentations, and act on behalf of the user when this is needed.
- The agent must be collaborative, ie, understand user desires, form intentions that the user shall achieve his goals and, where appropriate, exhibit helpful behaviour.

A major related issue is the ability of the interface agent to balance and intermix between reactive and social deliberative behaviour. In the first case, the agent takes the initiative to act by itself and presents information or performs on behalf of the user by reacting to user actions and to events happening in the interface. Deliberating socially, the agent collaborates with the user towards achieving a shared goal.

To implement a collaborative agent so as to address the above-mentioned key issues, we have adopted the SharedPlans model for collaborative planning activity proposed in [Grosz, B.J., Kraus, S.: 'Collaborative plans for complex group action'. In: Artificial Intelligence 86, pp. 269-357, 1996]. SharedPlans is a general model that provides the basis for the design of collaborative agents. It adopts a mental state view of plans and enables collaborating agents to treat their mental states (ie, their beliefs, desires and intentions) in an integrated way.

To address the key issues mentioned above and comply with the generic principles and the stages of the collaborative activity [Wooldridge M., Jennings N.: 'The Cooperative Problem Solving Process'. In Journal of Logic and Computation, 9(4):563-592, 1999] we have used the ICAGENT agent development framework. This is a generic framework for the implementation of intelligent agents that act in dynamic and unpredictable environments [Kourakos-Mavromichalis V., Vouros G.: 'Balancing Between Reactivity and Deliberation in the ICAGENT Framework', In: Balancing Reactivity and Social Deliberation in Multi-agent Systems, LNCS, Springer-Verlag, 2001 (to appear)]. The ICAGENT framework provides advanced plan management abilities, has been developed based on the SharedPlans model for collaborative activity and provides the basis for the implementation of collaborative agents.

The collaborative interface agent recognizes the situations where users need help. This is currently done with specialized situation recognition rules.

The agent possesses knowledge on interface objects and domain tasks, ascribes mental attitudes to users, forms intentions to act and intentions to exhibit helpful behaviour, and utilizes instructional and presentation strategies for instructing users on how to achieve their goals.

Future activities

Future activities concern:

- further implementation of advanced collaborative features according to the SharedPlans model of collaborative activity, emphasizing on the efficiency of the implementation
- the implementation of advanced plan recognition techniques for the interpretation of agents' actions and for understanding agents' intentions during action
- evaluation of the effectiveness of collaborative interface agents on users competency in using a software system.

Please contact:
George Vouros — University of the Aegean, Greece
Tel: +30 273 82226
E-mail: georgev@aegean.gr

Interface Development Toolkits for Non-visual and Switch-based Interaction

by Constantine Stephanidis and Anthony Savidis

Commercial user interface development toolkits do not provide support for interaction techniques alternative to mouse/keyboard based graphical interaction. Two toolkits, the HAWK toolkit for non-visual interaction and the SCANLIB toolkit for switch-based interaction have been developed at ICS-FORTH.

The HAWK toolkit provides a set of standard non-visual interaction objects and interaction techniques that have been specifically designed to support high quality non-visual interaction. HAWK is appropriate for developing not only interfaces targeted to blind users, but also interfaces for a variety of situations in which dialogues not relying on visual communication and standard input devices are required (eg, driving, telephone-based applications, home control auditory interaction). A key notion in HAWK is that of a container interaction object. In the HAWK toolkit there is a single generic container class that does not provide any pre-designed interaction metaphor, but supplies appropriate presentation attributes through which alternative representations can be created. The container class has four attributes that enables each distinct container instance to be given a metaphoric substance by appropriately combining messages and

sound-feedback (both names and sound effects can have a metaphoric meaning). In addition to generic containers, HAWK provides a comprehensive collection of conventional interaction objects directly supporting non-visual dialogue, namely menus (exclusive choice selector object), lists (multiple choice selector object), buttons (push button analogy for direct command execution), toggles (radio button analogy for on/off state control, edit fields (single line text input) and text reviewers (multi-line read-only text editor, with mark-up facilities).

HAWK also supports a variety of interaction techniques, namely synthesized speech, Braille (2-cell transitory Braille, 40-cell Braille), and digitised audio for output, and standard keyboard, joystick used for gestures independently of visual interaction, touch-tablet (for programmable commands via associated regions) and voice recognition for input. The HAWK toolkit provides all the programming features met in currently available toolkits, such as hierarchical object composition, dynamic instantiation, call back registration, and event handling. The navigation dialogue enables the blind user to move within the interface structure composed of organizations of containers and contained objects in an easy way, through the provision of multi-modal control facilities. For instance, visiting contained objects is possible through joystick-based gestures, voice commands, keyboard short cuts, or via pressing specific regions of the touch-tablet. Container objects may contain other container objects realizing different metaphoric representations, thus supporting fusion of different metaphors in the context of non-visual interactive applications.

The HAWK toolkit has been used in the development of the AVANTI browser (see ERCIM News no. 41), of the NAUTILUS information kiosk (Project Nautilus, funded by the Hellenic Ministry of Development) and of a non-visual digital library for the Hellenic Blind Association.

Switch-based Interaction

In the SCANLIB interface development toolkit, the basic Windows object library has been augmented with scanning interaction techniques. Interfaces implemented through SCANLIB directly support motor-impaired user access, as well as access in other situations in which the keyboard and mouse input devices can not be used. Apart from enabling intra-application interaction control (eg, having access via switches to all interface elements of any interactive application), SCANLIB also supports inter-application interaction control (eg, enabling users to move across different applications). In SCANLIB, basic object classes are classified into five categories, each requiring a different dialogue policy to be designed:

- Top-level windows: these are the objects directly providing window management operations. Since in Windows it is impossible to 'access' (in a programmable manner) the built-in controls (icons) for window management, all top-level windows in SCANLIB have been augmented with an additional toolbar accessible via scanning.
- Container objects: these are object classes with instances present as intermediate nodes in object hierarchies, able to encompass an arbitrary number of contained object instances. Container objects enable sequential scanning of all contained objects.
- Text-entry objects: objects requiring text to be supplied impose the need for keyboard emulation. When the user focuses (via scanning) on a particular text-field object, a special on-screen keyboard automatically appears, through which text input is enabled. Such an augmented dialogue is realized transparently to programmers, for which the augmented text-field object at the programming level appears as a conventional Windows text-field object.
- Composite objects: typical examples of composite objects are the scroll-bar class (composed of two arrow buttons and a slider) and the combo-box class (composed of a drop-down button, a label, and a menu object). Composite objects enable sequential scanning of all component-objects.
- Button categories: these are simple interaction objects supporting direct user actions for executing associated operations or changing state of Boolean parameters. Typical examples from this category are: push buttons, check boxes and radio buttons.

The scanning interaction techniques are based on two fundamental actions: SELECT, and NEXT. Depending on the type of switch equipment required, four scanning modes are supported:

- 1-switch (for SELECT action), time scanning (automatic NEXT action, after a specific time interval)
- 2-switches (one for SELECT action, one for changing scanning direction), time scanning (as before for automatic NEXT action)
- 2-switches (SELECT action, NEXT action), no time scanning
- 5-switches (SELECT action, NEXT action, change direction, exit directly container, re-start with container).

Programming control is provided for the extra attributes introduced by the augmented version of Windows objects, thus enabling objects to differentiate scanning style depending on scanning mode (one of the five alternatives), time interval (when time scanning is supported), and highlighter presentation parameters.

The SCANLIB toolkit has been used in the development of the AVANTI browser (see ERCIM News no. 41), and of the GRAFIS word processor for people with motor impairments (see ERCIM News no. 38).

Please contact:
Constantine Stephanidis — ICS-FORTH
Tel: +30 81 391741
E-mail: cs@ics.forth.gr

Specifying User's Knowledge: A Cognitive Ergonomic View on IT Design

by Gerrit van der Veer

There are strikingly different views on what should be the focal points in designing technology for human users: graphical design, interaction design, engineering of state-of-the-art requirements, or a cognitive psychological approach. At the Vrije Universiteit Amsterdam scientists developed an eclectic approach towards the design of interactive systems named DUTCH (Designing for Users and Tasks, from Concepts to Handles).

The core of this approach is to define all relevant knowledge a user should possess to use the new system. Such a full specification of all aspects of the technology to be (re)designed, from the point of view of the prospective user, is called UVM (user's virtual machine).

In DUTCH we consider interactive systems where various types of users may have different roles, each of which needs different knowledge to use the system. Specifying a complete interactive system will amount to specifying various more or less overlapping UVMs, one for each role, based on the user's viewpoint. Adaptability is clearly separated from specifying the knowledge needed to perform role specific tasks. Whereas for the user the system appears monolithic, the designer has to find out what types of knowledge will have to be specified. In the first place the system should speak a language that is understandable for the user in relation to his intentions and tasks. In interactive systems, however, much more complex knowledge is needed. Users need to be able to instantiate a mental model of what is going on in interaction, based on knowledge available in human memory and information provided by the system at the relevant moment. This knowledge and information must be specified in an early phase of the design. We should systematically take care of the different types of knowledge that will be engineered at a later stage (probably by very different specialists who may never meet each other, like software engineers, hardware designers, graphical designers, workflow managers, and systems architects).

What is going on behind the Screen

Users hardly care about 'what's inside'. Only if certain internal processes take more than a few seconds, or are unpredictable for the user, the system should make the user aware of exactly this, nothing else. On the other hand, certain things going on behind the screen are needed to understand the interaction and perform the intended tasks successfully (eg, "is this object still on my clipboard?"; "is the original file untouched after I performed a 'save as' operation?").

What is going on behind this Computer

Current interactive systems are seldom stand-alone. Even if we connect the machine to the outside world solely by exchanging diskettes, users need to know the reason for certain unexpected events, and take measures for safe, reliable, and efficient interaction with external agents. In the case of more sophisticated connectivity, users need to understand aspects of local and world wide connections and safety procedures (fire walls, signatures, authentication). In certain types of tasks, the actual location of certain processes should be known and understood ("is my pin code transported through the network to the bank, or is it validated locally at the terminal?").

What Organizational Structure is behind my Computer

In the envisioned interactive systems, a user needs to understand that other users are players in the same game, influencing the total process, for example by managing their workflow process or monitoring their actions, or by having different rights ("who may inspect or change or delete the e-mail I have just sent?").

What Task Domain is available through my Computer

If a fashion designer uses an interactive system, several tasks could be delegated to the computer with considerable added value, for example to search multimedia databases of video clips from cat walk shows, images of the people and the location the clothes are intended to serve, interviews with potential clients, and works of art that are characteristic of the cultural ambiance of the intended buyers. So far, possibly because the right media are not yet available, other parts of the task domain are not yet supported successfully, for example assessing the texture quality of fabric, or impressionistic sketching of the process of moving clothes at a stage before the design of shape has developed to enough detail for a mannequin to be actually dressed.

What are the Process and Time Aspects of delegating Tasks to my Computer

IT time scales are completely different from those of human information processing. Because of this, systems often react to user requests either amazingly fast or annoyingly slow and unpredictable. A more fundamental problem occurs if the computer is a front-end to a process where time characteristics are of a non-human scale. Here either the feedback is slow (eg, monitoring chemical plant processes) but the user is supposed to anticipate the outcome of previous actions in order to optimize the continuation of the process, or the process is fast in performing irreversible tasks (eg, delegating certain flight operations to the auto pilot) and, consequently, the user is expected to assess the expected consequences in all details that affect safety and reliability. Actually, both cases are requesting the same type of insight and knowledge from the user. Various recorded cases where pilots ended up 'fighting' the auto pilot are illustrations of designs where this type of knowledge was never considered. The same can be found in recorded disasters with nuclear power plants where potential consequences of task delegation were never modelled, and failure indications were completely inadequate and even disturbed human problem solving processes.

Specifying the above-mentioned types of user knowledge requires a modelling language that is precise and unambiguous enough to be understood by engineers without deep insight in cognitive science. At the same time this language should enforce complete specification of what matters to the user and what is relevant from the point of view of human information processing. We developed NUAN (new user action notation, adapted from Hix & Hartson) into precisely the modelling language and supporting tool that is needed for this.

Specifying the details of technology, from the point of view of DUTCH design, is equivalent to specifying all relevant knowledge needed for each type of user. Designing technology is thus translated into building prescriptive knowledge models, intended for the implementing engineer. Exactly this knowledge should also be available to be instantiated into a mental model whenever interaction requires it. Of course, in prescribing this knowledge, the cognitive ergonomic insights in human information processing are of paramount importance. The system should be understandable, usable, and require the lowest level of effort on the part of the user.

Link:
http://www.cs.vu.nl/~gerrit/

Please contact:
Gerrit van der Veer — Vrije Universiteit
Amsterdam, The Netherlands
Tel: +31 20 444 7764
E-mail: gerrit@acm.org

Human-Computer Interaction Challenges in Nomadic Computing

by Reinhard Oppermann

In Nomadic Computing, mobile users are supported by contextualised information presentation and interaction. At the GMD Institute for Applied Information Technology, prototypes and services are currently being developed in the framework of two projects: 'Crumpet', a European project with five partners, focuses on localisation of the user and personalisation of information; 'SAiMotion', a co-operation between GMD and Fraunhofer Gesellschaft, concentrates on context modelling and Human-Computer Interaction.

Human-Computer Interaction is one of the key issues for information systems in general, and for nomadic information systems in particular. Since information technology is used as a means for accomplishing tasks, when information processing is embedded into a more general frame of tasks, and computing technology is embedded into other equipment and technical infrastructure, Human-Computer Interaction is of central value for effectiveness, efficiency and user satisfaction. This is particularly true as mobile information and communication devices become smaller and more restricted with respect to information presentation, data entry and dialogue control. Mobile information and communication devices, like PDAs or cellular phones, are currently designed for dedicated tasks, but their application is becoming more and more integrated into a task flow where stationary and mobile devices are used together.

Nomadic computing starts with this state of technology, in which the activities of a user can be supported during the whole process of task accomplishment independently of specific devices, and access can be provided to public and personal information spaces. Nomadic Computing focuses on activities distributed in space (people roaming around), time (people with several overlapping tasks) and social communities (involving several people). Nomadic information systems allow the user to prepare an activity from home using the desktop, by browsing through and editing notes and messages in the public (the Web), corporate (business data) and personal information spaces (address book and personal papers). On tour, the ongoing activity is supported by a mobile device, and back at home or in the office the preferred device will again be the desktop computer. Continuous access to information spaces is provided via the Internet, wired or wirelessly.

The Human-Computer Interaction challenge comprises two classes of requirements. Firstly, interaction must be consistent from one device to another, using non-contradictory metaphors and methods. Secondly, interaction has to be appropriate for the particular device and environment in which the system is used. Note that the requirement does not call for identical metaphors and methods. The desktop computer allows for different interaction techniques than a palmtop computer or a WAP phone. Using the keyboard and a mouse may be the method of choice for the desktop, while using the pen and the microphone may be appropriate for the palmtop. The predominance of visual display and mechanical input with the desktop will be replaced by the audio display and speech control with the palmtop.

Crumpet and SAiMotion

In the projects Crumpet and SAiMotion we develop scenarios and prototypes for indirect interaction by exploiting the user's navigation in the physical space and providing easy and intuitive interaction methods. By roaming in the physical space the user indicates environments of interest and objects of relevance for current tasks. Tourist and mobile business applications are selected to demonstrate contextualised content selection and new interaction paradigms. Users are supported by logistic services in finding transport, shopping or accommodation services, partners or places for business, and scalable explanations of objects related to the interests and previous knowledge of the user. The implementation will be based on a standards-compliant framework, extended to support nomadic applications, devices, and networks for fast creation of robust, scalable, seamlessly accessible nomadic services. In particular the following services are foreseen:

- **Adaptivity:** Adaptive nomadic services responding to underlying dynamic characteristics, such as network quality of service and physical location
- **Platform independence:** Suitability for a wide range of terminal types, including next generation mobile phone/PDA/PC hybrid terminals
- **Navigation:** The support of users in every phase of mobile problem situations via dynamic presentation of relevant routing and navigation information will provide an effective navigation in space not depending on external media as maps, guidelines or verbal enquiries
- **Action assistance:** Hints regarding necessary actions to be taken to successfully achieve the defined goal, as well as relevant navigation and routing information are provided by the system
- **Supply of resources:** Automatic identification and supply of helpful and necessary resources, in the vicinity of the user, which are needed to complete the current task, eg indoor network resources (printer, photocopier) as well as outdoor resources (ticket counter, bank machine, bus stop)
- **Conflict monitoring:** A permanent monitoring of the action sequences of the user allows the recognition and prevention of possible conflicts
- **Context sensitive interaction:** Multimodal interaction and presentation mechanisms that adapt to environmental variables like daytime, noise, and situation of communication allow a context- and situation-adapted interaction including audio display and speech control
- **Anticipated information supply:** Proactive presentation of relevant information adapted to the present task sequence, interest and situation of the user and to the frame activity improves the efficacy of interaction.

Other research and development activities are planned to find and disseminate solutions for new types of devices, infrastructure and services which will increase the benefit for a wide range of users. Several partners are involved in these activities, including FhG Fraunhofer-Institut für Graphische Datenverarbeitung IGD), Fraunhofer-Institut für Arbeitswirtschaft und Organisation (IAO), Fraunhofer Institut für Zuverlässigkeit und Mikrointegration (IZM), Fraunhofer Institut für Integrierte Schaltungen (IIS); GMD Institut für Publikations- und Informationssysteme (IPSI); Emorphia (UK), University of Helsinki, Sonera Ltd (Fi), Queen Mary & Westfield College London, European Media Laboratory GmbH (D), Portugal Telecom Inovação, S.A.

Current pen based interaction.

Links:
http://fit.gmd.de/topics/icon/
http://fit.gmd.de/projects/crumpet/
http://fit.gmd.de/projects/SAiMotion/

Please contact:
Reinhard Oppermann — GMD
Tel: +49 2241 14 2703
E-mail: Oppermann@gmd.de

Reasoning about Interactive Systems with Stochastic Models

by Gavin Doherty, Giorgio Faconti and Mieke Massink

TACIT is a research network addressing its research towards integration and unification concerns and the use of novel technologies for the support of continuous interaction techniques. The network includes as partners three members of ERCIM (CNUCE, RAL, FORTH) together with the German Institute for Artificial Intelligence (DFKI) and the Universities of Bath, Grenoble, Parma, Sheffield, and York.

Interactive systems in the modern world are becoming both increasingly pervasive, and increasingly rich in the variety of tasks supported, the amount of information available, and the different ways in which the user can interact with them. Interacting with such systems involves multiple media, supporting a continuous flow of information. The shift towards more continuous interaction means that important properties of such systems are better expressed in terms of some quality of service parameter with time playing a central role (latency and jitter are often critical to the usability of a system).

With this motivation, an approach to modelling is to apply timed specification and analysis techniques to encode a set of assumptions about user behaviour in a user model. Although human behaviour is inherently non-deterministic it can be expected to follow probability distributions, and so an interesting possibility is to apply stochastic techniques to consider uncertainty in user models.

The following example illustrates how stochastic models may be used to represent both user and system behaviour. The example refers to the MagicBoard, an augmented reality interactive whiteboard based on a finger tracker and image projection, developed at the University of Grenoble (see figure 1).

What distinguishes this application from an electronic whiteboard is that the user draws directly on the board with real pens; what is drawn on the board is scanned by the system by means of a camera. The user uses gestures to select, copy, paste, etc. what is drawn on the board. The results of the operations are projected back on the board. Thus there is gestural input from user to system, and feedback

both of the user's input (real ink on the board) and from the system ('digital ink' projected on the board). User's gestures are captured by the system that recognise and track a user's finger. A cursor is displayed (by projection) on the board when the system is tracking and gives the user feedback about whether it is able to follow the position of the finger. We use the stochastic process algebra approach to model this aspect of the MagicBoard interface and examine the relation between system performance and the constraints placed on user's behaviour. The situation that we examine is one in which the users need to slow down their movement to allow the system to keep up.

In order to develop a model we make a number of assumptions about the real-time behaviour and its variability of both system and users. Part of the necessary information is obtained from the literature. In particular, the delay generated by the tracker is assumed to be modelled by an exponential distribution because each lag occurs independently from any previous lag (memoryless property).

The timing aspects related to the behaviour of the user are of a different nature. We need to know how much time it takes for users to move their finger from one point to another on the whiteboard in an unconstrained way. Further we need to know how much the movement slows down when the user observes a lag in feedback.

In order to have reasonable data for our illustration, we require data on the human performance for moving a finger over large distances and in different directions. We have available experimental data collected at the University Joseph Fourier

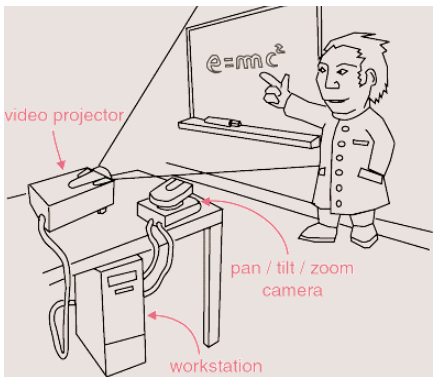

Figure 1: The Magic Board concept.

in Grenoble, showing that the variability in time of fingers movement from one place to another on the whiteboard follows a lognormal distribution. Since the distances that are covered are relatively long, the initial part of the movement is performed very quickly; from the motor skills literature we take this as the initial ballistic part of the movement. The last part of the movement is performed more slowly and we take this as corresponding to the visually guided part of the movement to the target.

Following the (uninterruptible) ballistic part of the movement, the users check whether the cursor is managing to follow their finger. A delay may be introduced at this point before the final part of the movement. Finally, we must formulate our assumptions about the threshold of time for the user to take account of the lag (taken as cognitive delay) and the delay introduced by the user taken as a combination of cognitive and motor delay. For these delays we use the bounds data from a model of the human information processor. A similar argument holds for the delay introduced by the user, which we estimate to be uniformly distributed.

Based on these assumptions, we construct a stochastic process algebra model, presented in Figure 2, that describes the relevant parts of system and user behaviour. We used the model to obtain a histogram of the distribution of the time it would take to reach the target for 10000 different 'runs' of the simulation of the model. Figure 3 presents the simulation results for this model showing the distribution of time until the target is reached for a variety of system performances.

As we can see, there are two modes, corresponding to the waiting and non-waiting conditions. When system performance is good, the non-waiting mode dominates (curve on the left); as performance degrades, it shifts to a bimodal distribution (curves in the middle), and as it degrades further the waiting mode dominates (curve on the right). The shift to a bimodal distribution corresponds to a system that has an average delay of between ca. 60 and 200 ms. The tracking example above shows how stochastic models could help to visualise the possible impact of assumptions made about user and system behaviour on the overall interaction in the early stages of interface development.

However, the modelling of assumptions about user behaviour requires a very careful approach. There are a number of problems and limitations that have to be taken into account. The best choice of modelling approach may not always be evident, and tool support is still an active area of research.

Conclusions

Stochastic modelling, simulation and analysis using stochastic automata is still

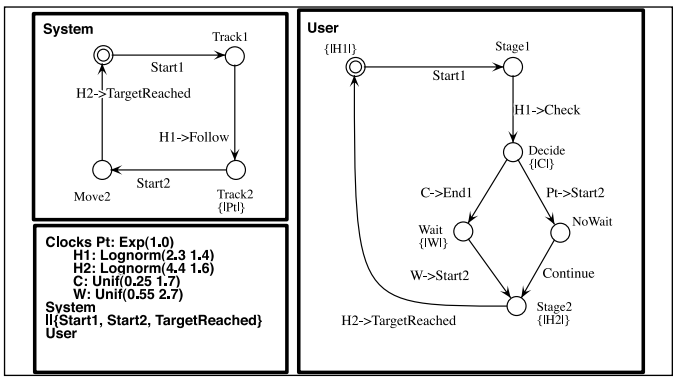

Figure 2: Finger tracking model.

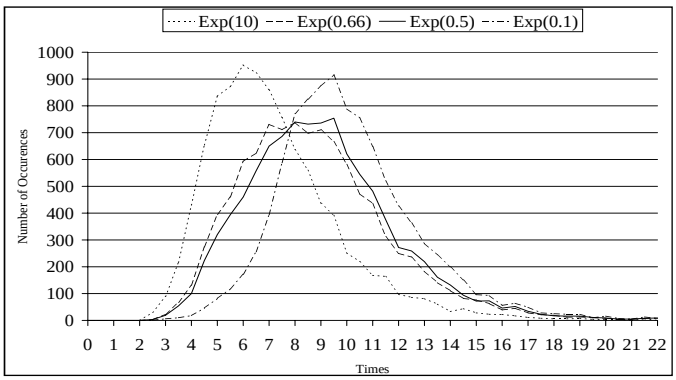

Figure 3: Results for different tracking performance

a relatively new field. The expressiveness of the specification languages, the theories concerning analysis of specifications, and their incorporation into automated support are still at an early stage of development.

We believe such techniques have an exciting potential for modelling performance in interactive systems, taking into account the abilities and limitations of both user and system. Such models allow us to generate a richer set of answers to design questions, which enables meaningful comparison of the results of an analysis to human factors data and other empirical evidence. Two specific areas for investigation are to look at the different ways in which the approach can be used, and at sources of performance data and how they can be integrated.

Link:
http://kazan.cnuce.cnr.it/TACIT

Please contact:
Giorgio Faconti — CNUCE-CNR
Tel: +39 315 2990
E-mail: Giorgio.Faconti@cnuce.cnr.it

Hand-Gesture Recognition for HCI

by Atid Shamaie, Wu Hai and Alistair Sutherland

The goal of this project in the School of Computer Applications at Dublin City University is to develop a system to recognize hand gestures from real-time video images, even in the case of occlusion. Our system uses a multi-scale approach, Principal Component Analysis (PCA) and graph matching.

The main applications of gesture recognition are communicative (eg Sign Language recognition) and manipulative (eg controlling robots without any physical contact between human and computer). Irish Sign Language (ISL) is the native language of the Irish deaf community and contains thousands of gestures. Each gesture consists of a specific hand-shape and a specific hand-motion (see Figure 1).

Hand-shape recognition

First we gather a data set of all the hand-shapes we wish to recognise. A naive approach to recognizing a new image D would be to simply compare it with all the images stored in the data set and find the target image T with the closest match. But because there are so many images in the data set this will take far too long. We can reduce the time by using a multi-scale approach. We divide up the data set into groups of images, which are similar to one another by blurring the images at different levels so that small differences between similar images will be eroded. Thus a whole group of original images may become reduced to just one image, which represents the entire group (see Figure 2).

As a first step, we need to search only those images, which remain distinguishable after blurring and find the target image T1 which matches D best. As a second step, we reduce the blurring slightly and then search only those images in the group represented by T1. We find the image T2 which matches D best. We then reduce the blurring still further and search only those images represented by T2. We find the next target T3, and so on, until the blurring has been reduced to zero. Suppose we can search 10 images at each step. It takes only 4 steps to find the best match out of 10000 images.

Hand-motion recognition

PCA is a technique which allows us to represent images as points in a low-dimensional space. If each image is composed of 32x32 pixels whose values vary from 0 to 255, then each image defines a point in 1024 dimensional space.

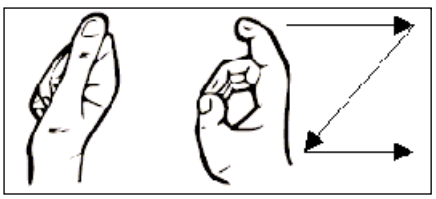

Figure 1: The hand on the left is 'a' in ISL. The right one is 'z', which includes movement.

Figure 2: Blurring the image at different levels.

If we grab a sequence of images representing a gesture then this sequence will generate a sequence of points in space. However, this set of points will usually lie on a low-dimensional subspace within the global 1024D space. The PCA algorithm allows us to find this sub-

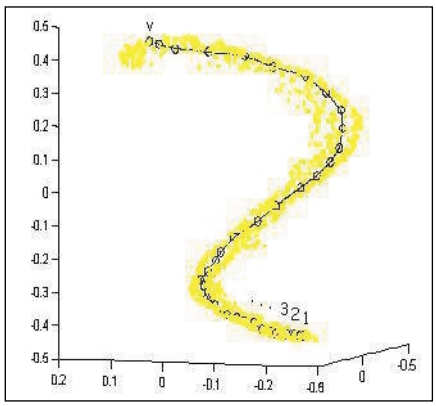

Figure 3: A graph representing a gesture in a subspace.

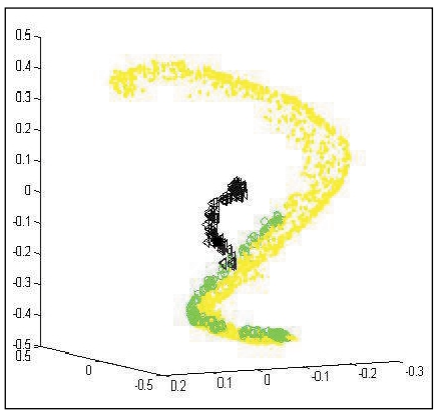

Figure 4: Projection of two different new gestures into a subspace. The green graph matches the known graph (yellow) but the black one is not.

space, which usually consists of up to three dimensions. This allows us to visualise the sequence of points representing the gesture (Figure 3). We can represent this sequence of points by a graph. We can find a different sub-space and graph for each gesture in the vocabulary of our system.

In the recognition phase we capture an image sequence of a new gesture and try to classify it as one of the known gestures in the system's vocabulary. We project the new gesture into each of the sub-spaces and compare the graph of the new gesture to the graphs of the known gestures. We have developed a graph-matching algorithm that finds the best match based on finding the shortest edges of a complete bipartite graph (see Figure 4).

One of the advantages of this algorithm is its ability to recognise partially-occluded hand gestures. Even if there is a gap in the new gesture sequence due to temporary occlusion it is still possible to compute a match. In addition, it is possible to identify the start and end points of the new gesture. It is even possible to use this algorithm when the gesture is made in reverse temporal order, since the graphs are the same in each case.

Future Work

Since speed is crucial in this application we may have to parallelise our algorithm if the number of gestures gets large. This algorithm is inherently parallel at different levels. Not only projecting the unknown input gesture into n subspaces and forming n graphs could be distributed to several processing units but the graph matching algorithm could also be run on several processing elements to reduce the running time.

The phenomenon of 'co-articulation', in which one gesture influences the next in a temporal sequence, is very important in fluent Sign language. Recognition of co-articulated gestures is one of the hardest parts of gesture recognition. We believe our graph-matching algorithm can be adapted quite naturally to cope with this problem. Finally detection of several gestures connected sequentially in a sentence and understanding of the whole sentence is ultimate goal of this research.

Link:
http://www.compapp.dcu.ie/~ashamaie/mvg

Please contact:
Alistair Sutherland — Dublin City University, Ireland
Tel: +353(1) 7005511
E-mail: alistair@compapp.dcu.ie

Improving Speech Recognition for Communication-oriented Activities

by Andrew Sears

Speech recognition (SR) can be a powerful tool, especially for individuals with physical impairments that limit their ability to use a traditional keyboard and mouse. While SR technology continues to improve, users still experience significant difficulty creating and editing documents. To address these difficulties, researchers at the Laboratory for Interactive Systems Design at University of Maryland (UMBC) and the IBM TJ Watson Research Center are investigating the processes by which users interact with SR systems, the difficulties users encounter, and techniques that will make SR more effective.

Our focus is on dictation-oriented activities (eg, writing email, letters, memos, or papers) as opposed to command-and-control activities (eg, turning lights on and off, answering the phone). Our goal is to improve the users' experience as they interact with SR systems—allowing for greater productivity and increased satisfaction. This is accomplished by investigating how users interact with these systems, where they experience difficulties, why those difficulties occur, and the consequences they experience as a result of these difficulties.

Many researchers are investigating techniques to reduce the number of recognition errors, resulting in substantial improvements in the underlying SR algorithms. In contrast, we are interested in assisting users in correcting those recognition errors that still occur. Our initial study included individuals with high-level spinal cord injuries as well as traditional computer users with no physical impairments that hindered their ability to use a keyboard and mouse. Both groups of users completed a variety of dictation-oriented tasks using a state-of-the-art SR system. Our results confirmed that these individuals spent less than 35% of their time dictating and over 65% of their time correcting errors in the dictation. Of particular interest was the fact that these users spent over 32% of their time navigating from one location to another within the document. The results from this preliminary analysis provide strong support for the view that more effective navigation and error correction will support substantial increases in productivity.

A more detailed analysis of the difficulties users experienced while navigating within the documents they created revealed important patterns. Nearly 18% of all navigation commands fail. Over 99% of these commands fail due to recognition errors, users' issuing invalid commands, or users' pausing in the middle of issuing a command. Further, over 99% of these failures lead to one of three results: modifying the content of the document, moving the cursor to the wrong location, or no changes at all. By understanding the underlying reasons for these failures, and the consequences users experience as a result of failed commands, we can prioritize future efforts. Our initial efforts to make navigation more effective have included modifying the way commands are processed and changing the commands that are available to users. We have evaluated these changes through a longitudinal study involving fifteen participants. A preliminary analysis of the results suggests a strong connection between navigational-efficiency and

productivity. A detailed analysis of the results from this study is underway.

While SR algorithms continue to improve, new applications will be found that involve more complex tasks, noisy environments, and users with more diverse speaking patterns. As a result, we believe that recognition errors will continue to be a significant problem for the foreseeable future. To allow users to experience the full potential of SR for dictation-oriented applications, we must improve both the underlying recognition algorithms and the processes by which users interact with SR systems. Our initial study confirmed that users experience significant difficulty correcting recognition errors when they do occur—with difficulties navigating within documents accounting for almost one-third of the users time. By documenting the causes and consequences of navigational difficulties, we were able to identify specific changes that eliminated some failures and changed the consequences of others. We were also able to revise the navigation commands to provide greater power while reducing the complexity of the commands. These changes have been evaluated and results will be reported shortly.

Future studies will continue to focus on the processes by which users interact with SR systems to provide additional insights into the difficulties users experience, where these difficulties occur, why they occur, and the consequences users experience.

Link:
http://research.umbc.edu/~asears

Please contact:
Andrew Sears — University of Maryland, Baltimore County, USA
Tel: +1 410 455 3883
E-mail: asears@umbc.edu

Bimodal Voice/Video-Based Speech Interfaces

by Pavel Zikovky and Pavel Slavík

Synthesizing the movement of an artificial face has not only the role of 'feeling' for the user, it also transfers information. Moreover, it is proved that a whole speech can be recognized from the movement of speaker's lips (for example, by deaf people). So, applying bimodality (visual and audible) to speech signal increases the bandwidth of information transferable to the user.

The strong influence of visible speech is not limited to situations where auditory speech degradation occurs. Moreover, a perceiver's recognition of sound depends on visual information. For example, when the auditory syllable /ba/ is dubbed onto video face saying /ga/, the subject perceives the speaker to be saying /ga/. Another benefit of facial animation is that the only information transmitted over the Internet is textual, and therefore low-sized and fast. This text is input to speech and visual engines on client side.

Lips animation

Lips animation is done by setting the face to the position of the appropriate viseme (viseme corresponds to phoneme) and then morphing to the following one. The viseme parameter allows choosing from a set of visemes independently from the other parameters, thus allowing viseme rendering without having to express them in terms of other parameters, or to affect the result of other parameters, insuring the correct rendering of visemes

Adapting Face Presentation to Different Platforms

As users can access discussion (chat) over the Internet, they can obviously access it from different HW and SW platforms. Moreover, the computers may have different display capabilities and computational power. So, aiming to provide optimal solutions on a variety of different computers, we will use different algorithms for displaying the face. For example, a face displayed on a handheld device that has slow processor and low-resolution display will be displayed as a simple 2-D model, while on a well-equipped PC the same face will be displayed in a full color 3-D model. The following explains how displaying will be changed and divided across platforms and devices.

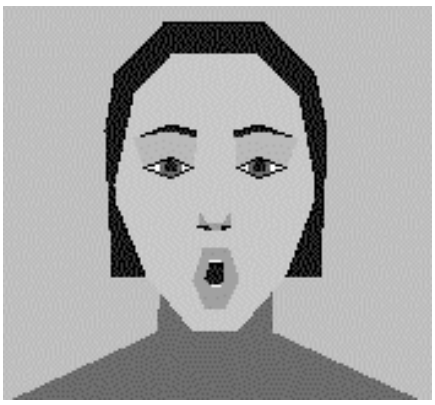

Viseme 'o'.

The advantage of handheld (such as Palmtops, Psions etc.) devices is that they can be carried anywhere by anyone, eg without the need of external power and any special skill. But nowadays they are usually not equipped with appropriate displays and computational power to display any complicated graphics. So, for those devices simple 2D drawing is the best solution. As can be seen on Figure 1, simplifying the face in 2D set of triangles can be quite meaningful. The disadvantage of this 'oversimplifying' is that it cannot express the position of the participant in space. On the other hand, it can be very useful while used on handheld devices as simple autonomous bimodal agent (such as a visitor's guide).

When users access the system from a common low-end computer, notebook or high-end portable device, there is no need to reduce displayed information only to simple 2D polygons due to computational power. So, it is possible to use a simple polygonal 3D model of the head, which gives us additional informational space to use. It gives us the chance to give an expression of more users organized in a space, eg around the table, etc.

On high-tech devices we can do both an excellent 3D animation, which is only a better remake of a previous possibility, or excellent full-color 2D animation — such as television — which is our goal. The animation of the face is done by morphing real visemes (photographs of a real person) of the speech as it flows in real time. This is computationally expansive and requires a lot of memory, but gives a perfect video of a person speaking.

Conclusion

A research has been conducted about capabilities of possible devices, which lead us into creating three different clients for different types of machines and environments. There is the possibility to adapt the look of a face, and this can be used in multi-user environments. The simplest way of face representation on handheld devices gives us the possibility of creating simple portable bimodal interfaces. As the visual quality can be adapted according to computer and platform and the only transmitted information is text, it makes our attitude accessible within huge number of devices which are not capable of common videoconferencing. While thinking about real bimodal chat over the Internet, we have used Jabber (www.jabber.org) as a network communication layer.

Please contact:
Pavel Zikovky and Pavel Slavík — Czech Technical University, Prague
Tel: +420 2 2435 7617
E-mail: xzikovsk@sun.felk.cvut.cz, slavik@cs.felk.cvut.cz

The Virtual Glove Box — A New I/O Device

by Detlef Krömker and Christian Seiler

The Virtual Glove Box is a project which started in 2000 and is currently under development at the Fraunhofer Applications Centre for Computer Graphics in Chemistry and Pharmaceuticals in Frankfurt a.M., Germany. The core of the project is a new Virtual Reality I/O device in which virtual objects are shown in a stereoscopic display and can be manipulated with both hands by the user through haptic feedback devices.

The employment of simulation technology together with the exploration of 3D virtual environments creates a unique advantage in diverse fields such as planning, rapid prototyping, marketing and presentations, or information dissemination and training. This is particularly true if the focus is on inherently three-dimensional structures and their spatial layout. Such structures range from the macroscopic (eg assembly of parts, architecture etc.) to the microscopic (nanomachines, molecules, etc). However the intuitive usage of 3D computer graphics is often hampered by the fact that 2D interaction metaphors are used for interacting with 3D content. This makes it non-intuitive to manipulate three-dimensional content and is a reason why most people find it difficult to navigate and to orient themselves in a virtual 3D world. With the motivation to leverage the use of 3D computer graphics and virtual reality for different purposes and different user groups, we introduce an innovative dedicated I/O device. This interaction device overcomes limitations of traditional interaction devices, even those that are dedicated to 3D (eg the space mouse, or the PHANToM).

In order to overcome the problems and obstacles described above, we introduce a new I/O device called the 'Virtual Glove Box'. A glove box is an apparatus used in chemistry or biology to work on a targets in a closed atmosphere, without contaminating the substances or endangering the user. A glove box usually consists of a transparent hull in which the experiments can be performed. The user can reach into the box through gloves which are attached at the inside of the box. The gloves can be reached through corresponding holes in the box.

We used the glove box metaphor to build our Virtual Glove Box. The Virtual Glove Box combines a stereoscopic display with haptic displays. The stereoscopic part consists of a two projector system with polarised lenses which continuously display images for each eye. The user views the scene through matching polarised glasses. Directly below the display frame are two openings through which the user can reach into the box. The user cannot see his real hands but rather a computer generated representation which mirrors the user's motions. We achieve haptic feedback by using a Cyber Grasp exoskeleton for each hand. In this way the user can actively reach into the virtual domain to grab and move the virtual objects there. One of the technical obstacles in this set-up is to avoid the user reaching into the projector beam. Another problem with an exoskeleton-only approach is that gravity cannot be displayed in this set-up. We are currently evaluating different approaches to solve this problem, for example, the inclusion of a robot arm.

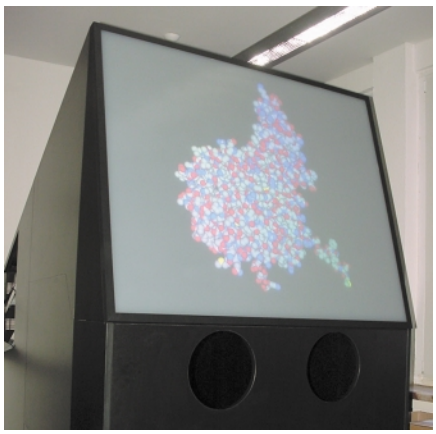

Virtual Glove Box prototype.

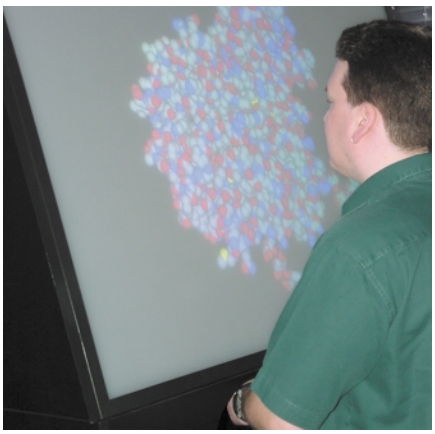

User in front of the Virtual Glove Box.

The usage of a real life interaction metaphor like a glove box greatly increases the user's acceptance for such a new device. The Virtual Glove Box enhances the immersion since the haptic exoskeletons worn by the user are hidden from view by the apparatus, and a holistic approach is thus presented for immersive VR devices with haptic feedback. The Virtual Glove Box offers a multitude of different fields of application wherever interaction with virtual objects is required. The foremost area of application as described above is the planning, modelling and simulation of structures on a microscopic level and smaller, like molecular structures of three-dimensional objects from production plants, machinery, pipeline systems etc. The usage is not limited to real objects, since abstract systems like work flow diagrams can be similarly experienced. A very important aspect of the device is the communication presentation and visualisation of new ideas or solutions, which opens the usage of the Virtual Glove Box to many different sets of users. Experts benefit from intuitive interaction because they can concentrate on their domain competences (eg design or mechanics) rather than having to learn to interact with complicated authoring systems. Users from management or marketing, and customers, share the same benefit: things can be easier to communicate and understand if you have a 'hands-on' experience.

We have described a new I/O device for intuitive human computer interaction. The device itself offers many advantages over conventional I/O devices as described above. Beside the implementation of the device there is another important aspect of the Virtual Glove Box project: the user interacts with virtual objects. Therefore the description of such an object must include not only visual properties, as in other VR applications, or the haptic description, but also a behavioural description. We are currently improving a component-based description model to incorporate the different aspects of a virtual object (visuals, haptics, simulation, etc.) into a software framework.

Please contact:
Detlef Krömker — Fraunhofer Applications Center for Computer Graphics in Chemistry and Pharmaceutics, Germany
Tel: +49 69 97 99 51 00
E-mail: krömker@agc.fhg.de

Novel User Interfaces for the Future

by Veli-Pekka Räty

VTT Information Technology is developing user interfaces with which people interact with movements, gestures and touches — instead of conventional computer controls. An example of VTT's latest innovations is an interface for computer games. The interface, developed in the Lumetila (Virtual Space— User Interfaces of the Future) project, replaces conventional keyboards, mice and joysticks with person's body movements in interacting with the computer game and with other players.

Today, games use the same controls as word processing, spreadsheet, database or Internet surfing. The recent development of computer game controllers is trying to take advantage from force feedback joysticks and steering wheels. However, these solutions still rely on the old way of controlling the game by using mainly one's hands and fingers. Also, the player is alone in front of his computer.

Players Share Experiences

In fact, the novel bodily and spatial interface is in a room where the body movements of the players are recognised by pressure sensitive sensors installed in the floor. From the sensors the information is transmitted to the computer, which processes the information and shows the game graphics on the large screen.

The room is also equipped with a surround sound system and light effect devices. These special effect devices and a real time 3D graphics engine together with the team play are essential in creating the immersion and experience in the Virtual Space. The fact, that the players need not wear any virtual reality devices, like data glasses or gloves, helps in creating a total immersion and enables the players to work as a team. The boundaries between the physical space and the interactive virtual space that is presented to the players has thus been weakened. This results in an interactive natural environment with shared experiences between the players.

A group of school children steering the diving bell Nautilus.

Nautilus — a Submarine Adventure

For testing the Lumetila concept VTT Information Technology created a game, the LumePong, based on the well-known Pong game from the early 1970s. The test results and experiences were used in developing, planning and designing a new game called Nautilus.

With Nautilus a group of players can make a virtual trip into the underwater world. Their task is to rescue a dolphin that is trapped in a ship wreck. Along the journey to the wreck the team gets score points when collecting star fish and sea horses. Also, there may be dangers under water.

The game is based on cooperation of a group of 5-8 players. When the group goes to the right the submarine turns to the right. If the group steps towards the front edge of the room the speed of the vessel increases. As the group moves backwards the speed of the vessel decreases and the vessel reverses. The whole underwater environment can be seen on the screen. Furthermore, strong light and sound effects are used to enhance the immersive experience, for instance in the case of a collision with a shark. In addition to recreation, the trip involves also physical exercise.

With Nautilus the research group has received essential information for developing advanced user interfaces for entertainment applications, for instance those for science centres and theme parks. This technology and knowledge is also useful in other applications at home and at work.

In the Lumetila project VTT Information Technology works in collaboration with the media company Cube Ltd., Nokia Research Center, theme park Tampereen Särkänniemi Ltd. and the University of Lapland. Tekes, the National Technology Agency of Finland cofunds the project.

Focus on User Friendliness

In all its research activities VTT Information Technology focuses on users and their context of usage. From the very beginning and throughout the whole design process the main guidelines are: who are the users, where are they, and what is the nature of their everyday life. In this Human-Centred Design (HDC) approach, feedback from the users is an essential source of information. The HDC approach is also iterative. The current context is only the starting point for design.

VTT has already applied the HDC approach in many Internet services and mobile systems. Lumetila game concept and user interface is an excellent example of comprehensive application of the Human-Centred Design approach.

Links:
Intuitive Environments:
http://www.vtt.fi/tte/research/tte5/tte54/
Lumetila:
http://www.vtt.fi/tte/projects/lumetila/

Please contact:
Veli-Pekka Räty — VTT Information Technology
Tel: +358 3 316 3340
E-mail: Veli-Pekka.Raty@vtt.fi

Roomware: Towards the Next Generation of Human-Computer Interaction

by Norbert A. Streitz

Roomware® consists of computer-augmented room elements with integrated information and communication technology facilitating new forms of human-computer interaction. They are part of an approach that the 'world around us' is the interface to information and for the cooperation of people. The Roomware® components were developed at GMD's Integrated Publication and Information Systems Institute (IPSI) in Darmstadt.

The next generation of human-computer interaction is determined by a number of new contexts and challenges. One major challenge is to overcome the limits of desktop-based information environments currently in use. In the past, the introduction of information technology caused a shift away from real objects in the physical environment as information sources towards monitors of desktop computers at fixed places as the interfaces to information. Accordingly, user-interface designers developed the known types of human-computer interaction for the desktop paradigm. In contrast to this, we developed a new approach that emphasises again the relevance of physical objects. In this approach, the 'world around us is the interface to information' where traditional human-computer interaction will be transformed to human-information-interaction and human-human communication and cooperation. One result of this research are so called 'Roomware®' components for the Workspaces of the Future in so called 'Cooperative Buildings'. They require new forms of interaction with information, wich are supported by the software we developed.

Roomware®

The Roomware® components were developed in the AMBIENTE-division at GMD-IPSI in Darmstadt as part of the i-LAND environment (Streitz et al, 2001). Roomware® results from the integration of information technology into room elements as, eg, walls, doors, and furniture. Roomware components are interactive and networked; some of them are mobile due to independent power supply and wireless networks, and are provided with sensing technology. The figure shows already the second generation of roomware components which was developed together with partners from industry in our R&D Consortium 'Future Office Dynamics'. In the following, some examples of roomware are described.

DynaWall® and InteracTable®

The DynaWall® in the AMBIENTE-Lab is an interactive wall covering one side of the room completely. The size of 4.50 m width and 1.10 m height and the very smooth integration of this very large display into the architectural structure creates the impression that you are really writing and interacting with a wall or wallpaper. The surface is touch-sensitive so that you can write and interact on it with your bare fingers or with a normal pen (no electronics needed). Several people can write/ interact in parallel in (currently three) different areas of the DynaWall. Beyond these physical affordances, our BEACH software enables very intuitive interaction based on gestures that are

The second generation of Roomware® components.

The Disappearing Computer Initiative

As part of the European IST programme, 'Future and Emerging Technology' launched the proactive initiative 'The Disappearing Computer' (DC). After a careful review and selection process, 16 projects were finally accepted for funding. An overview of the projects can be found at http://www.disappearing-computer.net/projects.html.

The projects started on January 2001. The duration of the projects is between two to three years. 37 institutions from academia as well as industry in 13 countries are participating in the 16 projects. The total effort of all projects is close to 300 person years, 50 % of it being funded by the European Commission.

The objectives of the 'Disappearing Computer' initiative are to explore how everyday life can be supported and enhanced through the use of collections of interacting artefacts. Together, these artefacts will form new people-friendly environments in which the 'computer-as-we-know-it' has no role. The aim is to arrive at new concepts and techniques out of which future applications can be developed. Specifically, the projects of the initiative will focus on three inter-linked objectives:

- Developing new tools and methods for the embedding of computation in everyday objects so as to create artefacts.
- Research on how new functionality and new use can emerge from collections of interacting artefacts.
- Ensuring that people's experience of these environments is both coherent and engaging in space and time.

Beyond the individual projects, the DC initiative is planning a range of inter-project collaboration activities as, eg, Disappearing Days workshops, Troubadour grants, Research Ateliers, Jamborees as conference-like major events, etc.

The DC-initiative started their collaboration and networking activities by electing a steering group of four people: Norbert Streitz (chair) (GMD-IPSI, Darmstadt, Germany), Lorna Goulden (Philips Design, Amsterdam, The Netherlands), Spyros Lalis (ICS-FORTH, Heraklion, Greece) Paddy Nixon (University of Strathclyde, Glasgow, UK). The initial kick-off meeting of all DC-projects took place on January 31 in Darmstadt starting to form the DC community. The steering group can be contacted via dc-sg@disappearing-computer.net. The DC website has been launched and will soon become the primary place of information and communication for the DC community.

Links:
The Disappearing Computer Initiative: http://www.disappearing-computer.net
EC Future and Emerging Technologies: http://www.cordis.lu/ist/fethome.htm

Please contact:
Norbert Streitz — GMD
Tel: +49-6151-869 919
E-mail: streitz@darmstadt.gmd.de

reflecting actions with physical objects in the real world (eg, "take and put", "throw", "shuffle" , ...). When throwing objects (with different accelerations), the speed and thus the flying distance is dependent on the initial momentum provided by the user. People can interact this way immediately after having seen it once.

A similar way of interaction is provided by the InteracTable®, another roomware component we developed. It has a display size of 65 cm x 115 cm and a diameter of 130 cm. Beyond the type of interactions available at the DynaWall, it provides additional ones required by horizontal and round or oval-shaped displays. To this end, we developed in BEACH special gestures for shuffling and rotating individual information objects or groups of objects across the surface so that they orient themselves automatically. This accommodates easy viewing from all perspectives. Furthermore, one can create a second view of an object and shuffle this to the other side so that all team members have the correct view at the same time. Other examples of roomware components are the CommChairs® and the ConnecTables®. More details on the different roomware components are given in [Streitz et al (2001). Roomware: Towards the Next Generation of Human-Computer Interaction based on an Integrated Design of Real and Virtual Worlds. In: J. A. Carroll (Ed.), Human Computer Interaction in the New Millennium. Addison-Wesley].

The Passage mechanism

Passage is a mechanism for establishing relations between physical objects and virtual information structures, thus bridging the border between the real world and the digital, virtual world. So-called Passengers (Passage-Objects) enable people to have quick and direct access to a large amount of information and to 'carry them around' from one location to another via physical representatives that are acting as physical 'bookmarks' into the virtual world. It provides an intuitive way for the 'transportation' of information between roomware components, eg, between offices or to and from meeting rooms.

A Passenger does not have to be a special physical object. People can turn any

object into a Passenger: a watch, a ring, a pen, glasses, a wooden block, or other arbitrary objects. The only restriction Passengers have is that they can be identified by the Bridge and that they are unique. Passengers are placed on so-called Bridges, making their virtual counter parts accessible. With simple gestures the digital information can be assigned to or retrieved from the passenger via the virtual part of the Bridge. The Bridges are integrated in the environment to guarantee ubiquitous and intuitive access at every location in a building (=> Cooperative Building).

Future Work

At the beginning of 2001, we started a new project called 'Ambient Agoras: Dynamic Information Clouds in a Hybrid World'. It is funded by the European Union as part of its proactive initiative 'The Disappearing Computer'. 'Ambient Agoras' will provide situated services, place-relevant information, and a feeling of the place to the users. It aims at turning every place into a social information marketplace (= agora in Greek) of ideas and information where one can interact and cooperate with people.

Links:
http://www.darmstadt.gmd.de/ambiente
http://www.future-office.de
http://www.roomware.de
http://www.Ambient-Agoras.org

Please contact:
Norbert A. Streitz — GMD
Tel: +49 6151 869 919
E-mail: streitz@darmstadt.gmd.de

Human Computer Interaction and Virtual Worlds

by Jutta Becker, Ivo Haulsen and Herbert Rüsseler

The development of scalable Human Computer Interfaces, and the definition of adaptive interaction paradigms, have led to both the usage of standard input devices for 3D interaction purposes and the integration and configuration of innovative interaction devices. The ViSTA team of GMD-First is currently engaged in some HCI related projects for exploring adaptive interaction paradigms. Standard input devices with haptic feedback and sound are used within the scope of VR learning and training environments, and novel multi-touch-sensitive input devices are being developed to improve tomorrow's human/computer interaction with Virtual Worlds. These projects are funded by the Federal Ministry of Education and Research.

Virtual Worlds gain their acceptance through a large number of interaction possibilities. The development of web standards like VRML, X3D and Java represented an important step towards scalable virtual environments. Exploring virtual worlds within immersive environments, such as CAVE systems, demands new interaction paradigms and device technology with respect to common desktop standards.

In general, a vocational learning environment will be centred around a PC or laptop computer. Keyboard, mouse or touchpad are used, and occasionally additional devices like joysticks or trackballs are available. According to a majority of different learning contexts, interaction methods must be adapted to the actual user and his/her learning requirements. Different learning scenarios provide the user with interactive simulations or predefined animations. Besides navigation through and exploration of objects and environments, a variety of manipulation possibilities assist in the training.

In common HCIs for VR systems, navigation and manipulation interaction methods are predefined and cannot be changed without new compilation. The scalable interaction model allows the configuration of input devices by reading an XML configuration document, and defining the input channels and their transformation, without new compilation. Mouse movements can be used alternately for translation, rotation or other parameter variations. Additional buttons or wheels are programmable without developing new functions. According to the user's role and classification (eg expert/beginner, technician/manager), certain interaction and configuration possibilities are either hidden or made available.

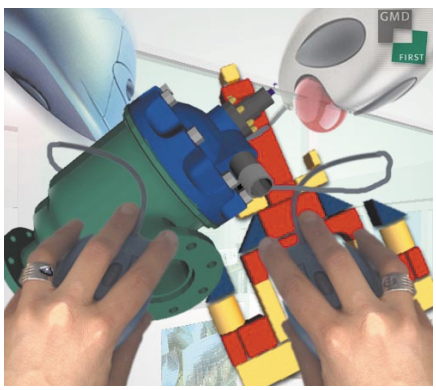

Standard input devices and multi-channel interaction paradigms improve HCI with virtual worlds.

We integrate additional spatial audio and tactile feedback into the training environments. For instance, the collision of objects during manipulation is combined with audio and force feedback. The new IFeel-Technology™, integrated in standard mouse devices, is available for every PC and laptop computer.

In other project areas, we are developing, with our partners FHG Institute IAP, pressure-sensitive devices based on polymer materials, advanced electronics and intelligent interaction paradigms. As well as the exploration of organic materials for future low-cost multisensorial devices, we will also be examining emotion and gesture processing.

Links:
http://www.first.gmd.de/vista
http://www.first.gmd.de/vista/hci

Please contact:
Jutta Becker — GMD
Tel: +49 30 6392 1776
E-mail: bonjour@first.gmd.de

Ivo Haulsen — GMD
Tel: +49 30 6392 1777
E-mail: ivo@first.gmd.de

User-centric Multimedia on the Web

by Lynda Hardman

First generation web pages tend to be text-centric and created by hand for a broad audience. Second generation web pages offer some form of user-tailoring by extracting relevant content dynamically from a database and displaying it in a generic template. While this approach is already in use for creating textual documents, research carried out at the Multimedia and Human-Computer Interaction group at CWI aims at providing end-users with on-demand, user-tailored, synchronized hypermedia presentations.

The main goal of the research is to enable the dynamic generation of hypermedia documents that take into account user preferences and make optimal use of system and network resources. In order to achieve this we require to elevate multimedia to the same status as text in the way it can be included in documents and presented to a user. Multimedia document processing has a number of requirements that are fundamentally different from text which make it more difficult to incorporate in a 'traditional' electronic publishing environment. For example, multimedia document formatting cannot be based on text-flow layout rules. These differences make it extremely difficult to apply most commonly used solutions because they are too text specific. The reasons for wishing to tailor text documents to the needs of individual users, however, still apply to multimedia. Current trends are that the Web is accessed using appliances other than a standard PC, the primary mode need not be textual and the contexts in which the information is used are more diverse. Users are also becoming more demanding and expect content to be supplied in their own language and to take their cultural background into account. Our group has developed multimedia models and techniques that go beyond current text-centric approaches and can be employed to cater for the required diversity.

The approach that we are currently investigating is to generate individual, tailored documents on demand, based on the prevailing conditions, rather than having to rely on predicting these beforehand. The goal of our research is to construct an environment which allows variations in end-user platform, network conditions and user preferences to influence the multimedia document generation process.

The Cuypers Architecture

The Cuypers environment is a prototype hypermedia generation environment which is able to select media items from a database or the Web and combine them in a multimedia presentation. During the generation process, network conditions, discourse and interaction models, user preferences, and other varying conditions can be taken into account. The result of the process is the generation of a multimedia presentation that can be transmitted over the Internet and played to the user by a standard Web client. After the relevant media items have been retrieved, they need to be grouped into a coherent presentation that conveys the intended semantics of the author and is also adapted to the specific requirements of the user. In order to achieve an adequate level of adaptation, presentation independent abstractions are required to be able to generate a wide variety of hypermedia presentations while maintaining the underlying semantics.

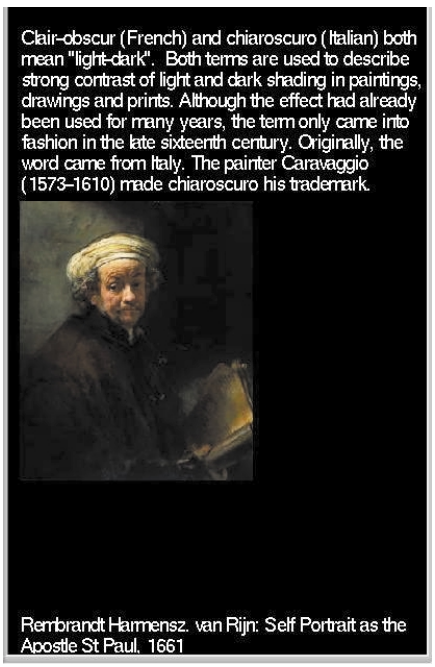

Presentation explaining the chiaroscuro technique as used by Rembrandt, generated by the Cuypers system.

There are various levels of such abstractions, ranging from very high-level abstractions, that focus on the semantics but say little about the presentation-oriented details, to the presentation models used by concrete multimedia file formats.

Towards the Third Generation: Multimedia on the Semantic Web

Third generation Web pages will make use of rich markup, based on XML, along with metadata, eg encoded in RDF (W3C's Resource Description Framework), to make the content not only machine readable but also machine processable. This will allow more directed search of documents, but also the development of new applications making use of the semantic information in the documents. E-commerce applications, for example, can make use of agreed-upon semantic mark-up to produce integrated information from several sources. In the medical world, patient records can be annotated using standard domain descriptions to allow their exchange among hospitals, and even countries.

In our work, we are interested in semantic markup to the extent that it potentially enriches the created hypermedia presentations and facilitates the presentation generation process. When media items are selected from databases we need information, for example, on the potential role the item may play in the integrated presentation. One of our research questions is to what extent this is possible, and what types of metadata actually improve the generation process.

Link:
http://www.cwi.nl/ins2/

Please contact:
Lynda Hardman
and Jacco van Ossenbruggen — CWI
Tel: +31 20 592 4147
E-mail: Lynda.Hardman@cwi.nl,
Jacco.van.Ossenbruggen@cwi.nl

Multimodal Interfaces for Multimedia Applications: Consequences on the Interface Design

by Laurent Romary and Nadia Bellalem

Under the auspices of the European project MIAMM, two teams of INRIA collaborate to define new types of multimodal interactions involving strong language recognition and understanding capacities.

Considering the ever growing size of information available on line, coupled with the wider variety of potential users which are likely to demand some kind of access to it without necessarily being computer experts, it is mandatory to provide access means which are basically based upon human spontaneous communication abilities. Such interfaces must obviously incorporate dialogue component based on speech and gesture, without losing fundamental qualities like simplicity, efficiency, speed and robustness. Conversely, the system must be able to answer either in natural language or by means of an adequate presentation of available information in a graphical way, probably coupled with force feedback devices to get a better grasp on the actual organisation of the information which is being presented.

The aim is thus to design a human-computer flexible interface which is able to adapt to the communicative capabilities of the user by interpreting the user's requests in the application context. The development of such interfaces may have strong consequences on more fundamental research on man-machine communication:

- it implies having robust language processing modules, from speech recognition to deep language understanding, comprising the interpretation of multimodal referring expressions (eg demonstrative and gesture) and the contextual understanding of the action to be taken or answer to be given
- it forces the design of new models for the analysis and understanding of gestural trajectories, there again implying the contextual interpretation of possible designations
- these two aspects have to be combined in generic communication models where the complementarity of the two modes (language and gesture) is understood at a semantic level
- dialogue management is even more crucial in a multimodal environment since it contributes to the building of the context (cf. Eg elliptical phenomena) but also ensures a good continuity of the dialogue by ensuring that the right communication act is understood and produced
- finally, one has to design an optimal graphical representation of the task where objects are clearly identified and organized in such a way that the lay-out of the scene as perceived by the user will be predictable.

In the context of a multimodal interface comprising a haptic device, it is essential to analyse how such a presentation can integrate the third dimension (depth). It should be both more natural from a perceptual point of view, and provide situations where direct action on the scene is preferable to communication with the virtual assistant (situations of emergency, re-organization of the scene etc.). This opens a wide range of possible research that has been barely tackled so far.

These various aspects have been for long part of the core activities of the 'Langue et Dialogue' team at Loria, mainly in the context of multimodal reference interaction with 'classical' designation device, but our strong involment in the future MIAMM European project (Multidimensional Information Access using Multiple Modalities). This project aims at providing a multimodal interface based on a haptic device to provide access to multimedia data. In collaboration with the Parole team at Loria, we will develop the following aspects:

- robust linguistic processing module putting into competition both stochastic models and symbolic ones (based in particular on the TAG (Tree Adjoining Grammar) formalism)
- provide a unified semantic of 3-D gestures as provided by the haptic device, which integrates the relation to perceptual grouping factors
- multimodal dialogue management where both contextual interpretation and interlocution will be managed in a uniform framework to ensure the close interaction between these two functions
- software integration, based on the definition of an multimodal exchange format MMIL (Multi-modal Interface Language), expressed in XML, which will gather up the various results yielded by each module in the system.

The MIAMM project will start in September 2001 and also involves DFKI, Sarrebruck; TNO, Netherlands Organisation for Applied Scientific Research, Delft; Sony Int. Europe GmbH, Stuttgart and Canon Research Center Europe Ltd., Guilford.

Link:
http://www.loria.fr/equipes/led/

Please contact:
Nadia Bellalem — LORIA
Tel: +33 3 83 50 20 33
E-mail: nbell@loria.fr

Constraint-Based Navigation in Three-Dimensional Scenes

by Christoph Wick and Stefan Pieper

Efficient navigation of three-dimensional image data requires visualisation and interaction techniques derived from and suitable for the tasks. Scientists at the GMD Institute for Applied Information Technology (FIT) present two examples for a design principle that improves navigation in complex spatial environments. This approach is called constraint-based navigation in three-dimensional scenes.

Medical informatics often deals with volume data sets acquired by CT, MRI or 3-D ultrasound scanners. In the field of interactive computer graphics, a lot of research has been conducted on generic visualisation techniques such as volume and surface rendering. Interactive computer graphics have also provided basic mechanisms for navigation in such data sets with 2-D pointing devices. Various examples can be found in the Open Inventor library or in VRML viewers.

With the EchoCom application we have introduced a tool to explore 3D cardiac ultrasound volumes and to train the handling of an ultrasound transducer. To be usable as a web-based application, an appropriate mouse interaction technique had to be introduced. A major design goal of the interface was to provide something that resembles the handling of a real transducer and gives maximum orientation in the complex spatial environment. Common 3D interaction tools like a virtual trackball or CAD-like interfaces proved to be counter-intuitive and require a remarkable learning effort from the user.

Our analysis of traditional training in echocardiography showed that transducer handling is always a two-step process: depending on the desired view the transducer is first positioned on one of a few standard points on the patient's chest. After that the transducer is swept, rotated, and tilted until an optimal view is reached. With this in mind, we designed an interface to navigate with a virtual ultrasound transducer as 'naturally' as possible.

A surface model of the thorax serves as a reference scenario. If the user clicks on a spot on the thorax model, the tip of a virtual transducer jumps to that spot and positions itself perpendicularly to the thorax. If the user drags the mouse from here, it follows the mouse pointer. The transducer tip never loses contact with the surface and thus no meaningless positions are achieved. Another option for setting the transducer's initial position is to select one of the standard views defined in echocardiography and continue exploration from there. Apart from moving the transducer over the chest the user can sweep, tilt, and rotate the transducer in a second mode. For that purpose, three intuitive and easily accessible areas on the transducer handle can be 'grasped' with the mouse. Each area lets the user perform one of the three basic orientation operations (see figure).

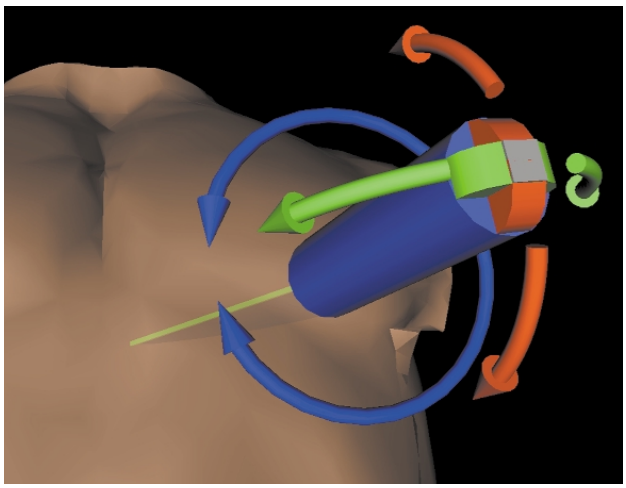

The virtual transducer is positioned on the thorax model and can be swept, tilted, and rotated by dragging at the three touch-sensitive areas (blue, red, and green).

Initial informal evaluations with both novice and trained cardiologists, some familiar and some unfamiliar with 3D interfaces, have shown that such a specialised, constrained interface is easy to use, intuitive, and suitable for the task.

In the TriTex project another application is about to be developed as a visualisation tool for 3D medical and geologic data. These data originate from medical 3D image scanners or seismic surveys mapping earth's sub-surface structures. Because of the data's complexity, manual, slice-by-slice analysis results in very time-consuming and error-prone work. In order to facilitate the manual interpretation, an automated texture analysis component processes the three-dimensional data sets and identifies sub-volumes of homogeneous structural qualities, thus classifying and segmenting the raw data.

The analysis results can be used as an augmentation in a mixed-reality visualisation. However, only with an adequate navigation mechanism will this augmentation guide the user to a more efficient exploration. The key concept to make use of is the relevance of certain border surfaces and their enclosed or adjacent volumetric structures. Thus, the navigation can concentrate on letting the user pick bodies and surfaces of interest and, taking the picked object as a reference structure, allowing him to explore the immediate vicinity.

At a planar border surface as reference structure, for example, the view can be set to a perpendicular cross-section and the navigation can be constrained to translations along this surface. Certain spherical structures can be explored through diametric cross-sections, thus reducing the navigation to only the rotational parameter.

Both examples share a common design principle: efficient navigation in complex environments, such as three-dimensional image data, requires visualisation and interaction techniques derived from and suitable for the tasks. Generic navigation techniques must be constrained to reduce complexity in interpretation of data as well as in interaction with it. Abstract reference structures - standardised, registered models or objects automatically extracted from the raw data - provide spatial constraints, while the remaining degrees of navigational freedom are determined by requirements to visualise proximal visual contents simultaneously. Constraint-based navigation proved to be a successful tool for different applications that were designed to address highly specialised needs. However, to develop an adequate constraint-based navigation scheme it is always necessary to perform an in-depth field analysis. One result of this analysis will be an exact understanding of how information has to be visualised and how reproducible manipulation of views can be accomplished.

Link:
GMD Institute for Applied Information Technology:
http://fit.gmd.de

Please contact:
Jürgen Marock — GMD
Tel:+49 2241 14-2155
E-mail: juergen.marock@gmd.de

Implementation of a Laser-based Interaction Technique for Projection Screens

by Michael Wissen

A laser-controlled interaction technique consisting of a large interactive screen, 'The Interactive Wall', together with a laser-based pointing instrument, has been developed by Fraunhofer Institute for Industrial Engineering (IAO), Stuttgart.

The Interactive Wall forms a 6.5 x 1.6m continuous workspace (see Figure 1). Its dimensions offer users a generous presentation space, as well as various forms of interaction in support of group project applications. Contemporary interaction devices, including pointing devices like the mouse, trackball, trackpoint and touchpad, which operate on the principle of relative motion, deliver less than satisfactory results: the user cannot directly address co-ordinates, but must instead move the pointer from its current position to the target position through relative motion of the device. On large projection surfaces such as the Interactive Wall described here, the disadvantages of relative-motion pointing devices become instantly visible. It takes too long to select widely separated objects, and relative-motion pointing is contradictory to the intuitive use of a projection screen. Since the Interactive Wall works much like a blackboard, the pen metaphor is more appropriate.

Application fields for large interactive screens are to be found primarily in the support of creative processes in conference rooms and in co-operative visualisation and exploration of content, information structures and processes. The system offers users considerable input support through gesture recognition based on a writing instrument.

A pen-like pointing device for use in conjunction with a laser has been developed which can perform mouse functions such as pointer movement and clicking. The mouse pointer follows the visible laser dot on the screen, thereby enabling immediate movement between distant locations by direct pointing. The device also allows mouse pointer control from long distances (up to 8 metres) and from acute angles (25 degrees).

Set up and Operation

At the heart of the Interactive Wall system is a graphics computer with three display adapter cards, each of which is connected to one of three projectors, which combine to create a single continuous workspace. The projectors are each assigned to one of three cameras, which in turn are

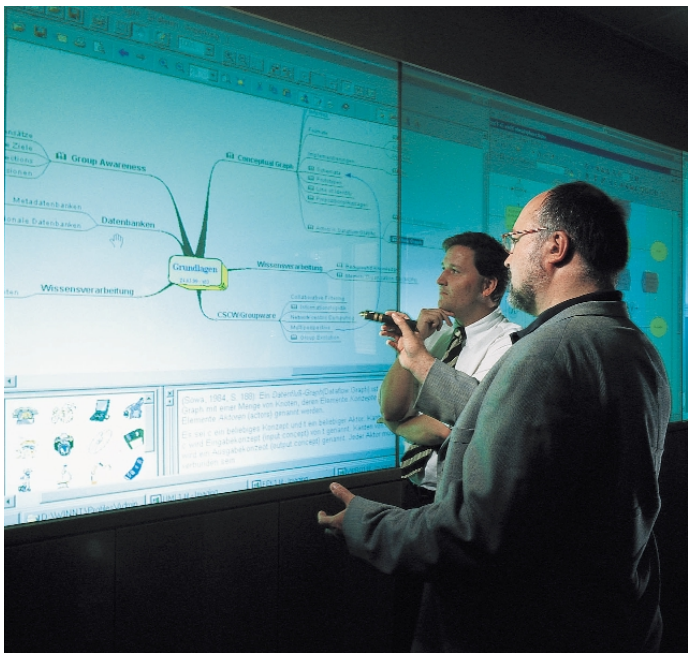

Figure 1: The Interactive Wall in the Knowledge Media Lab

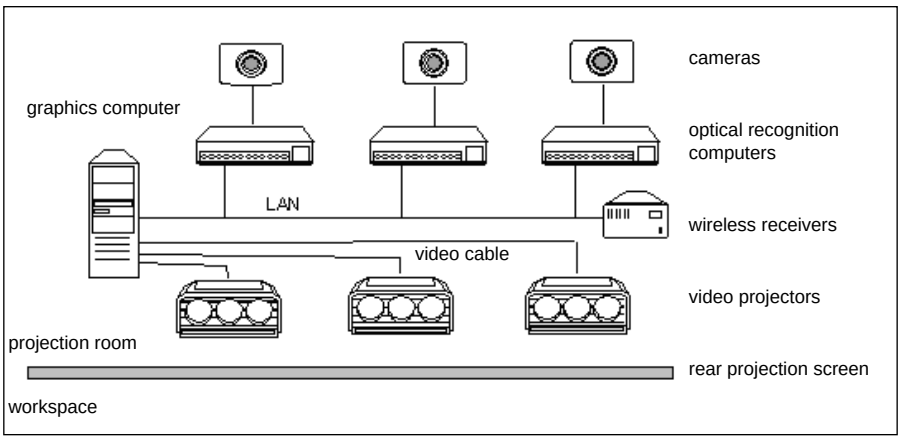

Figure 2: Interactive Wall schematic layout.

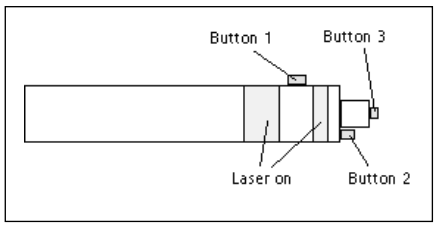

Figure 3: Laser pointer schematic configuration.

connected to three optical recognition computers (see Figure 2).

The camera system checks the projection surface for the appearance of the laser dot. If found by one of the three cameras, the laser's absolute coordinates are calculated and transmitted over a LAN system to the graphics computer, which then positions the mouse pointer at the appropriate spot.

The laser pointer consists of a standard laser encased in a pen-shaped housing with three integrated buttons (see Figure 3). For the laser to remain activated, finger contact must be maintained across two metallic rings at the end of the device — if contact is broken, the laser is deactivated, ensuring a longer service life for the pointer.

An interaction device must not hamper the user's freedom of movement — this is achieved through the use of a laser which can be accurately positioned and recognised even over relatively long distances, and by using wireless technology for the transmission of button command signals.

While the first two buttons emulate the action of the left and right mouse buttons, the third button at the tip of the device serves to fix the mouse pointer on the surface of the interaction screen (similar to left-clicking with a mouse to set cursor position). Thus the user can sketch or write with the laser pen as on a blackboard.

The individual computers communicate with one another using TCP/IP. The graphics computer requires only the installation of a driver to translate incoming data from the optical recognition computers or wireless receiver into the appropriate system messages. It is thereby possible to remotely operate any computer with a network connection to the system components described herein. The desktop display of the remote computer appears on the projection surface, and the co-ordinate data from the mouse pointer is received via the network and converted by the driver into pointer movements.

Interaction Problems

The use of a laser pointer for GUI interaction presents various problems, such as hand tremor, interaction objects that are too small, and some disadvantageous characteristics associated with using a projection system. To some extent, these effects can be mitigated through the control software — for example, various stabilisation strategies were implemented to minimise the effects of natural hand tremor. The desktop designs of Windows and other GUI operating systems are also problematic, as they were developed for mouse interaction, are relatively small, and are difficult to interact with from great distances. Since an application GUI might stretch out over several meters on the projection display, the ability to execute each function should if possible not be restricted to a single point on the desktop. The best example of this is the familiar 'Start' button from the Windows operating systems, usually found in the lower left corner of the display.

Tests show that the use of a laser pen facilitates significantly faster interaction when compared to traditional input devices, provided that the user spends a few minutes beforehand becoming familiar with the basic operation of the device.

Prospects

The interactive wall implementation introduced in this article includes several laser pens, but at present they can only address a single mouse pointer. However, the system can be expanded without a great deal of technical effort to give every user a 'personal' laser pen. Distinguishing between individual laser pens can be accomplished through the employment of pulsing laser beams. In this case, the camera computer would receive not only positioning data from the laser dot, but also information identifying the laser pen so that it can position the corresponding mouse pointer. These additional virtual mouse pointers can be visualised as desired.

Another possibility that would enable users to work intuitively with the Interactive Wall is the detection of hand shadows on the projection surface. This would involve the user touching the screen directly, and would only function with rear-projection systems and under certain lighting conditions.

Please contact:
Michael Wissen — Fraunhofer Industrial Engineering IAO, Germany
Tel: +49 711 9 70 23 80
E-mail: Michael.Wissen@iao.fhg.de

Físchlár on a PDA: A Hand-Held User-Interface to Digital Video

by Hyowon Lee and Alan F. Smeaton

At the Centre for Digital Video Processing, Dublin City University, we are working on diverse access to Físchlár, a web-based digital video processing and management system that allows its 1,000 users within the campus to record, browse and playback broadcast TV programmes for learning and for entertainment purposes. Since late 2000, we have been developing and testing innovative video content browsing interfaces for mobile devices such as PDA (Personal Digital Assistant), to provide our users with mobile access to Físchlár.

As on-going research in digital video management, we have developed an experimental video system called Físchlár which allows a user to initiate the recording of TV programmes from one of 8 major TV channels in Ireland. The requested programme, when broadcast, is recorded in MPEG-1 format and is analysed for automatic indexing with camera shot/scene detection, and extraction of representative keyframes to be used at a user' content browsing interface through a web browser. Clicking on a displayed keyframe on the screen starts streaming the video from the clicked moment onwards, thus allowing the user to watch the recorded content. The system has been deployed since July 1999 in the University's computer labs and residence through a LAN to students and staff, and currently more than 1,000 users have registered and have been using the system to record and browse TV programmes in their own time. This allowed us a campus-wide user study of internet-based video browsing and usage behaviour.

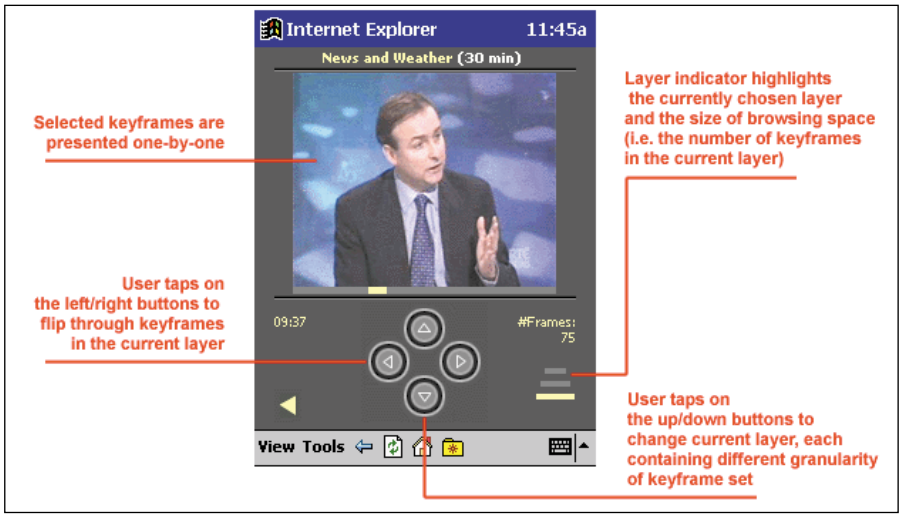

Browsing interface for small-screen, button-operated PDA platforms.

Applications of Físchlár System

Using the system as an experimental testbed, we have been exploring various aspects of video services and possible applications including Físchlár-PTV, a personalised recommender system which uses case-based reasoning and a population of users to recommend TV programmes which users are likely to want to watch. Also under development is work to allow our current users to access the system through wireless LAN using PDAs and mobile phones. A design framework has been developed to help us design various video browsing user-interfaces to support users' differing needs and preferences; we used this framework to develop several distinctive keyframe browsing interfaces currently used in the Físchlár system which are suitable for desktop PCs with relatively large monitors and the use of a mouse.

Design Framework for Video Browsing Interface

The purpose of this design framework is to allow us to design and provide many alternative browsing interfaces to different users for different tasks. This framework enumerates possible design options for different categories. In summary they are:

- Layeredness—consideration about different details or granularity of the browsable keyframe space. Design options in this category are Single layer, Multiple layer with/without navigational link, etc.
- Temporal orientation—what kind of time information is to be provided on the browser. Options in this category are Absolute time (eg timestamp in each keyframe), Relative time (eg timeline bar), etc.
- Spatial vs. temporal presentation—whether to present miniaturised keyframes spatially or one by one temporally on the screen.

From this framework, 'designing a browsing interface' is to simply select a possible option from each category. For example, selecting Single layer / Relative time / Temporal presentation specifies one distinctive browsing interface, currently available on Físchlár's desktop browser as one of the browser options. This method allows us to come up with various interfaces that were not thought of before, for the same underlying Físchlár system.

Video Browsing Interface on a PDA through Wireless LAN

Along with desktop-based video browsers currently deployed to our user population, we are developing mobile applications to allow our users to access the system through a wireless LAN within the campus. While working on technical development to allow keyframes and streamed video data from Físchlár to communicate with mobile devices, we are designing and building suitable browsing interfaces for such devices. We use the design framework to make suitable decisions for small-screen, button-

operated PDA platforms (Psion Revo and Compaq iPAQ), and one of the example designs shown in Figure 1. In this design, we also take into account the particular usage environment of a PDA: using the device on a busy street, in a shaking bus, holding it with 1 or 2 hands. Interfaces for mobile devices should be designed in such a way that the frequency of user interaction could be high but the range of interaction choices should be limited. The design in the Figure features only 4 main buttons to navigate through and within layers with a thumb while holding the device with the same hand. With this kind of design, the interaction becomes similar to that of video game consoles, a style that requires less intensive and constant visual attention at any given moment, careless tapping is acceptable with simple, easily reversible actions. As our wireless LAN becomes more used on campus, our students and staff will use this interface to access available programmes from Físchlár with their own PDAs.

While working on more useful video indexing techniques such as more precise shot/scene detection, object identification and tracking, and teletext-based searching and alerting, we are also considering mobile phone access to Físchlár, heading more towards ubiquitous access to video information.

Links:
http://lorca.compapp.dcu.ie/Video

Please contact:
Hyowon Lee and Alan F. Smeaton —
Dublin City University, Ireland
Tel: +353 1 700 5262
E-mail: {hlee, asmeaton}@compapp.dcu.ie

Designing for Tangible Interaction: The BUILD-IT Project

by Morten Fjeld and Matthias Rauterberg

One of the most challenging research questions in human-computer interaction is: what will be the next generation of user interfaces? How can we interact with computers without a keyboard, monitor, and mouse? To find possible answers to these questions, scientists at the Swiss Federal Institute of Technology (ETH) in Zurich and at the Technical University Eindhoven investigate the theoretical grounding, the practical design, and usability aspects of Tangible User Interfaces (TUIs). As a result, the BUILD-IT Project has been developing a planning tool, based on computer-vision technology, to support complex planning and composition tasks.

In the year 1995, inspired by the work of Pierre Wellner in the context of the augmented reality design paradigm (Interacting with paper on the DigitalDesk. Communications of the ACM. 36(7), pp. 86-96), a research line for further investigations of TUIs was established at the Swiss Federal Institute of Technology (ETH) in Zurich.

To compare the advantages and disadvantages of a computer vision based TUI a field study was carried out. During five days of the largest computer fair in Switzerland, four different interaction styles were realized and presented for public use: (1) a command language interface, (2) a direct manipulative mouse interface, (3) a touch screen interface, and (4) our new computer-vision based TUI interface. The aim of the field study was to establish which was the best interaction style. The statistical analysis of all empirically recorded user actions showed clearly that the most appropriate interaction style for complex problem solving tasks is the TUI. We conclude that TUI's are promising candidates for the next generation of interaction styles.

Encouraged by these results, we started a project with a TUI called the BUILD-IT system. This is a planning tool based on computer-vision technology to support complex planning and composition tasks. The system enables users, grouped around a table, to interact in a virtual scene, using physical bricks to select and manipulate virtual models. A plan view of the scene is projected onto the table. A perspective view of the scene, called side view, is projected on the wall (see Figure 1). The plan view contains a storage space with originals, allowing users to create new models and to activate tools (eg, navigation and height tools). Model selection is done by putting a representative 'brick' at the model position (see Figure 2). Once selected, models can be positioned, rotated and

Figure 1: A design team using BUILD-IT.

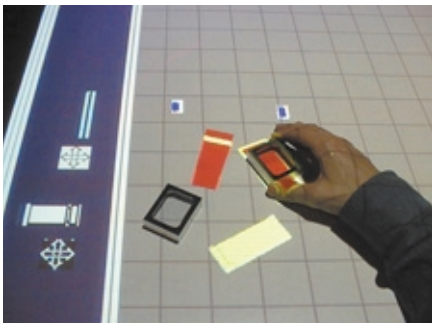

Figure 2: Model selection with a 'brick'.

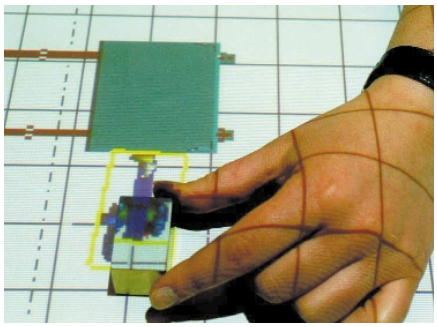

Figure 3: Multiple bricks for two-handed interaction.

fixed by simple brick manipulation (see Figure 3).

Our design practice is grounded on a work-psychological tradition called activity theory. This theory is based on the concept of tools mediating between subjects and objects. Our design philosophy and practice is exemplified by the realization of the BUILD-IT system. Guided by task analysis, a set of specific tools for different 3D planning and configuration tasks was implemented as part of the actual system. We investigate both physical and virtual tools. These tools allow users to adjust model height, viewpoint, and scale of the virtual setting. Finally, our design practice is summarised in a set of design guidelines.

Using the BUILD-IT system as a research platform for tangible interaction, our exploration takes the following path: We first introduce some of the problems related to working in physical and virtual environments, then indicate a few guidelines to achieve what we call natural interaction. Then we give more details about the interaction content, which are configuration and planning tasks. We then come up with new implementations for three-dimensional (3D) navigation. As a particular use of hand-held tools, we introduce alternative ways to control model height in the BUILD-IT system. We also discuss the outcome of our design activity and suggest ways to advance the issues presented before. One continuation of this research line is established with the Visual-Interaction-Platform (VIP) at the Technical University Eindhoven.

Acknowledgements
The BUILD-IT project was financially supported by the Institute for Construction and Design, the Institute of Hygiene and Applied Physiology, the Institute for Work Psychology of the ETH Zurich, and by the Swiss Commission for Technology and Innovation (KTI project) together with several Swiss and German Companies.

Links:
Man Machine Interaction at ETH Zurich:
http://www.iha.bepr.ethz.ch/pages/forschung/MMI/
BUILD-IT Project: http://www.fjeld.ch/hci/
BUILD-IT Publications:
http://www.fjeld.ch/pub/
VIP Project:
http://www.ipo.tue.nl/projects/vip/
HCI Publications:
http://www.ipo.tue.nl/homepages/mrauterb/

Please contact:
Morten Fjeld — ETH Zurich
Tel: +41 1 632 39 83,
E-mail: fjeld@iha.bepr.ethz.ch

Matthias Rauterberg — IPO-Center for User System Interaction, The Netherlands
Tel: +31 40 247 5215
E-mail: g.w.m.rauterberg@tue.nl

The W3C Web Accessability Iinitiative User Agent Accessibility Guidelines 1.0

by Ian Jacobs and Daniel Dardailler

The User Agent Accessibility Guidelines (UAAG) 1.0 is the third of a trilogy of accessibility guidelines published by the Web Accessibility Initiative (WAI) of the World Wide Web Consortium. These documents were designed to present a consistent model for Web accessibility in which responsibilities for addressing the needs of users with disabilities are shared (and distributed among) authors, software developers, and specification writers.

Ms. Laitinen is an accountant at an insurance company that uses Web-based formats over a corporate intranet. She is blind. She uses a screen reader in conjunction with a graphical desktop browser and a speech synthesizer. She uses speech output, combined with navigation of the important links on a page, to scan documents rapidly for important information, and has become accustomed to listening to speech output at a speed that her co-workers cannot understand at all.

For Ms. Laitinen it is critical that her desktop browser communicate with available assistive technologies (screen reader, speech synthesizer), so UAAG 1.0 includes requirements related to communication (through programming interfaces) and to the implementation of system conventions (which increase the likelihood of interoperability). Communication with her assistive technology does not suffice to make her browser more accessible, however. Some of her other needs include:

- the ability to operate it through the keyboard (since a mouse is almost useless to her)
- the ability to move focus to links and form controls (so that her assistive technologies know 'where she is looking')
- access to descriptions of images and video (since these text descriptions can be read by her speech synthesizer)
- documentation on the accessibility features of her browser.

'User Agent Accessibility Guidelines 1.0'explains what software developers can do to improve the accessibility of mainstream browsers and multimedia players so that people with hearing, cognitive, physical, and visual disabilities will have improved access to the World Wide Web. UAAG 1.0 is developed by the W3C User Agent Accessibility Guidelines Working Group, whose participants include software developers, users with disabilities, and international experts in the field of accessibility technologies.

UAAG 1.0 is the third of a trilogy of accessibility guidelines published by the Web Accessibility Initiative (WAI) of the World Wide Web Consortium.

The other two documents are:
- 'The Web Content Accessibility Guidelines (WCAG) 1.0', which explains to authors how to create accessible Web content.
- 'The Authoring Tool Accessibility Guidelines (ATAG) 1.0', which explains to developers how to design authoring tools that are accessible to authors with disabilities, and that produce accessible Web content (ie, content that conforms to WCAG 1.0).

Details on the Guidelines

UAAG 1.0 explains the responsibilities of user agents in meeting the needs of users with disabilities. A user agent that conforms to these guidelines will enable access through its own user interface and through other internal facilities, including its ability to communicate with other technologies (especially assistive technologies such as Ms. Laitinen's screen reader).

Checkpoints and Guidelines

The heart of UAAG 1.0 consists of nearly ninety 'checkpoints', each of which includes one or more requirements. The checkpoints are ranked according to their importance to accessibility (priority 1 for most important, then priority 2 and 3). Here is one example of a checkpoint:
- Checkpoint 1.1 Ensure that the user can operate the user agent fully through keyboard input alone. [Priority 1]

Checkpoints are organized into twelve 'guidelines'. Each guideline expresses a general principle of accessible design. Here is one example of a guideline:
- Guideline 4: Ensure user control of rendering. Ultimately, the user must have final control over the rendering and behaviour of both the content and the user interface. While it is important to preserve the author's intentions as much as possible, without final user control, some content may not be usable at all by some users with disabilities.

Techniques

A user agent may satisfy the requirements of UAAG 1.0 in many different ways. The checkpoints of UAAG 1.0 have therefore been written to be independent of specific markup languages (eg, the Hypertext Markup Language (HTML) or Scalable Vector Graphics (SVG)) and operating systems. To assist developers in understanding how to satisfy the requirements for specific technologies and operating systems, the User Agent Accessibility Guidelines Working Group has published a separate document entitled 'Techniques for User Agent Accessibility Guidelines 1.0' (http://www.w3.org/TR/UAAG10-TECHS/). This document includes references to other accessibility resources (such as platform-specific software accessibility guidelines), examples, and suggestions for approaches that may be part of satisfying the requirements of UAAG 1.0.

The 'User Agent Accessibility Guidelines 1.0' are currently at the 'last call' stage of the W3C Recommendation track. This means that the UAWG has requested technical review from the Web community, and intends to request advancement to Candidate Recommen-dation once review comments have been processed. The UAWG anticipates that UAAG 1.0 will become a W3C Recommendation in late 2001. At that time, we look forward to it being an important resource for promoting Web accessibility.

Links:
World Wide Web Consortium:
http://www.w3.org/
Web Accessibility Initiative:
http://www.w3.org/WAI/
W3C User Agent Accessibility Guidelines Working Group:
http://www.w3.org/WAI/UA/
Web Content Accessibility Guidelines 1.0:
http://www.w3.org/TR/1999/WAI-WEBCONTENT-19990505/
Authoring Tool Accessibility Guidelines 1.0':
http://www.w3.org/TR/2000/REC-ATAG10-20000203/
User Agent Accessibility Guidelines 1.0:
http://www.w3.org/TR/UAAG10/

Please contact:
Ian Jacobs — W3C
E-mail: ij@w3.org

Daniel Dardailler — W3C
Tel:+33 4 9238 7983
E-mail: dd@w3.org

Supporting Usability through Software Architecture

by Len Bass and Bonnie E. John

Some facets of usability such as cancellation, undo, progress bars, and others require software architecture support. Because they reach so deeply into the architecture of a system, these facets must be built into the system from its inception rather than added after an initial system design and user interface has been achieved.

Achieving usability for a computer system depends on a number of factors. The functions provided by the system must accomplish the intended tasks, fit within the work context of the user, be understandable and clearly visible through the user interface. The input and output devices and layout must be correct for the target user group(s) of the system and their physical work environment. The software architecture must be designed to support the user in the achievement of the necessary tasks. In recent work undertaken at Carnegie Mellon University, we have focussed on determining the facets of usability that require software architectural support.

Software architecture comprises the earliest design decisions for a system and, consequently, they underlie subsequent decisions and are the most difficult to modify. In previous work, the focus was on supporting the iterative design process by allowing the presentation and input/output devices to vary. This was accomplished by separating the

presentation from the remainder of the application. Separation is an architectural technique that restricts the impact of modifications. Separating the presentation thus makes modification of the user interface relatively easy after user testing reveals problems with the interface. This approach has been quite effective, and is commonly used in practice, but it has problems. First, there are some aspects of usability that are not affected by separation and, secondly, the later changes are made to the system, the more expensive they are to achieve. Forcing usability to be achieved through modification means that time and budget pressures will cut off iterations on the user interface and result in a system that is not as usable as it could possibly be.

Our approach is more proactive. We identified a set of specific usability scenarios that have software architectural implications. This set currently contains 26 scenarios. Although we make no claim, yet, that these scenarios are exhaustive, they seem to cover much of what people typically mean by usability (eg, efficiency, learnability, some aspects of user satisfaction). For each one of these aspects, we identified an architectural pattern that enables the achievement of the scenario. We then organized these scenarios and patterns in a fashion we will discuss after we introduce an example.

Consider cancellation. Users require the capability to cancel an active command for most commands. It does not matter why they wish to cancel, they may have made a slip of the mouse or changed their mind, they may have been exploring the effects of an unknown command but do not want to wait. Once a command has been cancelled, the system should be in the state it was in prior to the issuance of the command. One architectural pattern for achieving cancellation is for the system to maintain a thread of control separate from that executing the command that listens for the user to issue a cancellation command. Furthermore, the component that implements the issued command must save its state prior to beginning work on the command so that this state can be restored if the user does, in fact, cancel the command. The architectural pattern we have produced for the cancellation scenario discusses these and other aspects of supporting cancellation.

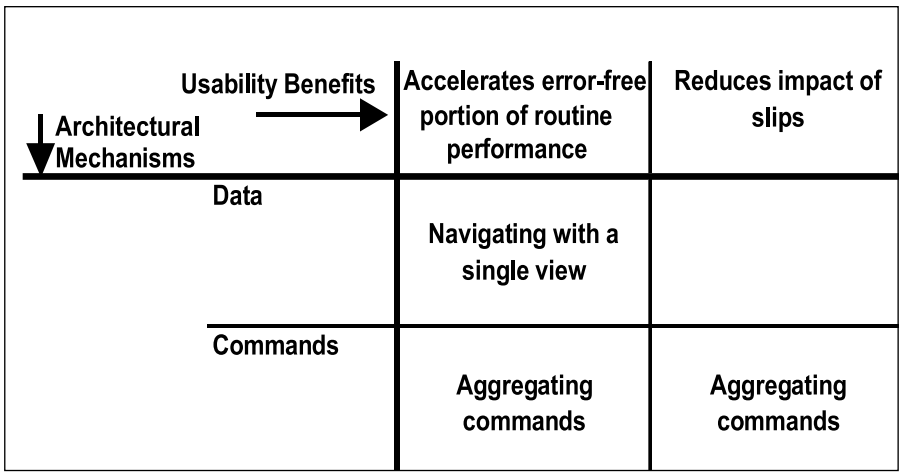

A sample of four cells of the benefit/mechanism hierarchy. The full hierarchy is 9 benefits by 13 mechanisms, or 117 cells.

Other examples are undo, progress bars for feedback to the user, and propagation of known information to input forms. This last scenario would support, for example, requiring the user to input name or address only once and then automatically filling in this information in portions of the user interface that require it.

Once we generated our 26 scenarios, we then organized them in two separate fashions. We organized them into a hierarchy that represents benefits to the organization from applying these scenarios. For example, 'accommodates users' mistakes' and 'supports problem solving' are entries in this hierarchy. Each scenario was placed into one or more positions in the hierarchy. The second organization was into a hierarchy of software architectural mechanisms. 'Recording state' and 'preemptive scheduling' are two of the items in this hierarchy. Again, each scenario was placed into one or more positions in the hierarchy.

The hierarchies were then treated as the two axes of a matrix and each scenario was placed in the appropriate cells. That is, if a scenario provided a single benefit to the users and required a single architectural mechanism to be in place, then it would show up in a single cell in the matrix indexed by its position in the two hierarchies. The figure gives a small portion of this matrix.

For the designer, the matrix can be used in two fashions. If the desired benefits to the organization are known, then the matrix can be used to determine which scenarios are relevant to that benefit and which software architectural mechanisms should be included in the system to achieve those benefits. If one scenario is to be included in the system, then an examination of the matrix will yield other scenarios that could be incorporated into the system at a small incremental cost.

We have used these scenarios in three different design or analysis situations on two commercial and one prototype systems. The systems were a large financial system, a driver information system, and a multi-user meeting support system. That is, we have applied these in a variety of different domains. In each case, the scenarios discovered issues in the user interface design that would have otherwise gone undiscovered, at least through initial implementation.

Readers wishing more information about this work are referred to 'Achieving Usability Through Software Architecture', CMU/SEI-TR-2001-005.

Link:
http://www.sei.cmu.edu/publications/documents/01.reports/01tr005.html

Please contact:
Len Bass — Carnegie Mellon University, USA
Tel: +1 412.268.6763
E-mail: ljb@sei.cmu.edu

Combining Formal and Empirical Evaluation Methods Regarding the Usability of Interactive Systems

by Nico Hamacher

TREVIS is a tool for analysing the usability of interactive systems. This analysis can be done in various ways, eg a formal evaluation using normative user models based on the GOMS theory, and an extensive action log analysis. This tool supports the system engineer to evaluate interactive systems and produces suitable analysis data as a base for decisions while the systems are being developed.

Today, usability engineering of interactive systems is being paid increasing attention. To analyse the usability of interactive systems an empirical evaluation is commonly used. This type of evaluation requires a prototype and a couple of qualified testing subjects. Mostly, this procedure is very expensive and time-consuming. Furthermore, this evaluation is feasible only in late stages of the development process when a prototype is available, so that the results of this evaluation often are not available for redesign. Even though if this empirical evaluation cannot be omitted, because of the generation of plenty of useful information for improving usability, an earlier evaluation would be very helpful for the design and specification of the system.

With normative user models, usability analysis in an early stage of the development process is possible using the GOMS theory, first introduced by Card, Moran and Newell. GOMS is an abbreviation for the components of the model: Goals, Operators, Methods and Selection rules. Over the years it has been shown that GOMS is able to sufficiently describe the interactions of an operator. Its simple and plain structure makes this method easy to understand particularly for development engineers who do not have a psychological background.

An analysis of the user models generates qualitative as well as quantitative predictions, eg execution and learning time. The execution time describes the time to reach the goal whereas the learning time specifies how much time an operator needs to learn the whole task.

System Architecture

TREVIS combines the different usability analysis so that this tool can come into operation in various stages during the development process. The tool includes the following main modules: the user model editor, the device model converter and the analysis module.

Based on the task sequences as one result of the requirement analysis, the user models can be created manually in the user model editor. The tool supports this process, eg, by offering a graphical editor and a library for reusing components. Moreover, user models can be stored in projects, where the project represents the interactive system and the user models describe the tasks which have to be performed.

In contrast to other GOMS-editors which only allow to edit the user models manually, TREVIS is able to generate user models semiautomatically. A device model contains details about the inner work of the device. Using the device model converter the task sequences can be generated semiautomatically. The user models can also be created from these sequences.

In the analysis module the following four different analysis methods are included,

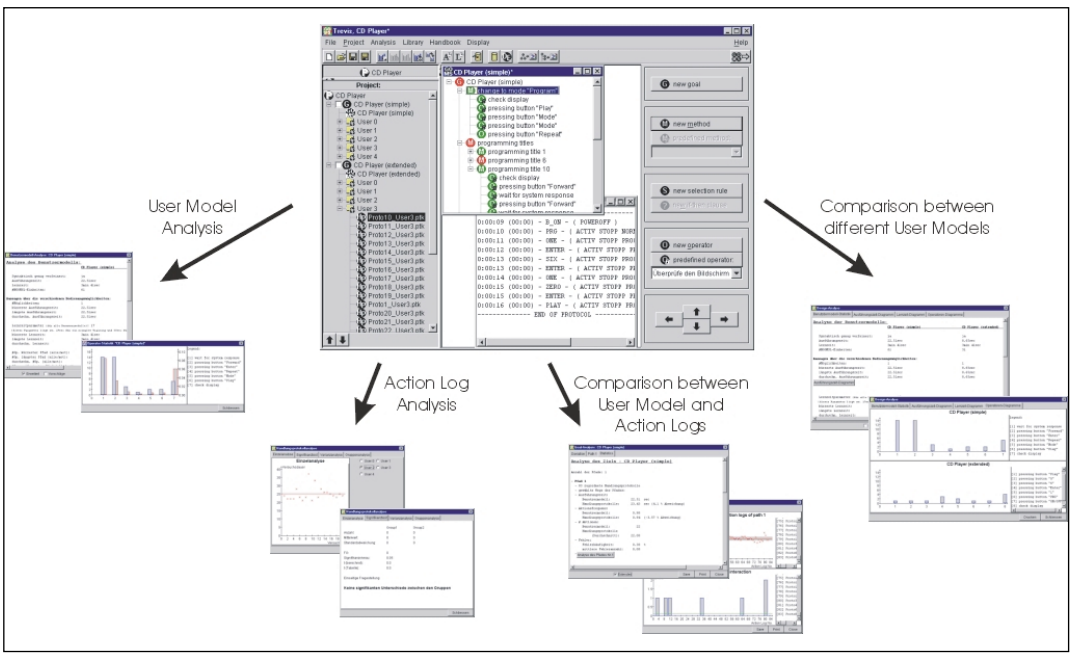

The graphical user interface of the tool TREVIS with the four main analysis methods.

which depend on the development phase in which TREVIS will be used:

- The user model analysis generates qualitative as well as quantitative predictions, eg execution and learning time (as already described with NGOMSL).
- A comparison between different user models is implemented in the design analysis module, which can be used as a basis for design decisions. Although this comparison presentation is a helpful functionality, no other tool includes it.
- In the action sequence analysis, action sequences resulting from testing a prototype can be imported and analysed. A grouping of different sequences is possible, eg to perform an analysis of significance. With this feature, TREVIS is also applicable in late stages of a development process.
- A fourth method analyses these action sequences in comparison with the user models. This analysis shows the differences between the actions specified in the user models and the activities the users performed while interacting with a prototype. Based on this analysis, predictions about effectiveness and efficiency can be made.

Some results of the different analysis methods are shown in the figure.

With the possibility to analyse both user models and action sequences, TREVIS is the first tool that can be utilised in early as well as in late phases of the development process. Beside this, a comparison of user models with action sequences has never been implemented before.

Future Work

The tool has reached a first demo status and was tested with a couple of smaller projects. Now, a large-scale evaluation of TREVIS is planned in order to show that the tool can successfully be applied in a real system development process. Furthermore an integration of some interfaces to other popular development tools, eg prototyping and specification tools, is scheduled.

Link:
http://www.techinfo.rwth-aachen.de/Forschung/MMI/Trevis/index.html

Please contact:
Nico Hamacher — Aachen University of Technology, Germany
Tel: +49 241 4402 6105
E-mail: hamacher@techinfo.rwth-aachen.de

Evaluation in Interactive Safety-Critical Contexts

by Fabio Paternò and Carmen Santoro

Interactive safety-critical applications have specific requirements that cannot be completely captured by traditional evaluation techniques. At the Human-Computer Interaction group of CNUCE-CNR scientists have developed a method that performs a systematic inspection-based analysis, aimed at improving both usability and safety aspects of an application, by analysing a system prototype and the related task model. The goal is to evaluate what could happen when interactions and behaviours occur differently from what has been assumed by the system design.

Research on model-based design and evaluation of interactive applications aims at identifying models to support design, development, and evaluation of interactive applications. In particular, task models describe activities that have to be performed in order to achieve a user's goals, where a goal is a desired modification of the state of an application.

Various approaches have been studied to specify tasks. In our work we consider task models that have been represented using the ConcurTaskTrees notation and the associated tool freely available at http://giove.cnuce.cnr.it/ctte.html. In ConcurTaskTrees the activities are described hierarchically, at different abstraction levels, and represented graphically in a tree-like format, using a rich set of operators to describe different temporal relationships (concurrency, interruption, disabling, iteration, option and so on).

Task models can also be useful in supporting design and evaluation of interactive safety-critical applications. The main goal of such applications is to control a real-world entity, fulfilling a number of requirements while avoiding that the entity reaches hazardous states. Many examples of safety-critical systems exist in real life (air traffic control, railway systems, industrial control systems) and a number of crucial issues arise. One example is when it is impossible to undo a user action; in this case, how to appropriately design the user interface to cope with user errors acquires special importance.

The Method

In our work, we consider the system task model: how the design of the system assumes that tasks should be performed. The goal is to identify the possible deviations from this plan. A set of predefined classes of deviations are identified by guidewords, which are words or phrases referring to a specific type of abnormal system behaviour. Interpreting the guidewords in relation to a task allows the analyst to systematically generate ways the task could potentially deviate from the expected behaviour, analysing the impact on the current system and generating suggestions on how to improve the current design. The method consists of three steps:

- Development of the task model of the application considered; in order to identify how the system design requires that tasks are performed.
- Analysis of deviations related to the basic tasks; the basic tasks are the leaves in the hierarchical task model,

tasks that the designer deems should be considered as units.
- Analysis of deviations in high-level tasks; these tasks allow the designer to identify groups of tasks and consequently to analyse deviations that involve more than one basic task.

It is important that the analysis of deviations be carried out by interdisciplinary groups where such deviations are considered from different perspectives. The analysis follows a bottom-up approach (first basic tasks, and then high-level tasks are considered) allowing designers initially to focus on concrete aspects and then to widen the analysis to more logical steps.

The user interface of the evaluated prototype.

We investigated the deviations associated with the following guidewords:
- None, the unit of analysis has not been performed or it has been performed but without any result. This is decomposed into three types of deviation: lack of input, missing task performance, missing result.
- Other than, the tasks considered have been performed differently from the intentions specified in the task model. Three sub-cases can be distinguished (less, more or different) and each can refer to the analysis of the input, performance or result of a task.
- Ill-timed, the tasks considered have been performed at the wrong time: we distinguish between early or late performance with respect to the planned activity.

For each task analysed the following information can be stored in a table: Task; Guideword; Explanation; Causes; Consequences; Protection; Recommendation. The explanation is classified in terms of which phase of the interaction cycle (according to Norman's model of intention, action, execution, perception, interpretation and evaluation) can generate the problem.

For example, consider the Check deviation task (the user checks whether aircraft are following the assigned path in the airport) and the class of deviation None. For example, the No input case (lack of information about the current state of the aircraft) can have multiple causes (controller distracted, system fault) but the same consequence: the controller does not have an updated view of the air traffic. The No performance case can occur for example when controllers do not interpret correctly the displayed information on the traffic, whereas in No output case the controllers find a deviation but they forget about it as they are immediately interrupted by another activity.

For high-level tasks, the interpretation of each guideword has to be properly customised depending on the temporal relationships between its subtasks (a paper with details is available on request).

Experience

The method described was tested in a real case study in the European project MEFISTO (http://giove.cnuce.cnr.it/mefisto.html). CNUCE-CNR was the coordinator of this project which has seven other partners: University of York, University of Toulouse, University of Siena, Alenia Marconi Systems, DERA, CENA and ENAV (the Italian association of air traffic controllers).

We applied the method to a prototype provided by Alenia Marconi Systems for air traffic control in an aerodrome. The main purpose was to support data-link communications handling aircraft movement, where data link is a technology allowing asynchronous exchanges of digital data coded according to a predefined syntax.

Controllers exchange messages with pilots by interacting with graphical user interfaces (see figure) showing in real-time the traffic within the airport and using the so-called enriched flight labels, which only display essential flight information permanently (standard mode) and show additional information interactively (selected mode).

We carried out the exercise with a multidisciplinary team in which software developers, final users (air traffic controllers in this case) and experts in user interface design were involved. During these exercises, many interesting issues arose that generated a list of suggestions for improving the user interface in order to better support usability and safety.

Our experience has shown the effectiveness of the method despite some social constraints that often occur in software development enterprises (developers tend to defend every decision taken, users tend to digress, pressure of time, etc).

Link:
http://giove.cnuce.cnr.it/

Please contact:
Fabio Paternò or Carmen Santoro — CNUCE-CNR
Tel: +39 050 315 3066
E-mail: {F.Paterno, C.Santoro}@cnuce.cnr.it

Conducting Audio and Video Quality Assessment based on HCI Evaluation Principles

by Angela Sasse

The Computer Science Department at University College London is well-known as a leader in computer networking research. It is therefore not surprising that a significant part of the HCI research activity in our department focuses on design and usability issues surrounding communications technologies. Since 1990, we have conducted user trials and assessments of applications such as videoconferencing. The need to provide meaningful assessment of such emerging technologies has led to the development of new evaluation techniques.

With the increasing popularity of Internet multimedia and moves towards introducing charging, more application designers and service providers are faced with making decisions about audio and video levels required for a particular application. Currently, most designers and service providers resort to one set of 'standard' instruments — the ITU MOS (Mean Opinion Score) scales — to determine whether a certain level of quality is good enough. However, these instruments were developed for traditional image and sound quality assessment, where the question is whether users can detect a particular type of degradation in short clips when focussing on spotting degradations. The essence of a user-centred approach, however, is the recognition that users use systems to achieve a wide range of goals, and that the tasks they complete in pursuit of goals have very different requirements. The low framerate which may spoil a users' viewing pleasure when watching a feature film may be just right for a multi-way videoconference where the collaborative task is centred around a drawing on the whiteboard, and a high framerate of head-on-shoulders videos is distracting. Also, aiming for 'perfect' audio and video may put such services out of reach of many individual and collective users, eg in education and healthcare. The key question is what quality is good enough for users to complete a particular task without adverse effect.

There is no standard set of HCI methodologies for conducting this type of assessment. Over the past few years, we have developed a methodology for conducting audio and video quality assessment based on HCI evaluation principles. A meaningful assessment of usability acknowledges that any interaction involves a user and a system working together on task in a particular physical and/or situational context to achieve a goal. Any particular characteristics of these may generate specific requirements or preferences regarding the level and type of audio and video quality, and have to be replicated in the assessment situation as far as possible. An experiment in a quiet lab situation may not be good predictor of the audio quality required in an environment with high background noise.

Evaluation has to consider at least 3 criteria — task performance, user satisfaction, and user cost. The criterion of user cost had been somewhat neglected in HCI evaluation over the past 2 decades. Audio and video quality assessment has, to date, largely relied to subjective measures of user satisfaction. Subjective assessment — user opinion—is cognitively mediated. Users may not always be able to detect or attribute the effects of audio and video quality correctly. As part of a BT/EPSRC funded PhD studentship, Gillian Wilson has been investigating ways of determining the impact of quality levels on users' physiology.

We need to obtain objective measures of user cost to determine a particular level and quality might be harmful to users; we monitor physiological signals that are indicative of stress and discomfort. When a user is presented with insufficient audio and video quality in a multimedia conference, he/she has to expend extra effort at the perceptual level. If they struggle to decode the information, this should induce a response of discomfort/stress, even if the user remains capable of performing his/her main task. We measure Heart Rate (HR), Blood Volume Pulse (BVP) and Galvanic Skin Response (GSR). We also administer subjective assessments of user cost, ie scales of discomfort, in order to aid interpretation of the physiological results.

Four main studies have been performed to date. The first study investigated the subjective and physiological impact of high and low video frame rates. Twenty-four participants watched two recorded interviews. Results showed that there was a significant increase in stress responses at 5fps from 25fps. Interestingly, only 16% of participants noticed that the frame rate had changed. We investigated 6 types of audio degradations: 20% and 5% packet loss, too loud and too quiet, 'bad' microphone and echo. The most stressful condition, bad mike, was not subjectively rated as poor. The three most stressful and subjectively poor conditions, the network caused only one: 20% packet loss. We showed that problems caused by hardware and user behavior affect users as much as network problems. The results to date are promising, and we are continuing with a study of repeated exposure to different levels of quality while watching entertainment movies streamed over controlled network conditions.

Links:
http://www.cs.ucl.ac.uk/staff/A.Sasse/

Please contact:
Angela Sasse — University College London, United Kingdom
Tel: +44 20 7679 7212
E-mail: a.sasse@cs.ucl.ac.uk

Usability on a Multiple Terminal Media Portal

by Antti Tammela, Sari Lehtola and Katja Rentto

Complex, multiple service portals are very popular entry points into the internet. To attain more and more users new features have been included to the portals. They combine both audio and video material into the news, chat, community, etc. The IMU system developed at VTT integrates related articles, video clips and media events. Video articles, mostly TV news, and text-based articles are treated and presented in a similar way.

The goal is to integrate content for a multiplicity of distribution channels and receiving terminals. Through PC, TV, WAP and MP3 terminals, the trial users could access the same news content at work, at home in the living room and on the move. One purpose of our research was to examine how the features and functions on PC service works on TV service. A TV Set Top Box was built from a computer to study how these media portal services could be used in a living room with a tv remote control.

Figure 1: The IMU publication, where the channels (common and personal) are presented on the left. The channel 'TV-news' is selected. A News banner with the top headlines is displayed on the right.

Figure 2: TV-IMU is used with a TV remote control. The menu 'Common channels' is selected with the remote control red button. The channel 'Kotimaa' is selected with the arrow keys.

Personalisation of News, and Event and Media Calendar

IMU user interface offers a set of channels (Figure 1). One channel type are common channels to which the reader can subscribe (Top News, the weather, etc). The coverage of these channels cannot be changed by the user. In addition to these fixed channels, the user can tailor so-called personal channels according to his or her interests. The reader can, for example, make a sports channel by combining the sports sections from one TV station and one newspaper.

The PC and TV interfaces are very similar (Figure 2). The same information can be retrieved via the TV and the PC. Of course, the smaller resolution of the television screen, affects the way in which the elements are shown. For example, the news stories cannot be scrolled as on the PC, and have therefore to be paginated. The remote control device has to be taken into account, when designing the TV interface. Most of the commands should be given with the buttons on the remote control.

Description of Users and the Usability Evaluation

Ten different user groups tried IMU. Out of these 410 potential users, 325 actually used the system. There were also users who tried more than one terminal. Most of the users were males. There were only 50 females who enrolled on trial. Potential users were mostly quite young. 324 of all were under 35 years old. There were also quite many students (160). Most of the potential users were also very experienced computer users. 217 used Internet more than 10 hours per week.

Easiness of Use

PC and WAP systems were considered to be easy to use, TV navigation on the other hand was considered to be difficult to use. The problem on TV was mainly navigating and commanding with remote control that was felt to be difficult and unfamiliar.

Based on the test results, the TV service as such was too difficult for novice users. The features must be cut down, or alternatively the service must be structured to a different user levels like level for novice users and level for expert users. But then again, such a result was to be expected—one meaning of the research was to find out how the same functionality works in TV and in PC.

User Acceptance

PC-IMU service was considered to be reliable and fast. Some users thought that TV-news didn't appear fast enough. The delay between the actual TV broadcast and the time at which the program was available in IMU service was considered to be too long. Overall test users evaluated IMU to be a relatively useful system. Test users were interested in using IMU if such a system exists. As a competitors to IMU, users mentioned newspapers and other Internet based services. This could partly explain the fact that test users were not very willing to pay for such a service. Test users also hoped that IMU would be updated faster and there were other kinds of content, for example foreign sources.

Link:
IMU: http://www.vtt.fi/imu2/tutustu.html

Please contact:
Antti Tammela — VTT
Information Technology
Tel: +358 3 3163 349
E-mail: antti.tammela@vtt.fi

Enabling Computer Access for Motion-impaired Users

by Simeon Keates, P John Clarkson and Peter Robinson

Conventional computer interfaces and input systems present serious difficulties to users with impaired motion. The Engineering Design Centre of the University of Cambridge, in conjunction with the Computer Laboratory, has been researching methods of enabling computer access for Motion-Impaired Users (CAMIU) for a number of years, and the research is currently funded by both the UK Engineering and Physical Sciences Research Council and Consignia plc.

The primary aim of this research is to enable the design of accessible computer input systems and interfaces for all motion-impaired users.

Current practices in the design of interfaces (both input hardware and software) are often based on user models and descriptions derived almost exclusively from studies of able-bodied users. However, such users are only one point on a wide and varied scale of physical capabilities. Users with a number of different physical impairment conditions have the same desire to use computers as able-bodied people, but cannot cope with most current computer access systems. Such conditions include Cerebral Palsy, Muscular Dystrophy and spinal injuries or disorder. Frequent symptoms include tremor, spasm, restricted movement, and reduced strength. The general population is also growing older. By 2020, almost half the adult population in the United Kingdom will be over 50, with the over 80's being the most rapidly growing sector. Associated with the ageing process are decreases in physical capability, such as reduced hearing and vision. As computer usage spreads throughout the population, the profile of the user base that must be accommodated is changing.

The issue of how to enable computer access for motion-impaired users is being addressed by a number of key research activities. Central to this research is the idea that an interface needs to attain user acceptability, both socially and practically. To achieve practical acceptability the interface must offer satisfactory functionality (utility), usability and accessibility. Social acceptability requires the user to want to use the system and not feel stigmatised.

Characterising the Users

The first objective of the research is to obtain an understanding of the needs and capabilities of users with motion impairments. Research is being performed to study the impact of each of the principal motion impairment symptoms on computer interaction for both disabled and elderly users. The knowledge gained from this work provides insight to the ways in which computer access may be improved. User modelling techniques, such as the Model Human Processor, are being applied and have helped identify fundamental differences between how able-bodied and motion-impaired users interact with computers. New models and methods of characterising user performance are being developed and calibrated, particularly for cursor control and on-screen selection tasks.

Evaluating Novel Interaction Techniques

Research is being performed to evaluate the acceptability of a number of input systems that have been suggested to offer benefits for motion-impaired users.

For example, gestural input is a popular research area for computer access for motion-impaired users owing to their inherent flexibility and ability to be tailored to particular users. The theoretical justification for using gestural input with motion-impaired users is that the freeform nature of gestures should allow users to select movements that are comfortable and repeatable to make. Gestures are also considered to be natural and intuitive as they constitute a component of everyday interpersonal communication. However, user trials showed that motion-impaired users found gestures to be slow and both physically and cognitively demanding to produce.

Haptic feedback is another input option under investigation. Force feedback, for example implemented as a gravity well enclosing a target, can have a greatly beneficial effect, improving interaction times by up to an order of magnitude under certain conditions. However, haptic feedback must be implemented appropriately. For instance, certain types of vibrotactile feedback such as adding vibration to a target, have been shown to have a detrimental effect on the interaction.

Developing a New User-centred Design Approach

To ensure that users who do not conform to the able-bodied stereotype are not excluded from the use of computers, it is necessary to develop and adopt a methodological design approach for implementing inclusive design of computer interfaces.

There are several existing approaches for designing more inclusive interfaces. However, there are shortcomings of each of these approaches that prevent each of them from being used to provide the definitive design approach that designers can use in all circumstances. The principal weaknesses stem from the targeted nature of the approaches. The existing design approaches are often targeted at specific population groups or impairment types. Alternatively, they may be targeted at specific cultures.

The prescribed ways of application of the existing methods are often vague. For example, Universal Design, with a few notable exceptions, is more of an ethos than a rigorous, systematic design approach. There are very few structured descriptions of the implementation of Universal Design in more detail than broad design objectives. Consequently,

while in combination the existing approaches may offer complete coverage of the population needs, individually they do not. Therefore, there is a need for a new approach that draws on the strengths of the existing inclusive design approaches and offers practical and measurable design criteria.

A new design model, the 7-level approach, has been developed as part of the CAMIU project. The approach has been developed to emphasise both aspects of acceptability, the need for practical and social acceptance. The 7-level approach has been successfully applied to a number of case studies, including the design of interfaces for software products, robotics applications and public access kiosks.

Link:
http://rehab-www.eng.cam.ac.uk

Please contact:
Simeon Keates — University of Cambridge
United Kingdom
Tel: +44 1223 766 958
E-mail: lsk12@eng.cam.ac.uk

Visual Dysfunction and Human-Computer Interaction

by Julie A. Jacko

As part of a larger research agenda aimed at empowering all citizens access to electronic information, a research project at the Georgia Institute of Technology, USA, is focusing on empirically linking clinical diagnoses of visual dysfunction with human-computer interaction. The project is funded through 2004.

Funded by awards from the National Science Foundation (Presidential Early Career Award for Scientists and Engineers) and the Intel Corporation, the project is motivated by the knowledge that there is a critical need for all citizens to be empowered to access information electronically. Visual impairment remains a major impediment to electronic access of information. Moreover, in the United States alone, one in every six Americans by the age of 45 years will develop some type of uncorrectable visual impairment.

The research project investigates existing software technologies (ie, Windows and Windows-type platforms) by isolating specific features of computer interfaces like visual icon design, the use of background and foreground colors, and the construction of menu systems. Elements of the interaction such as a person's brain activity in the visual cortex and eye movement patterns on the computer screen are measured. Researchers have shown that visually impaired computer users perform visual search more slowly than their fully sighted counterparts. However, little is known about intermediate stages of visual search that exist between stimuli presentation and stimuli detection/identification. Therefore, the primary focus of the former research segment is to investigate two intermediate stages of visual search in visually impaired computer users: preattention and focal attention. This was accomplished through use of physiological measurements using electroencephalogram (EEG). Through the use of EEG, it has been evidenced that the additional time required by a visually impaired computer user to complete visual search is because of the extra time required for active search, once the visual cortex has already been engaged. Thus, a person's visual limitations are concentrated on the second stage of the process, focal attention.

Subsequent research, centered on focal attention, has involved investigations of eye movements. These investigations are ongoing and involve a remote-mounted infrared video eye gaze tracking system that is used in concert with software that enables isolating specific interaction scenarios during use of a graphical user interface. The eye gaze control unit and software record values of x and y coordinates for the participant's point-of-gaze, at a rate of 60 per second. Such data enable characterizations of specific performance metrics such as fixations and saccades, both indicative of the strategies utilized during visual search and processing (see the figure for a photo taken during experimentation).

The advances that have already been made enable the establishment of a research basis for this field of inquiry, offering well-grounded, empirical findings that, in some cases, support intuition and, until today, ad hoc solutions. This has been accomplished by establishing groundwork for inquiries aimed at understanding how basic human-computer interaction is linked to the visual capabilities of the low vision user. The research is uniquely collaborative, involving collaboration with an ophthalmologist (Dr. Ingrid U. Scott) at the Bascom Palmer Eye Institute (BPEI) of the University of Miami School of Medicine, USA, in order to couple

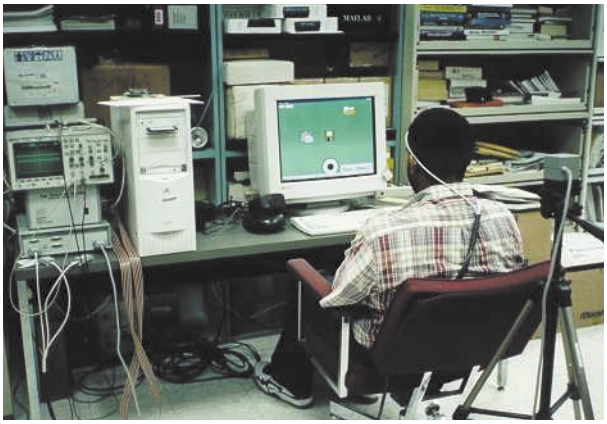

Experimental paradigm involves tracking a human subject's eye movements (research led by J.A. Jacko, Ph.D., Principal Investigator).

ophthalmologic expertise with her expertise in human-computer interaction. This collaboration also enables the engagement of people with impaired vision from the BPEI Low Vision Clinic in the research. Additional collaborators in this research include Dr. Armando B. Barreto of the Department of Electrical and Computer Engineering of Florida International University (FIU), USA. Dr. Barreto is Director of the Digital Signal Processing Laboratory at FIU.

Conclusion

To achieve universal access to electronic information technologies, designs must overcome barriers that have been perpetuated by traditional 'one-size-fits-all' philosophies in order to accommodate the disabled, the elderly and technologically unsophisticated individuals. In order for information technologies to be universally accessible, there must be a paradigm shift in human-computer interaction (HCI) that shifts the burden of interpreting behavior from the human to the computer. An ever-growing population of users who will benefit tremendously from this paradigm shift is those with impaired vision.

To facilitate shifting the burden of interpreting behavior from the human to the computer, the notion of adaptive interfaces has emerged. Ongoing and future work of this research team involve developing methodologies and tools necessary to implement adaptive, multimodal human-computer interfaces that are personalized for individual users representing a full spectrum of visual capabilities.

Link:
http://www.isye.gatech.edu/lhci/

Please contact:
Julie A. Jacko — Georgia Institute of Technology, USA
Tel: +1 404 894 2342
E-mail: jacko@isye.gatech.edu

Navigating Information Spaces

by David Benyon

New HCI challenges arise from the emergence of information spaces and the related concept of navigation. Research at Napier University focuses on this topic.

Evaluation is a fundamental part of human-computer interaction (HCI). Good HCI practice tells designers to evaluate: evaluate requirements, evaluate designs, evaluate prototypes. The purpose of evaluation is to improve the usability of a software system; that is to make it easy to use, easy to learn, effective and enjoyable. But what is usability and what makes one device easier to use than another? Traditional HCI theory has produced a number of evaluation techniques and guidelines. These are based on some basic psychological assumptions which date back to the sixties.

Broadly speaking it is assumed that people have goals that they are trying to achieve and the reason that they are using a software system is to enable them to achieve those goals. People have to translate their goals into actions and then evaluate the response of the system to determine if their goals have been achieved. People have, therefore, to bridge the 'gulf of execution' (translating goals into actions) and the 'gulf of evaluation' (perceiving and interpreting what has happened). A practical method of evaluation, the Cognitive Walkthrough method, is based on these ideas. There are other approaches to evaluation which involve coming up with a set of good design 'heuristics'. Don Norman spurned the cognitive psychology and Jacob Nielsen gave HCI the most famous set of heuristics. Norman & Nielsen now travel the world advising on usability.

These ideas of usability and evaluation developed during the hay-day of HCI, the decade from the mid eighties. These were the days before the Web, when people went to work and sat in front of computers. These were the days before everyone used Microsoft Office, when programmers developed applications specifically for their users. The twenty-first century is not going to be like that (at least not the first part!) and it is a moot point whether the concepts and methods of usability developed during the eighties will transfer. For example, the design of Web sites, virtual environments, distributed computing environments and novel settings such as households all throw up usability challenges that traditional methods seem ill-equipped to deal with. Navigation in Web spaces and other hypermedia spaces is a key feature, yet counts for only one of Nielsen's heuristics. 'Know your user' is another heuristic; yet how can you know the people who visit a web site? Navigation is even more of an issue in virtual environments, where there are many more degrees of freedom and where users more literally move to where they want to go. Ideas of tasks and goals do not transfer easily to environments such as the household where entertainment and relaxation are key, and style and fashion are significant factors. With distributed computing environments people are literally surrounded by devices that are communicating autonomously with each other. People will be wearing computers and exerting their identity. These large open 'information spaces' have quite different usability requirements and hence require different evaluation mechanisms. One very important aspect of these systems is that we need to consider accessibility in addition to usability.

The traditional view of HCI sees the person outside the computer, looking onto a world of information. An alternative is to see people as inside their information space. This in turn leads to the idea that people are navigating through this space. An information space is made up from information artefacts which are the devices we use to help us undertake activities in the 'real world'.

The concept of an information 'space' leads naturally to consider ideas of navigation. In the real world we navigate quite happily. We also know a lot about how to design spaces such as museums, galleries, and so on to assist navigation and to make navigation a pleasurable experience. Research looking at how to take knowledge of signage, maps and so on to evaluate and inform the design of information spaces is on-going.

More than that we recognise that in the real world we do not just follow signs to get around — we ask people, follow people, make recommendations to our friends. We are also looking at issues in 'social navigation'. Direct social navigation through discussions and indirect social navigation through awareness of others and of history.

One approach to dealing with the difficulties of navigation in information spaces is to provide more personalised help and assistance through the use of intelligent agents and recommender systems to help navigation. Other ways of helping people find their way is through signs, maps and landmarks. The new environments that we have to design for demand new approaches to the ideas of usability and of how systems are evaluated. The 'navigation of information space' view of humans and computers in the twenty-first century provides a new perspective on usability: design for people to live in their information spaces.

Please contact:
David Benyon — Napier University
Edinburgh, United Kingdom
Tel: +44 131 455 4294
E-mail: d.benyon@napier.ac.uk

Information Access and Refinement — a Research Theme

by Fredrik Olsson, Preben Hansen, Kristofer Franzen and Jussi Karlgren

In order to develop systems and methods for improved information access for different users in different situations, a new research theme at SICS focuses on combining knowledge of human behaviour, language technology and linguistic principles.

Texts are complex and underspecified structures that leave room for interpretation on the part of the reader. This may seem undesirable, and it is arguably more efficient in terms of information transmission to structure information properly at the time of production, eg by employing controlled language strategies or database solutions of various kinds. This approach not only removes some of the major drawbacks of text but also removes some of the major assets of text as an information transmission and storage device. It is all but impossible to predict what information needs and structures the text will eventually be used to fulfil; it is presumptious to impose the author's structures on all future readers; the portability of the information from one informational context to another will suffer. Text is crucially vague and indeterminate, and it is worth keeping that way.

Instead, we need systems that deliver tailored information just-in-time. Our research theme puts great emphasis on tailoring information access, which in turn necessitates understanding information needs.

Systems for accessing information employ a rather simplistic attitude towards notions such as text, user, information need, seeking process and relevance. Most of the research in the field today concentrates on tuning algorithms to standard problems. But each one of the above central concepts deserves more attention. Our research theme works on methods for the next generation of information access systems—which we believe will be less general and more crucially bound to usage and domain.

Information Access and Refinement—SICS Approach

Information access and refinement is a new research theme at the Swedish Institute of Computer Science (SICS). Information refinement is a growing field, especially as regards information extraction and summarisation technology. Likewise, the study of interaction and usage of information access systems has gained great research interest. There is some research done on finding supra-sentential structure in text. What makes our research theme unique is that we consider both the beginning and the end of the chain, and that we strive for a model that is primarily verifiable through study of human behaviour and based on language technology and linguistic principles. Engineering and system construction aspects are in second place. This is our way of avoiding local minima: linguistics may appear impractical and behavioural science may appear detached. But the conjunction of the two will, in the long run, provide us with solutions to practical problems.

Central Concepts to the Theme

Information access is about providing people with different tools and methods for granting reliable and simple access to the information they need, ideally with awareness of task and context of the access situation. What ties together the notions of user, information need, and contents of a text is relevance. A text can be said to be relevant if it satisfies some user's information need in some situation.

By information refinement we mean the process in which a text is further processed to find and compile the pieces of content that are relevant from a certain perspective. Example techniques are information extraction, information retrieval, automatic summarisation, and generation of reports, all of which are techniques which consider the contents of a text from the point of view of a user with a particular information need.

Steps to be taken

The commercial systems of today assume that texts are simply bags of words, that all users are alike, that the need for information is static, and that dialogues in information seeking systems are simplistic one-shot transactions where the user is happy to exchange a few words for several thousands of possibly relevant links. Commercial systems also assume that relevance is not depending on users, situations, perspective, temporal aspects or other extratextual factors. This is a useful starting point from which to consider improvements.

To be able to refine information and to improve today's systems, we need to consider at least three things:

- We need to gain and use information about the reader, in particular about how the reader understands his information need
- We need to gain and use knowledge about the text that the reader makes use of in order to satisfy the information need
- We need to examine how the reader meets the text, and how the presentation of information can be related to the information need.

It is only by taking into consideration these three items simultaneously and integrating parts of them that we will be able to come up with something interesting.

We are currently involved in several projects, national as well as international, in which the intersection of our ideas presented in this text is tested. We aim at developing an open architecture for information refinement as well as defining methods for identifying ways of improving information access to fit the needs of different users with specialised needs of information expressed in different situations.

Link:
http://www.sics.se/humle/iatema

Please contact:
Fredrik Olsson —SICS
Tel: +46 8 633 15 32
E-mail: fredriko@sics.se

WebPR — Adaptive Information in the Physical City

by Emmanuel Frécon, Anneli Avatare and Adrian Bullock

The WebPR application promotes awareness and cooperation between physical and digital users by embedding digital information in public spaces and making it accessible to people who are present in those spaces.

The ICE-lab at SICS is currently developing a range of illustrative electronic landscapes that combine the physical and the digital in order to realize new forms of interactive experiences. One of these is the WebPR application, an application that aims to bridge the gap between the digital world of the Web and the physical world of real life organisations, activities and people. WebPR builds on the work of WebPath (see 'WebPath—a 3-D Browsing History Visualisation' by Emmanuel Frécon in ERCIM News 41) and visualises and personifies Web traffic for the Web server of an organisation at some central point, eg a reception area. Our aim is that WebPR will bring back an embryo of collaboration between digital and physical inhabitants and a more human-oriented vision to the digital age of Web servers.

We assume that in the future there will be universal access to online information and a universal digital information society: the future Internet environment. One of the most ubiquitous electronic landscapes in use today is the World Wide Web. One vision of the future sees society living in a totally digital environment, 'jacking in' to alternate realities as in the film The Matrix and the writings of Gibson and Stephenson, for example. However, we suggest that even given the sophisticated digital worlds of the future, there will always be the need for the physical, both physical contact and physical settings, and that meetings and visits will long be an inherent part of a business. We want to be able to offer a channel through which the digital and the physical will be able to exchange enough information to keep contact and bring back a human-oriented vision to the digital age of Web servers.

In particular we focus on Web servers as our digital environments and reception areas and foyers as our physical environments. With the popularity of the Web increasing exponentially, Web presence has become more and more important and all organisation types are concerned, from small and medium businesses, through big conglomerates, to government agencies. The Internet is here to stay and the amount of information that can be found on Web servers increases at a pace similar to the increase of new connections. For organisations, more people connected means that the information presented has more potential and will reach out to the masses. For business-oriented organisations, this means an increasing number of customers. For public organisations, this means the possibility to better communicate with the citizens, inform them on the latest decisions and let them know what taxes are used for. In short, one of the main reasons for having an attractive Web server is to improve public relations in general.

However, the exponential popularity increase of the Web introduces an awareness gap between virtual visitors and employees. The day-to-day contact that used to be the building block of our working habits is disappearing in favour of anonymity. WebPR was developed in order to bridge this gap. Employees and visitors at the physical setting of an organisation will use WebPR to get a sense of the amount of (virtual) customers and connected persons, together with their interests and actions. This will be

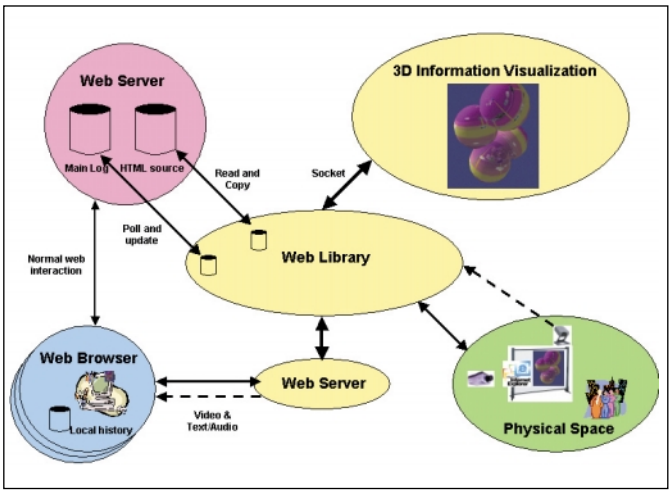

Figure 1: The WebPR architecture.

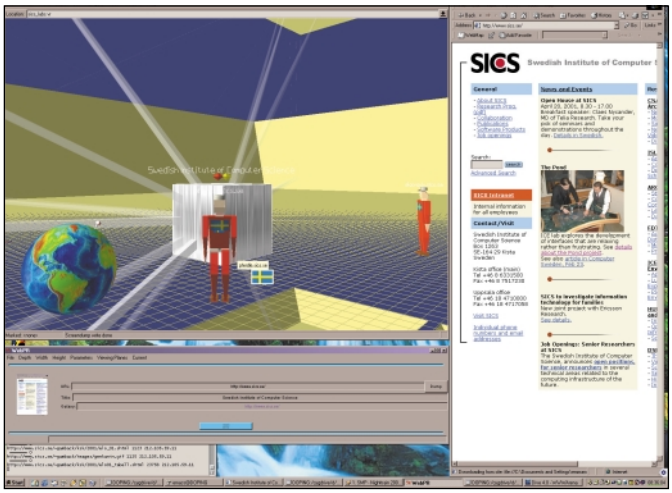

Figure 2: The WebPR application.

achieved through visualising and personifying Web traffic for the Web server of the organisation at some central places of its physical setting, eg foyers, receptions, etc.

At the heart of WebPR is the Web Library, a collection of programs written in TCL/TK, which govern and control all the interactions and interplay that takes place in the installation (see Figure 1). The other main component of the application is the information visualisation.

Figure 2 shows the WebPR application. It consists of a visualisation window, a web browser, and an interface to the Web Library. The main aim of the visualisation is to be sufficiently interesting and engaging to the passer by. It supports some control from occasional employees or visitors and in the future will offer an automatic mode (like a screensaver). It has to be appealing enough to be run discontinuously every day. We believe that the quality of the visualisation lies not only in the aesthetic of its appearance but also in the amount of information that it synthesises.

Initial internal trials have been made with the demonstrator, with mixed success, and we are currently reviewing the visualisation approach and method.

Links:
http://www.sics.se/ice/projects/infocity

Please contact:
Emmanuel Frécon — SICS
Tel: +46 8 633 15 34
E-mail: emmanuel@sics.se

Expert Contextual Online Help Strategies

by Noëlle Carbonell and Antonio Capobianco

Online help to the use of standard application software intended for the general public is still highly unsatisfactory. Despite the continuous efforts of researchers and designers over the past twenty years, most novice users still prefer consulting experienced users to browsing available online help or paper user guides. Specific training sessions and tutorials are still very popular among novices nowadays. A promising approach, based on the analysis of human expertise, is presented (project supported by the French Ministry of Defence and the CNRS).

To be efficient, usable and actually used, online help has to overcome one major obstacle at least, that is the 'motivational paradox'. Carroll and Rosson observed in the eighties that users in the general public are reluctant to explore new software and to learn how to use its functionalities efficiently. These users seem to be mainly concerned with achieving the tasks motivating their interaction with the software. Therefore, they are liable to ignore the facilities for autonomous learning provided by intuitive user interfaces in order to familiarise users with the operation of new software.

This paradox may also explain the failure of approaches which view online help as an interactive learning situation. In particular, it may account for the tendency of the general public to ignore online tutorials as well as online manuals in the form of databases, hypertexts or hypermedia.

To take this paradox into account, designers have to consider online help as a specific human-computer interaction situation. From this angle, contextual online help appears as an appropriate design framework. Providing users with the specific information they need to carry through their current task, and delivering this information at the right moment, contribute significantly to supporting users' interactive activity, thus meeting their expectations.

However, selecting relevant meaningful contextual information and exploiting it appropriately are still crucial research

issues, as illustrated by available implementations: products on the market are too crude to prove useful, and research prototypes too complex to be reliable, hence usable.

In order to improve the utility and usability of online help, we resorted to eliciting the help strategies of human experts from the analysis of expert-novice dialogues. Our study addresses two main issues:

- What type of contextual information do human experts actually use?
- How do they exploit such information for providing novice users with efficient usable interactive help?

Empirical Data and Analysis Method

We analysed expert-novice dialogues collected during an earlier study. Eight novice users performed twenty predefined formatting tasks on a given text, using Microsoft Word. They could communicate freely (over an intercom) with an expert user who helped them to carry out the prescribed tasks. Expert and novice were in different rooms; in one condition (A), the expert could view the novice's screen (via Timbuktu), while in the other one (B) she could not.

Dialogues were tape-recorded, and the subjects' screen displays videotaped. We used time-stamped written transcripts of the dialogues annotated with descriptions of the novices' actions and their results on the user interface.

Our analyses are based on the labelling of all speech acts in the textual corpus, using specific taxonomies which were evolved from a preliminary survey of the dialogues.

Results and Interpretations

Requests for contextual information represent over 14% of the expert's speech acts in condition B (vs less than 2% in condition A), and 87% of these requests aim at clarifying the software current state as displayed on the screen (54%), the progress of the current task execution (33%) or the subject's current intention (10%). Therefore, any help system which aims at emulating human experts' strategies should involve contextual information among its major knowledge sources.

A close analysis of the expert's speech acts shows that she most often uses contextual information for selecting the help information she gives to novices (over 90% of her help speech acts in both conditions). The main source of dynamic contextual information is the progress of the current task execution (over 40% of her help speech acts).

These results indicate that the expert encourages novices to adopt an autonomous 'learning by doing' strategy for familiarising themselves with the operation of a new software, by helping them to achieve the tasks which motivate their use of the software.

Such a strategy, which differs from standard didactic strategies implemented in computer aided instruction, will be well accepted by users in the general public (cf. the motivational paradox). It will also be easy to implement, as it exploits a model of the novice's current activity instead of a cognitive user model which is more complex to build, update and use.

Future Work

We are currently developing a prototype which will be used to evaluate (experimental evaluation with potential users) the usability and efficiency (in terms of learning) of this strategy, compared to other contextual help strategies.

Link:
http://www.loria.fr

Please contact:
Noëlle Carbonell and Antonio Capobianco — LORIA
Tel: +33 3 83 59 20 32, +33 3 83 59 30 83
E-mail: {Noelle.Carbonell, Antonio.Capobianco}@loria.fr

Scrutability for Personalised Interfaces

by Judy Kay

We are seeing a huge growth in personalisation and customisation of interfaces. This introduces the need for a collection of information about the individual user. Within the adaptive user interface community, this is called a user model. This is any individual, dynamic and explicit model the system holds about the user. It might represent the user's beliefs and preferences as well as other user attributes. The issue of scrutability of user models arises.

With the user model as the driving force for a system's adaptation to the individual user, we can illustrate the issues for scrutability of user models in the following scenario.

Jane starts Mynews, a personalised electronic newspaper. First she gets the headlines. Today, these include:

- Politician caught out by wife
- New movie release: Flipper saves the day
- New music release: Sydney Symphony Orchestra—Carr-Boyd's Prelude.

Mynews is supposed to select just those news items most likely to interest this user. It models Jane's interests so that it can collect reports on issues she will want to know about. Where there are several items about an issue, the user model is supposed to ensure selection of reports Jane will prefer. For example, she may like a particular critic's music reviews or one syndicate's reports for European news. In this scenario, Jane might ask:

- How did it decide I would be interested in a Flipper movie?
- Or a Carr-Boyd release?

- What or who is Carr-Boyd?
- How can I let it know that I am pleased to hear about the Flipper movie?
- Why did it give me that article about the politician?
- And especially at this time, during a national election campaign?
- Why hasn't it reported the new unemployment statistics that were due for release today?

This scenario would probably need a large model of user interests and preferences. The first two questions relate directly to parts of the user model representing Jane's preference for 'Flipper movies' and 'Carr-Boyd music'. The user should be able to find information to help answer the questions if they have access to the value of that part of the model and a record of the processes which contributed to that value. It may be that Jane has rated several Flipper movies and other animal-centred children's movies, usually giving very positive ratings. She may have also ordered several such movies from an on-line service that provides such information to the user model. If such information affected the values in the user model, Jane should be able to scrutinise the model to see that this was so.

The third question might be answered within the articles of the newspaper. Generally, such questions require an explanation of part of the user model ontology, in this case, a music ontology with 'Carr-Boyd' or 'Australian contemporary art-music composer'. The fourth question relates to user control of the model.

The last three questions are likely to involve a combination of the user model and the mechanisms controlling Mynews. These questions relate to issues of media control.

Why might we want scrutability? The strongest reason is that users should be able to feel in control. When the user knows that a screen has been adapted to them, it is natural for them to wonder just what has been adapted to them. What might a different user see? Are they missing out on something? This view is supported by the 1995 European Community Directive on Data Protection (http://www.doc.gov/ecommerce/eudir.htm, visited May, 2000) which states that personal data for customisation requires an 'intelligible form of the data undergoing processing and of any available information as to their source'.

We have been working in this area on several fronts. We have developed the accretion representation for user models. This ensures that the interface can explain what it believes about the user and why in terms of the sources of information and the way these are interpreted. From that foundation, we have been working on interfaces for viewing user models, both in detail as well as with visualisation interfaces which enable the user to gain an overview of a large user model. We also have projects which extend the scrutability beyond the user model and into the application. For example, in the context of interfaces to support learning we have explored a scrutable hypertext interface, and, in other projects, we are exploring a scrutable teaching strategy. In broader applications, we are building experimental systems that assist the user in managing documents with scrutable assistant programs.

Link:
http://www.cs.usyd.edu.au/~judy

Please contact:
Judy Kay — University of Sydney, Australia
Tel: +61 2 9351 4502
E-mail: judy@cs.usyd.edu.au

MKBEEM — Developing Multilingual Knowledge-Based Marketplace

by Alain Leger, Aarno Lehtola and Victor Villagra

The MKBEEM project (Multilingual Knowledge Based European Electronic Marketplace) develops a mediation system which adapts the language and the trading conditions of an Internet sales point according to its international customership.

One of the main objectives of a truly user-friendly Information Society is to focus on advanced human language technologies enabling cost-effective interchange across language and culture and more natural interfaces to digital services. Europe is made up from nationalities with a broad variety of cultures and languages. So, the information society must develop in such a way as to support and nurture these rich diversities, rather than act as a constraint to their evolution through the new media.

MKBEEM demonstrates the mediation system with two trials, namely sale of train travels with hotel and other bookings, and B-to-C mail-order sale of clothes and accessories. Figure 1 illustrates the system architecture, which is implemented as Java EJBs. The technical approach is highly based on combined use of human language processing and ontologies. This combination plays a central role in i) mapping the human language query onto an ontological representation; ii) producing composite services from CP catalogues by the broker; iii) mapping ontological queries onto the ontological representations of the catalogues, and iv) describing the domain knowledge. They are used for classifying and indexing catalogues, for facilitating multilingual man-machine dialogues and for inferring information that is relevant to the user's request.

Multilingual Cataloguing

Cataloguing concerns maintenance personnel of a content provider in two different roles. First, there are cataloguers, whose duty is to add new products or edit/delete pre-existing products in the mediated repertoire. Secondly, there are catalogue editors, whose role is to maintain the product models, attribute sets and value spaces, market specific product

category trees and terminologies. The MKBEEM system implements a multilingual cataloguing scenario for editing, checking and automatic translation of product articles and other product descriptions, for extracting product properties into structured database from textual product articles, for automatically deriving set of possible market-specific categories for a given product, and for processing queries to locate pre-existing product descriptions for maintenance.

The cataloguing uses a human language processing server (HLPS) and a domain ontology server (DOS). HLPS provides language identification, meaning extraction, machine translation, text checking and language generation services. Meaning extraction is a central service that is used for several purposes during cataloguing. It involves executing linguistic analysis on a language input and performing concept matching of the found linguistic constituents with respect to the domain ontology models in order to derive relevant concept bindings. The outcome, the ontologic formula, is represented as a logical clause in the CARIN language and can be used for properties extraction, for finding product categories and for query planning. Figure 2 illustrates the cataloguing UI of a content provider in MKBEEM Pilot 1 system.

Processing Cross-Lingual Information Requests

Cross-lingual information retrieval is the second key service feature of the MKBEEM service offering. MKBEEM makes use of the combined solution covering both query translation and result translation.

To this purpose, we have based our technical solution on the co-operation between Human Language Processing and Knowledge Representation (conceptual Interlingua reasoning) of the domain of discourse, namely the MKBEEM application domain. Ontologies form a knowledge base for MKBEEM system agents to infer information that is relevant to the user's dialogues with the MKBEEM system.

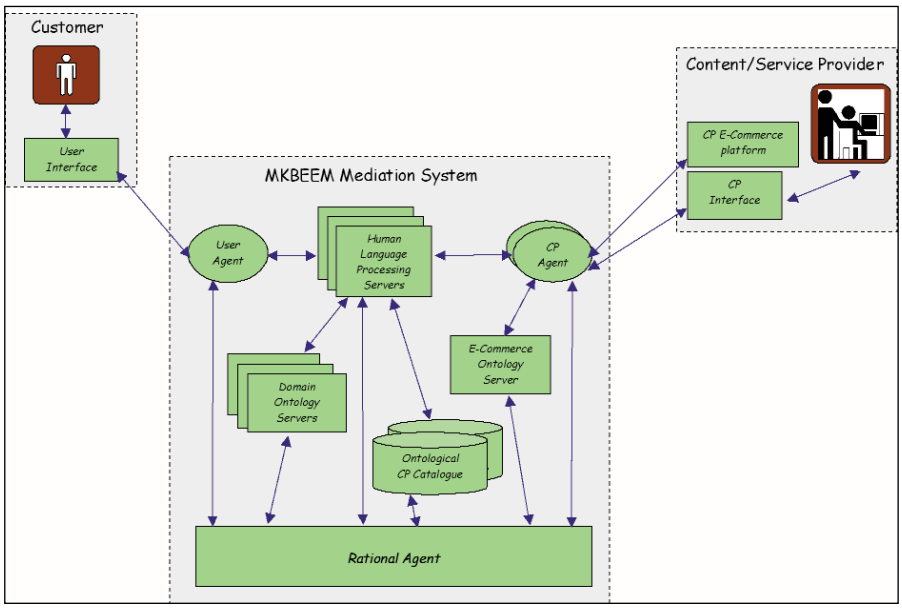

Figure 1: Mkbeem system mediates Internet services to international customers.

From the information retrieval aspect, ontologies improve the accuracy in fuzzy information search. Viewed from a natural language processing perspective, ontologies facilitate mono- and multi-lingual human-computer dialogues by paraphrasing the query of the user through context identification and disambiguation.

Finally, the use of an ontology aims also to solve the cross-language problem as there is not always 1-to-1 correspondence of terms between different cultures. The system should be able to find the same product even if the query is made in a language different from that of the product name / attributes. This possibility should be provided in a way as natural as possible to the user.

Multilingual Trading

The third main functionality foreseen for the MKBEEM system is related to the final step of an e-commerce transaction: the trading. Once a user has chosen his preferred offer provided by the system, he will be able to invoke the trading process in order to acquire the selected good. This implies one or several transactions with the e-commerce platforms of the content providers involved in that offer. As there can be offers that are satisfied by several content providers, the MKBEEM system will act

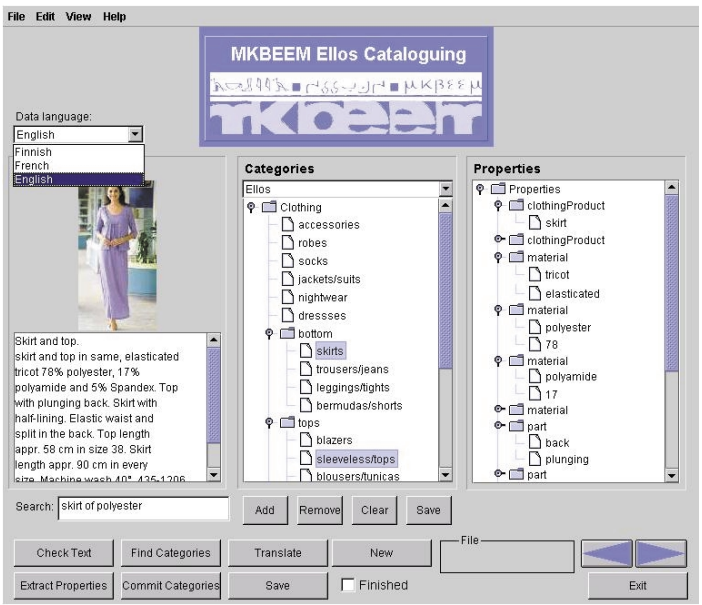

Figure 2: Cataloguing user interface applet of Mkbeem Pilot 1.

as a one-stop-shopping module able to mediate with the content providers. In this way, the MKBEEM system has to offer the following functionalities:

- Harmonise Business Rules. As the trading process implies some contracts to be agreed between the user and each content provider, the MKBEEM will try to harmonise these contractual issues depending on some parameters such as countries of each actor, preferences declared by each actor in their user profiles, etc. An e-commerce ontology will be developed in order to provide the knowledge to this process.
- Multilingual Trading. As the user and the content providers may use different languages, the MKBEEM system will act as a mediator for translating the e-commerce transactions between them.

The MKBEEM project integrates Knowledge Based processing (Knowledge Representation and Reasoning) and Human Language processing in providing multilingual e-commerce mediation services. The consortium aims at proving that the technology concept is robust for given domains, and thereby bringing advances in both technology and services. The consortium consists of the coordinator France Telecom, the user partners Ellos (Finland), SNCF (France) and FIDAL (France), and the R&D partners SEMA (Spain), UPM (Spain), NTUA (Greece), CNRS (France) and VTT.

Link:
Project home page:
http://mkbeem.elibel.tm.fr/

Please contact:
Alain Leger — France Telecom R&D, France
Tel: +33 2 99 12 4223
E-mail: Alain.Leger@rd.francetelecom.fr

Aarno Lehtola — VTT Information Technology
Tel: +358 9 456 6032
E-mail: Aarno.Lehtola@vtt.fi

Victor Villagra — Universidad Politécnica de Madrid, Spain
Tel: +34 9154 95700
E-mail: villagra@dit.upm.es

The A-Book: An Augmented Laboratory Notebook for Biologists

by Wendy E. Mackay and Guillaume Pothier

We all face a growing problem: managing the relationship between physical and on-line documents. A joint effort between Project MErLIn at INRIA Rocquencourt and the Institut Pasteur in Paris is exploring how to effectively integrate the two in the a-book: a computer-augmented laboratory notebook for research biologists.

Research in ubiquitous computing and augmented reality has challenged the view that a 'computer' means a keyboard, mouse and monitor. Instead, physical objects such as paper can serve directly as the interface to the computer. This encourages new forms of human-computer interaction that are more easily integrated into current work environments.

We are exploring this approach with respect to laboratory notebooks, which provide a fertile testbed for exploring how to manage information with both physical and electronic manifestations. Our research strategy involves participatory design with two types of users: biologists and archivists. We began with videotaped observations and interviews to understand their needs and then used video brainstorming and prototyping techniques to explore the design space for augmented laboratory notebooks. We are now developing a general software architecture to support persistent data and integrate paper and on-line documents.

Field Study Results

Our field studies examined the use of notebooks from several user perspectives. Individual biologists use notebooks to officially record their hypotheses, procedures and experimental results. They view them as personal documents that reflect their personalities as well as their current research needs. Laboratory notebooks are multi-media documents: in addition to writing and drawing, biologists paste in photographs, computer printouts with data analyses and even the physical results of experiments, eg, gels. Notebooks may also refer to external objects, such as test tubes stored in refrigerators. Although all are heavy computer users, most biologists appreciated the simplicity and flexibility of their paper notebooks. However, they found some tasks cumbersome, such as creating indexes or searching for a specific experiment, and they rarely summarized their data in the prescribed manner. Several also expressed frustration with finding relevant information in another colleague's notebook.

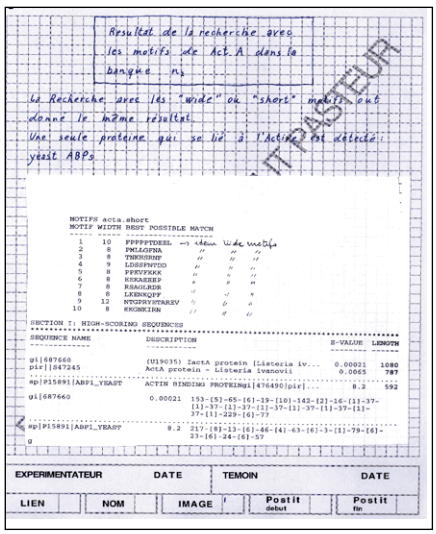

Figure 1: A-book prototype 1 captures handwriting with a CrossPad™. Menu commands are printed directly at the bottom of the page. Computer-generated research results are pasted onto the page, with hand-written comments above. The box indicates that the contents should be added to the on-line index.

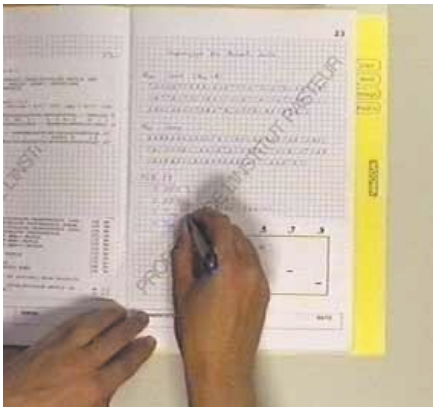

Figure 2: A-book prototype 2 captures hand-writing with a WACOM™ graphics tablet; menu commands appear on a separate sheet to the left.

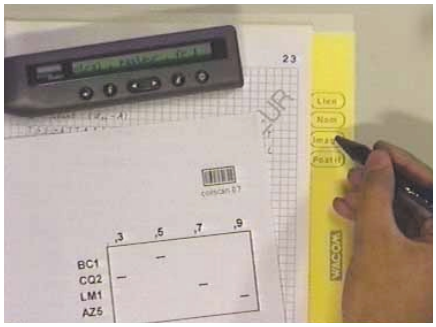

Figure 3: Adding a computer printout: the biologist reads the automatically-generated code with a Pocket Reader™ then 'presses' the image button to simultaneously insert it into the corresponding electronic page of the notebook.

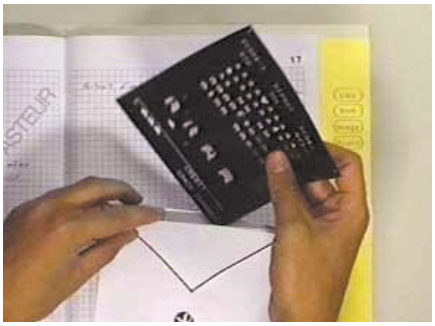

Figure 4: Adding a microscope photograph.

The Institut Pasteur strictly regulates laboratory notebooks, maintaining a staff to archive and manage them after biologists leave. From their perspective, the notebooks comprise the intellectual property of the organization: the archive staff are interested in both their research and historical value. They worry about storing paper-based data, since ink may spread over pages, media may decay, and links between notebooks and physical specimens may be lost. However, they are even more concerned with the growth of on-line data, which quickly disappears when stored in obsolete software formats on out-moded computer media.

A-book design

Our findings influenced the design of an a-book, or augmented laboratory notebook. After consulting with users, we built two working prototypes that test different technical and interaction strategies. Each prototype allows biologists to write on notebook pages as before, with ink on paper. However, the a-book simultaneously captures an electronic copy of each stroke and associates it with the correct page, providing a time-stamped, on-line record of the hand-written text. Biologists can easily specify links to on-line information, such as data results or digital photographs that are pasted into the paper notebook. If text is underlined, the a-book uses character recognition to interpret it. For example, a biologist might specify a particular web address or interpret a sequence of DNA for later analysis. Biologists can also label and categorize information for subsequent search, by drawing a box around a name, procedure, drawing or any other object. The a-book creates an on-line index and content summary and biologists can use either the paper or on-line versions to search for specific information.

Future research

We are developing a third prototype that incorporates user feedback and plan to test it with several volunteers in the summer of 2001. The new version will use a document-based software architecture that handles persistent data over long periods of time. Data, both physical and electronic, is represented as separable layers of information, each with different characteristics. For example, hand-writing, whether in ink or captured as on-line gestures, is unchangeable. However, additional layers may be added to provide annotations, explanations, or interpretations of the original text. The a-book challenges traditional views of software design and, by integrating multi-media physical and electronic information, offers a new vision for human-computer interaction.

Please contact:
Wendy Mackay — INRIA
Tel: +33 1 3963 5912
E-mail: wendy.mackay@inria.fr

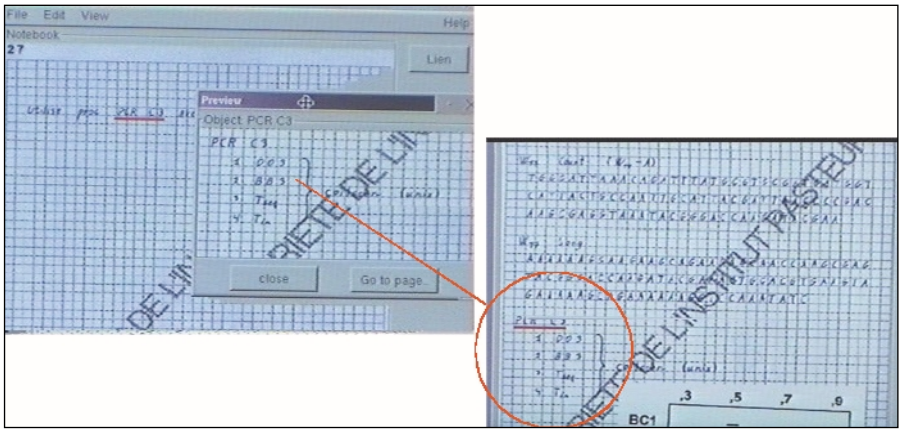

Figure 5: Following a link: Clicking on the underlined text (the procedure 'PCR 3') leads the biologist to the page where the procedure was originally defined.

Visual Knowledge Management with Adaptable Document Maps

by Andreas Becks and Matthias Jarke

Structuring and condensing corporate document collections is an important aspect of knowledge management. At the GMD Institute for Applied Information Technology (FIT), scientists have developed the corpus analysis tool DocMINER, which provides interactive visual access to text collections. Successful industrial applications include the critiquing of technical document collections such as user case descriptions in software engineering, or user manuals of complex engineering systems.

Since a great deal of corporate knowledge is contained in textual documents, techniques which provide analysts with task-adequate access to text collections play an essential role. Of particular importance is an effective support for explorative corpus analysis tasks, where the user is concerned with discovering patterns in the document space and getting an overview of available documents and their semantic relationships. This is an application area where information visualisation promises to be helpful: document maps visualise the overall similarity structure of a corpus of texts, using a suitable metaphor reminiscent of geographical or astronomical cartography.

In the DocMINER project (conducted at RWTH Aachen and GMD-FIT from 1998 to 2001), we have designed and evaluated a document map system for visually aiding text corpus analysis tasks in knowledge management. DocMINER differs from earlier document mapping efforts in that it is based on a careful analysis of why and how users actually employ technical document collections in their work. The resulting domain-specific task model not only serves as a guide for the technological development but also as a yardstick for evaluation.

DocMINER supports an adaptable framework for generating a graphical corpus overview, using a modular combination of algorithms from classical information retrieval, spatial scaling, and self-organising neural networks. The basic method allows a fine-grained (dis-) similarity analysis of specialised document collections. It can be tailored to domain-specific needs, since the module for assessing the similarity of documents is exchangeable. A semantic refinement extension based on fuzzy rules enables the analyst to incorporate a personal 'bias' into the map generation process. The system DocMINER — an interactive, map-centred corpus analysis and text-mining tool — tightly integrates the graphical display with explorative and goal-directed interaction methods. Its interface design was guided by Schneiderman's 'Visual Information Seeking Mantra': overview first, zoom and filter, then details-on-demand. System features include different zoom, scaling and sub-map functions, the means to define and assign document symbols, an annotation function, automatic map labelling and document group summaries, and a tight coupling with a query-driven retrieval interface.

So far, three application areas have been looked at via case studies in science and industry. Firstly, the UML-based large-scale software standardisation effort of a consortium of worldwide chemical industries employed DocMINER to analyse the consistency of collaboratively written user cases stored in GMD-FIT's BSCW internet workspace environment, and to assess the relevance of each scenario to the design of the different architectural components. Secondly, DocMINER was used by a software house supporting the steel industry in quality assurance for technical product documentation, i.e. checking the consistency of the topical structure of user manuals and defining single information sources. Thirdly, in a more scientific environment, DocMINER served as a forum for the discussion of relationships between sub-projects and terminology uses in Germany's Cultural Science Research Centre, 'Media and Cultural Communication', in Cologne. Each of these applications provides strong anecdotal experience that document maps improve the way in which the respective types of work are traditionally done. This was also confirmed by feedback from leading industrial technical documentation fairs.

To investigate how relevant to these successes is the document map visualisation itself, a controlled laboratory study with students and technical documentation experts was conducted, comparing the map interface with the usual title or abstract lists provided by search engines and similar IR tools, while

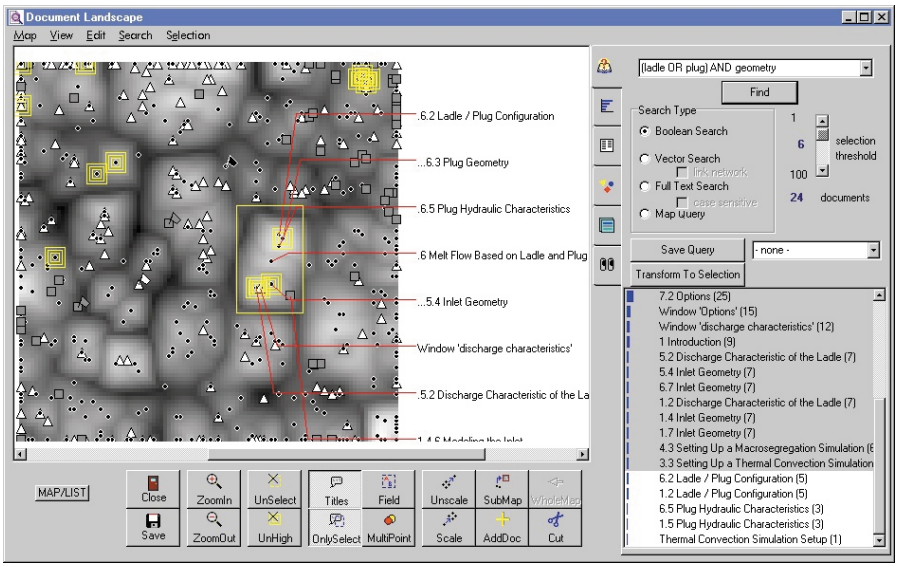

A snapshot of the DocMINER user interface.

leaving all the other features of DocMINER in place. The results clearly confirmed the task-adequacy of document maps: the computation of the overall similarity structure of the text corpus and its visualisation helps to significantly improve the effectiveness of typical task solutions. Furthermore, test subjects subjectively preferred the document map system in nearly all cases.

Summing up, from the viewpoint of target users the document map approach offers meaningful insights into a collection's structure, and allows one to effectively study relationships between single documents and document groups. It is particularly successful for supporting tasks that require a detailed structural analysis of document-document, document-topic or document-specification relationships.

With its text analysis features, DocMINER complements the range of visually oriented data exploration and information brokering tools developed in GMD-FIT's research department on Information Contextualisation (ICON), including the InfoZoom visual querying and navigation environment for relational databases, and the Broker's Lounge environment for the creation of context-specific brokering systems, both of which now enjoy a fair number of commercial applications. Further work will link these solutions with each other and with techniques for metadata management in the data warehousing and information flow management that FIT is developing in co-operation with a number of regional SMEs, eg for applications in financial controlling and e-learning.

Links:
http://www.docminer.de/

Please contact:
Andreas Becks, Matthias Jarke — GMD
Tel: +49 241 80 215 14
E-mail: {becks,jarke}@fit.gmd.de

Human-Computer Interaction Issues in FAIRWIS

by Maria Francesca Costabile and Paolo Buono

FAIRWIS (Trade FAIR Web-based Information Services) is an ongoing project at the University of Bari, funded by the European Union. This project aims at offering on-line innovative services to support trade fair business processes and a large number of exhibitors organised in a Web-based virtual fair.

Traditionally, the information media for supporting trade fair events are paper-based: booklets, flyers, maps, etc. are the means used to exchange information. In recent years, some Web-based information sites have been made available, providing information both on trade fair events and on companies participating in these fairs. However, these data are not organised in an integrated, homogeneous and comprehensive way, since they are usually presented in a rigid pre-designed company oriented style. Moreover, currently available Web sites exploit static data which is difficult to update and make available on-line in an appropriate format. FAIRWIS has a real time connection with an underlying database to guarantee coherence of data and up-to-date status.

In FAIRWIS, the whole concept of trade fairs is transferred into electronic form, and visualisation techniques, including virtual reality, are used in order to provide a feeling of 'reality' to the users of trade fair information systems. The figure is an example of a virtual reality visualisation of a fair venue, with all the stands. The project does not aim at substituting, but at enhancing the existing traditional approach of getting people together

Presenting data on the Web in a convincing and understandable way requires a lot of work when data change dynamically; in particular it is difficult to modify the graphical layout without disorienting the users. One of the main objectives of FAIRWIS is to facilitate human-computer interaction (HCI) and to allow easy access to the stored data. Several data visualisations are generated to present the retrieved information in appropriate ways. Thus, specific categories of users to whom FAIRWIS is primarily addressed, namely fair organisers, exhibitors, and professional visitors (people who visit the fair for business reasons and not only for fun), can get valuable help in the different phases of the decision making processes needed to improve their own business.

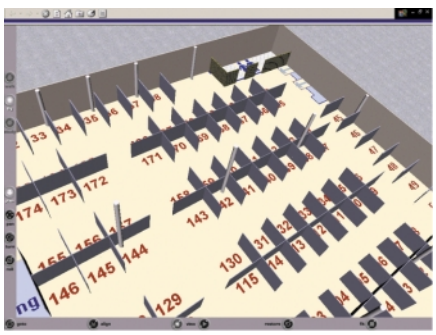

Example of a virtual reality visualization of a trade fair venue.

The FAIRWIS project and the system development have been carried out with a user-centred methodology, in order to build a system that satisfies clear usability objectives. The basic principles of user-centred design are:

- analyse users and tasks
- design and implement the system iteratively through prototypes of increasing complexity
- evaluate design choices and prototypes with users.

A user-centred approach requires understanding reality: who will use the system, where, how, and to do what. The system is designed iterating a design-implementation-evaluation cycle. In this way, it is possible to avoid serious mistakes and to save re-implementation time, since the first design is based on empirical knowledge of user behaviour, needs, and expectations. In accordance with this methodology, together with all partners in the project, we have devoted a lot of effort to collecting user information by observing users in their workplaces and interviewing them. It is

well known that this is not an easy task and is also very time consuming. However, a careful user analysis in the first phase of the system project and user involvement in the evaluation of the system prototypes are the only ways to ensure that the produced system will conform to the users' needs and expectations. FAIRWIS offers a modular solution, whose main modules are:

- FAIRWIS Core, which generates and manages an online dynamic web site
- graphical engine, which provides virtual reality representations
- user profile engine, which provides system personalization
- data analysis engine, which provides on-line marketing analysis on the basis of the stored data.

The figure shows some of the capabilities of the graphical engine. A characteristic worth mentioning is that an updating of the underlying database is automatically shown in the visualised scene.

The user profile engine has been developed to take into account that another important issue for improving HCI, especially when different types of users access the system, is to provide suggestions and indications automatically adapted to the user currently working with the system, in order to help him or her to find information of interest. The personalisation component implemented in FAIRWIS actually works as a recommender system that helps users navigating in the catalogues of the exhibited products.

The aim of the FAIRWIS marketing component is to manage and improve interactive relationships among FAIRWIS users. More specifically, company managers can plan company activity on the basis of historical data available in the database. An objective of the data analysis engine is to serve the needs of specific types of users, primarily fair organisers and exhibitors, by allowing them to easily retrieve information useful for their marketing activities. For this purpose, it exploits appropriate visualisation techniques, in accordance with the results of recent research in a specific field of HCI, that is information visualisation, which has proved that suitable visualisations can reduce the time needed to obtain information, and make sense out of it.

Acknowledgements

The support of the European Commission through grant FAIRWIS IST-1999-12641 is acknowledged. We are also grateful to the work carried out by the other partners of the FAIRWIS project.

Please contact:
Maria Francesca Costabile or Paolo Buono – Dipartimento di Informatica,
Universita' di Bari, Italy
Tel: +39 080 5443300
E-mail: {costabile, buono}@di.uniba.it

Anyone, Anywhere Access to Community-oriented Services

by Pier-Luigi Emiliani, Michael Pieper and Constantine Stephanidis

PALIO, a new project recently funded by the European Commission's Information Society Technologies (IST) Programme, builds on previous European research and technological development efforts to provide a novel understanding of anyone, anywhere access to community-oriented services. This brief article reports on the aims and objectives of the project, as well as its potential contributions to the universal access community in Human Computer Interaction.

The notion of universal access to the Information Society is rooted on the concept of universal design, as it has evolved over the years. Universal design refers to the conscious effort to consider and take account of the widest possible range of end user requirements throughout the development life-cycle of product or service. In recent years, it has been applied in interior and workplace design, housing and landscapes.

In the context of Human Computer Interaction, design for all implies a proactive approach towards products and environments that can be accessible and usable by the broadest possible end-user population, without the need for additional adaptations or specialized (re-)design. Building upon the results of earlier projects, such as ACCESS and AVANTI, we are now pursuing the universal access challenge at another level. Specifically, the PALIO project sets out to address the issue of anyone and anywhere access to community-wide services. This is an extension of previous efforts, as it accommodates a broader perspective on adaptation and covers a wider range of interactive encounters beyond desktop access. In what follows, we will briefly overview how this project addresses the issue of universal access and how it advances the current state of affairs by considering novel types of adaptation based on context and situation awareness.

The PALIO project is funded by the EC's IST Programme. The main challenge of the PALIO project is the creation of an open system for accessing and retrieving information without constraints and limitations (imposed by space, time, access technology, etc.). Therefore, the system should be modular and capable of interoperating with other existing information systems. In this scenario, mobile communication systems will play an essential role, because they enable access to services from anywhere and at anytime. One important aspect of the PALIO system will be the support of a wide range of communication technologies (mobile or wired) to access services. In particular, it will be possible for a user equipped either with a common cellular phone or an advanced WAP phone to access services wherever he/she is. The Augmented Virtual City (AVC) centre will adapt the presentation of

information to the different access technologies (Figure 1).

The PALIO system envisages the adaptation of both the information content and the way in which it is presented to the user, as a function of user characteristics (eg, abilities, needs, requirements, interests); user location with the use of different modalities and granularities of the information contents; context of use; the current status of interaction (and previous history); and finally used technology (eg, communication technology, terminal characteristics, special peripherals).

The PALIO information system consists of the following main elements (see Figure 1):

- A communication platform including all network interfaces to inter-operate with both wired and wireless networks
- The AVC centre which is composed of the main adaptation components, a service control centre, and the communication layers from and to the user terminals and the information services
- Distributed Information Centres in the territory, which provide a set of primary information services.

The AVC centre is the architectural unit that manages diversity and implements the mechanisms for universal access. The AVC will be perceived by users as a system which groups together all information and services that are available in the city. It will serve as an augmented, virtual facilitation point from which different types of information and services can be accessed. The context- and location- awareness, as well as the adaptation capabilities of the AVC, will enable users to experience their interaction with services as a form of 'contextually grounded' dialogue, eg, the system always knows the user's location and can correctly infer what is 'near' the user, without the user having to explicitly provide information to the system.

The envisaged main building blocks of the AVC are depicted in Figure 2, and can be broadly categorised (according to their role within the project) into the adaptation infrastructure; the service control centre; the software communication layer to and from the user terminals; the distributed information services and their integration infrastructure.

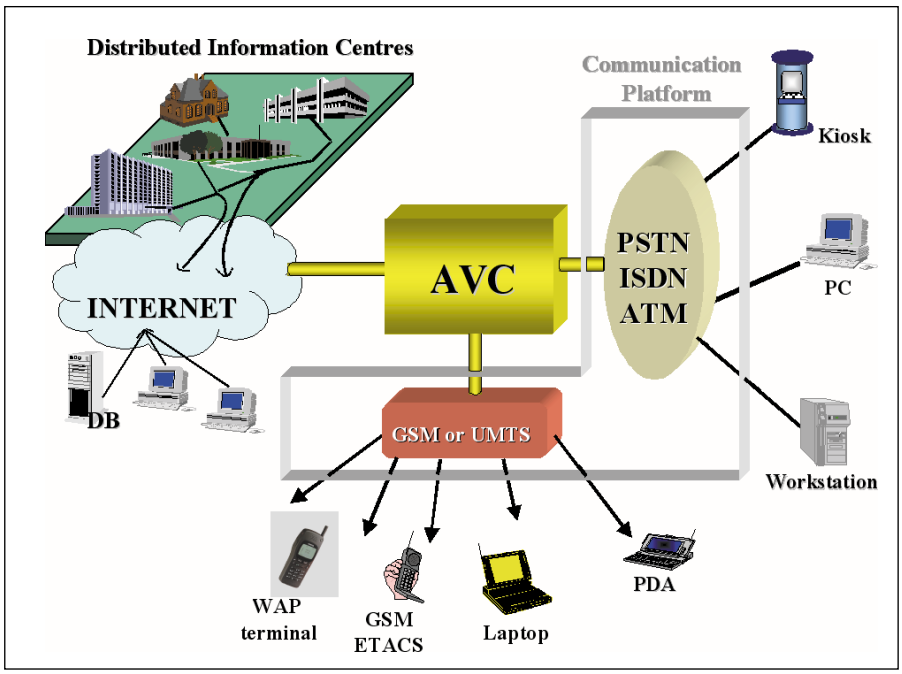

Figure 1: PALIO architecture.

It should be noted that one of the very important properties of the described architecture is the fact that there is continuous monitoring of the interaction between users and the PALIO system, so that modifications in the user behaviour, the context of use, or the status of communication facilities and devices can be identified, interpreted, and used to trigger the appropriate type of adaptations (eg, inclusion of links to relevant information, spatial or temporal restructuring of information elements, modification of presentation parameters such font, colour, voice gender, volume, etc). Central for the design of adaptations based on context-sensitive process-oriented knowledge is the Basic PALIO Tourist Scenario Framework, which comprehensively describes touristic situations, and forms the basis for deriving numerous information needs of tourists, and corresponding service requirements including situation-dependent user interface adaptations and accessible hardware devices. The framework distinguishes between exogenous, intervening and endogenous dimensions of scenario descriptions. The PALIO basic user model conceptualizes tourists' information processing behaviour in

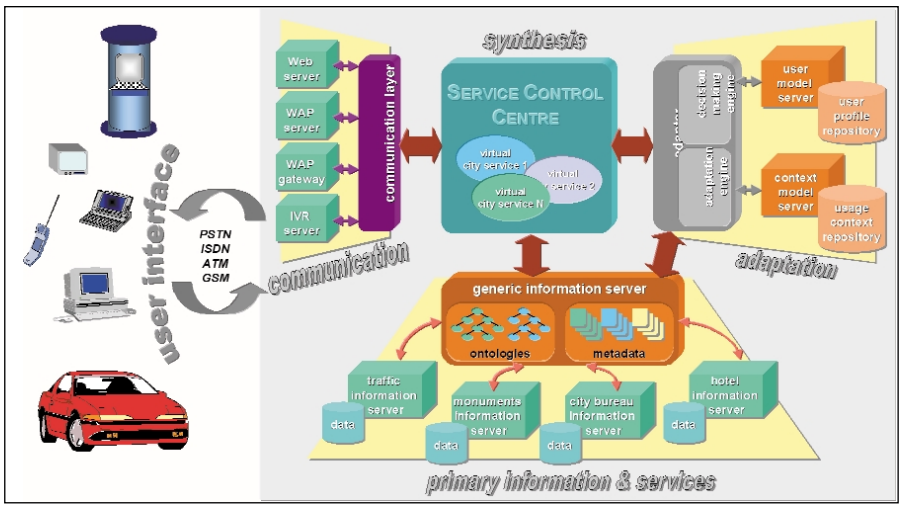

Figure 2: Augmented Virtual City Centre architecture.

terms of behavioural means-end analysis. Rules for interface and content adaptation in PALIO are inferred from the interplay between the Basic PALIO Tourist Scenario Framework and the PALIO basic user model.

Acknowledgements
The partners of the IST-PALIO project are: ASSIOMA S.p.A.(Italy)—Prime Contractor; CNR-IROE (Italy); Comune di Firenze (Italy); FORTH-ICS (Greece); GMD (Germany); Telecom Italia Mobile S.p.A. (Italy); University of Siena (Italy); Comune di Siena (Italy); MA Systems and Control Ltd (UK); FORTHnet (Greece).

Please contact:
Pier-Luigi Emiliani — IROE-CNR
Tel: +39 55 4235 235
E-mail: ple@iroe.fi.cnr.it

Michael Pieper — GMD
Tel: +49 2241 14 2018
E-mail: michael.pieper@gmd.de

Constantine Stephanidis — FORTH -ICS
Tel: +30 81 39 1741
E-mail: cs@ics.forth.gr

Áit Eile, Another World for HCI

by Paula Hicks, Rose Mc Grath, John Mc Keown, Kieran Mc Nally and Jane Grimson

Spending a lifetime with a serious or chronic health condition can be traumatic, stressful and physically and emotionally painful. This is especially true when the individual enduring the condition is a child or adolescent. The psychosocial issues suffered by these children include isolation, change in family dynamics, depression, withdrawal, loss of peer interactions to name but a few. The Áit Eile project proposes to tackle some of these issues by creating an online environment which will provide for communication, social support, relaxation, entertainment and distraction.

Áit Eile (Another World) is a two-year project run by the Centre for Health Informatics, Trinity College, which started in September 2000, and is funded by the Information Society Commission through the Department of Health and Children, and IBM. The main objective of the project is to develop a secure on-line virtual community for children in hospital, which would allow young patients to communicate with each other via email, live chat, and video-conferencing over a secure Internet connection. Through a web-portal, the system will also provide access to activities (eg arts and crafts), educational opportunities and links to favourite web sites pitched at an appropriate level, from hospital, home, or school. The aim is to empower seriously ill children to combat the medical and emotional challenges they face on a daily basis.

HCI and Usability Issues

The main HCI challenges which confront Áit Eile, emerge from the provision of a secure online community environment where children can chat, email and videoconference. Additionally they need to be able to access their local Bulletin Board for information, play games, and take advantage of educational facilities. The primary goal is to create a self-explanatory communication portal, which endeavors to make the children's interactive experience intuitive and satisfying. In designing the user interface, it is also necessary to take bandwidth availability into consideration. Children could be accessing the system via a high-speed ISDN line in the hospital or school, or from their own home via a slow dial-up modem connection,

The project has adopted a relatively standard approach to addressing usability issues. Interviews and focus groups involving the children themselves, the medical and nursing staff in the participating hospitals, teachers, educational psychologists, and play therapists, were used to capture the user requirements. A detailed task analysis was then carried out which resulted in a structured template that acts as a foundation for the user interface design. At this stage, a set of usability benchmarking standards were established against which progress can be measured and tested. In the case of video conferencing for example, the minimum quality of voice and visual data was outlined, with a preference given to the quality of voice over image.

Video conference prototype interface.

The conceptual and visual design phase is an iterative process, which aims to minimize the learning time necessary for the user to navigate the system, to maximize their retention of their newly acquired 'how to' knowledge, and most importantly to enhance their sense of satisfaction with the system. In this way the children will easily immerse themselves in the activities of the community without the risk of foundering on the technology. The visual design phase aims to provide an accessible multisensory environment in which the children have a sense of control. They will be provided with expressive tools with which they can develop their inherent creativity and sociability. A degree of personalization will be built into the portal offering different features to different age groups.

Testing and evaluation will be based on a series of prototypes, refined on the basis of feedback and experience, with each successive prototype offering increasing functionality and a variety of access modes. The first prototype, incorporating login, chat, games and video conferencing, is scheduled for live testing in early June 2001. This will involve children in two of the participating hospitals communicating with each other via the system. The two hospitals involved are already connected via a secure Intranet over ISDN, which significantly simplifies security considerations.

Future Plans

Áit Eile currently involves six hospital schools in Ireland: the National Children's Hospital; Our Lady's Hospital for Sick Children, Crumlin; Cork University Hospital; Temple Street Children's Hospital; the National Rehabilitation Hospital; and Cappagh Orthopaedic Hospital; as well as the Barretstown Gang Camp [www.barretstowngc.ie]. In the coming months the prototype will be extended and all the sites brought on-line. Initially children will access the system from a desktop PC generally located in the hospital school room. Within the two-year time-frame of the project, it is planned to open up access using laptops for children who are confined to bed, and for children at home. This will pose both technical and usability challenges, and, most importantly, will require careful attention to security and monitoring control.

A similar and highly successful project, Starbright (www.starbright.org) has been in operation in the US for a number of years and there is interest in developing a European-wide virtual community. It is hoped therefore to link Áit Eile with these other networks.

Please contact:
Kieran McNally — Trinity College Dublin.
Tel: +353 1 608 2381
E-mail: Kieran.McNally@cs.tcd.ie

Our Approach to Social Computing

by Kristina Höök

How can we empower people to find, choose between, and make use of the multitude of computer-, net-based- and embedded services that surround us? How can we turn human-computer interaction into a more social experience? How can we design for dynamic change of system functionality based on how the systems are used? Tackling these issues is the core of the work in the HUMLE laboratory, and our solutions are inspired by observing that much of the information seeking in everyday life is performed through watching, following, and talking to other people — what we name social computing.

If you enter a room in an unfamiliar environment, and you see that people are sitting on chairs, silent, looking down into their papers, you will do the same, silently awaiting what will happen next. Or imagine walking down a street in your hometown, trying to decide what to do. You notice a crowd outside your favourite café. Knowing that the café often has live music, you can guess that there must be a special event on tonight. You might decide that you're in the mood for a lively evening and join the line, or you might decide that you prefer a quiet night and go look for a different café. Or imagine you're in a library, looking for a book about interface design. One of the books on the shelf is much more worn and well-thumbed than the other, suggesting that lots of people have read it. You may decide it's a better place to start learning than the pristine books beside it on the shelf.

Unfortunately, in most computer applications, we cannot see others, there are no normative behaviours that we can watch and imitate. We walk around in empty spaces that very well might not have been visited by anyone else before us as far as we know. In our word processor, we might be lost for hours with no guidance whatsoever. On the web, there is no-one else around to tell us how to find what we are looking for, or even where the search engine is, there might very well be lots of services that we badly need, but that no-one tells us of, and once we find a service, we do not know how to interact with it, nor whether it can be trusted.

This is why we have been developing the idea of social navigation [A.J. Munro, K Höök, and D.R. Benyon, Social Navigation of Information Space. Springer Verlag, London, 1999]. By showing the traces of other actors or allowing actors to speak to one-another, we believe that they will more easily find what they look for. In the HUMLE laboratory we now want to expand this vision to entail all forms of social interaction — moving beyond the more narrowly defined concept 'navigation' towards interaction in general.

It feels like a truism to point out that people are social beings, that our curiosity

is easily raised by watching what others do, their relationship to one another, and the gossip around these events. Still we design computer interaction as if people were only work-oriented, objective, emotionally void, entities.

The aim of our social computing theme is to make computers to work like our everyday language or the way a city grows. Language has been around for a long time, it is democratic in the sense that anyone can add to it, words can shift in meaning as they are being used, it is a tool that never ceases to work, it is multipurpose, and beautiful. Language is therefore a dynamic, on-going, social process. The way a city grows is similar. Take for example Stockholm, it does not get released in new versions every year, instead it changes with what people do in various parts of the city: building new houses, taking new paths into usage, moving cafés and restaurants. Computer systems in general should have this quality: as they are used, actors change their functionality. Social computing is one way of achieving this. As a system is used, the usage will leave trails in the system that will guide future usage.

In social navigation there is a strong temporal and dynamic aspect. A person chooses to follow a particular path in the forest because she makes the assumption that people have walked it earlier. Forest paths are transient features in the environment; if they are not used they vanish. Their state (how well-worn they are) can indicate how frequently or recently they have been used, which is typically not possible with a road.

We see therefore that social navigation relies on the way that people occupy and transform spaces, leaving their marks upon them — turning a 'space' into a 'place' in the terminology of [Harrison, S., and Dourish, P., Re-Place-ing Space: The Roles of Space and Place in Collaborative Systems. ACM Conference on Computer-Supported Cooperative Work CSCW'96, Boston, 1996, 67-76]. In time, the social cues they leave behind can become sedimented and formalised, transformed into social practices (such as letting people get off the train before you get on), rules and regulations (such as those governing driving) or artefacts (such as signs and landmarks). Social computing, in the sense of our individual actions being designed around collective social behaviour, is not just something that is 'layered on top of' a space, but comes to transform both the space and the ways that people act within it. To design with such ideas is to leave yourself open to the possibility that actors will render your system unrecognisable by you and your co-designers.

Acknowledgement
The idea of social computing has been developed in interaction with the members of the HUMLE laboratory at SICS and Andreas Dieberger (IBM, Almaden), Paul Dourish (University of Irvine, USA).

Link:
http://www.sics.se/humle/socialcomputing/

Please contact:
Kristina Höök — SICS
Tel: +46 8 633 1514
E-mail: kia@sics.se

Wireless Wellness Monitor

by Juha Pärkkä and Raimo Lappalainen

The Wireless Wellness Monitor project was started in order to apply home networks to self-monitoring of health. A prototype of a wireless health monitoring system for home-use was built and feasibility of such a system was studied. In WMM, a weight management application was implemented on top of the system.

Our vision is that in the future, wireless home networks will be available in many homes. The core of the home network will be a *home server* with processing power comparable to PCs. The home server will also provide space for data storage and a gateway to the Internet. *Wireless terminals* will be able to access the home network locally (eg, by using Bluetooth) or remotely (eg, by using Internet). *Measurement devices* will be able to join the home network and send the acquired data to home server for storage, analysis and display. A wireless home network will allow peripherals (eg measurement devices), wireless terminals (eg smart phones) and data storage devices to communicate wirelessly.

Based on this vision, we built a prototype (see Figure 1) with currently available hardware and software. Self-monitoring of weight was chosen as the first application field. The fundamental idea behind the application is the *behavioural feedback model* (see Figure 2).

The WWM software provides four information sources to the user: (1) *Behavioural Feedback System* displays the measured data immediately after each measurement on user's wireless terminal and allows browsing of measured and analysed data whenever, wherever needed. (2) *Nutrition and Exercise Database* gives calorie values of different food items and mean calorie consumption of different activities. It also lets the user compare the foods and activities. (3) *Instructions Database* provides general information on weight management. (4) *Expert Service Interface* makes data exchange with professional healthcare services possible.

In the current system, when the user weighs himself, the result is stored automatically on the home server and displayed on the user's wireless terminal. Charts of long-term weight history, instructions and calorie comparisons can be viewed on the wireless terminal. The *instructions database* as well as the *nutrition and exercise database* can reside on the Internet, so that they can be updated regularly and the user can always use up-do-date information. The *Expert*

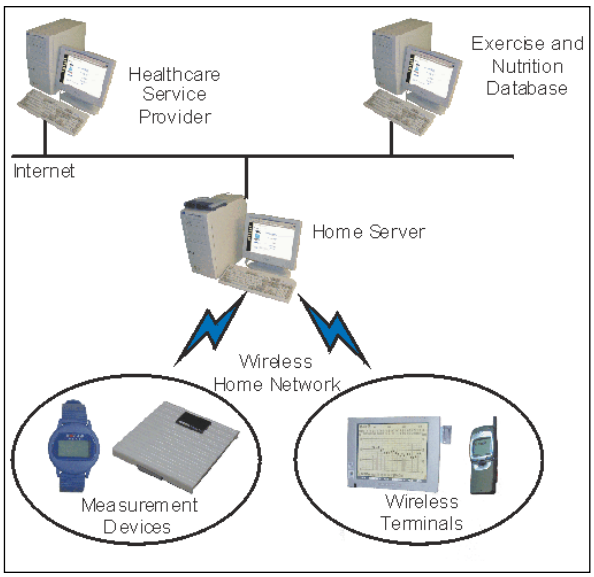

Figure 1:
The home network relies on the Wireless Local Area Network (WLAN) technology. The home server is a PC with WLAN access point. The measurement devices, a scale and a heart rate monitor, can join the home network via a proxy, which is a portable computer equipped with a WLAN card. Two wireless terminals are used, a portable pen computer with a WLAN card and a WAP enabled cellular phone. The core of WWM software runs on the home server as a collection of Jini services implemented using Remote Method Invocation (RMI).

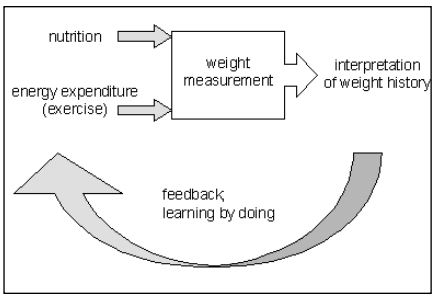

Figure 2:
Behavioural feedback model for weight control. The behavioural feedback system provides data of changes of weight and physical activity immediately and over long periods of time.

Service Interface is not yet fully implemented and the *Instructions Database* has not been integrated to the system.

A feasibility study was carried out as the software specification and a rough demo version of the software were available. The aim of the feasibility study was to find out how possible users evaluated the services that WWM provides. Especially, focus was on the content of the services, ie to find out if the services were what the candidate user needed and whether the services were appraised as useful and easy to use. 12 persons (9 women and 3 men, age between 23 and 61) participated individually in a structured theme interview. Three persons had no past or present weight problem, while the others had or had had overweight. Three men and two women had never tried to lose weight, but the remaining persons actively managed their weight or had done so at some point in life.

The services were seen to be useful and the interviewees did not see any critical problems associated with the services. *The Behavioural feedback system* and the *Expert service interface* were seen as the main motivating factors in the WWM. The interviewees thought that the most probable user groups would be the young people and women. General arguments for this was that the young people are more technology oriented, and that women are more interested in weight control issues than men. Four of the interviewees spontaneously mentioned the mobile aspect of the system as a benefit. The benefit they mentioned was the increased availability of advice (from *database services*) in difficult situations, eg, at dinner table or during grocery shopping.

Behavioural feedback system was seen to motivate long-term use, because it shows the consequences of one's behaviour. *Nutrition and exercise database* divided the interviewees into two groups more clearly than the other services. Interviewees who kept nutritional diary, regarded the service as highly useful as it would help them to find the information they needed. On the other hand, persons who were not on such a diet considered the service less useful. *Instructions database* was seen useful, but problematic because so much information is available in magazines and other literature. Frequent updating of the instructions database was expected. *Expert Service Interface* was found necessary for connecting the system to a professional healthcare service. Such connection was considered to make the system more humane. It was thought especially useful for beginners to learn how to do weight management right. Five interviewees saw consultation as anxiety arousing, because the user might have the feeling that he is being observed.

In WWM, the research concentrated on building a generic software framework to be used in different applications utilising a home network. A demo application for weight management was developed. New types of wellness monitoring applications utilising the framework will be developed in future projects. The new, wireless environment facilitates installing of different peripherals to the home server. Also the amount of user interaction needed in the process consisting of measurement, storage and display is minimised and data exchange with professional healthcare services is made possible.

The WWM project took place in January-December 2000. The partners involved were VTT Information Technology, Department of Psychology in University of Tampere, Polar Electro Oy and Nokia Research Center. The project was partly funded by TEKES.

Link:
http://www.vtt.fi/tte/samba/projects/wwm/

Please contact:
Juha Pärkkä — VTT Information Technology
Tel: +358 3 316 3346
E-mail: juha.parkka@vtt.fi

HCI Research at the University of Manchester Institute of Science and Technology

by Alistair Sutcliffe

HCI research in UMIST covers a wide area of the discipline and has the unique flavour of integrating usability engineering within a wider view of systems design and software engineering. The Interactive Systems Design group combines researchers in both software engineering, human computer interaction, linguistics, psychology and sociology giving it a unique capability for synthesising multi-disciplinary research contributions to advance basic science and the design process in socio-technical systems engineering.

The group's mission is to improve the design process of complex social technical systems for business effectiveness and enhanced quality of life, by developing methods and tools from a sound theoretical base. The Interactive Systems Design group is composed of two units, the Centre for HCI Design and the Centre for Expertise in electronic Commerce.

The Centre for HCI Design conducts a wide spectrum of research in two themes: cognitive engineering applied to interactive technologies (virtual reality, multimodality, multimedia) across a range of application areas (information retrieval, safety critical systems, distributed working and assistive systems) and requirements engineering including methods and tools for scenario based RE, requirements acquisition, collaboration in the RE process, theories of domain knowledge, reuse and component-based software engineering. Other interests are HCI theories of interaction, information extraction, knowledge representation, natural language processing, design of ubiquitous and mobile systems, analysis of multimodal communication in collaborative systems, and design of assertive systems for users with cognitive impairments, learning difficulties and blindness. Recent significant achievements have been research programme providing a significant input to the ISO 14915 standard on multimedia user interface design, and development of scenario based requirements analysis method which have been adopted by industry.

The Centre for Expertise in electronic Commerce and the eCommerce design group researches methods and tools to improve the design of web based systems. Interests range from surveys and experimental studies to understand the attractiveness of commercial web sites and the design features that influence people's attention and purchasing behaviour. Design guidelines have been developed to enhance the persuasiveness of eCommerce web sites. Other interests are development of smart card technology for customer profiling and personalised user interfaces, and investigating how smart card technology can make security more acceptable to users. The Centre collaborates closely with industry on a range of projects that create an integrated approach to the eCommerce design problem, eg starting with business modelling to develop Internet technology strategy, requirements analysis methods for site development and customer profiling, usability guidance for web site design, to architectures and tools for web site development. Recent achievements include development of a requirements capture method based on multidisciplinary product teams adopted by ICL in their 'Marketing to Design' programme and widely applied within other organisations and the Centre's role in developing competitive advantage for regional eCommerce has been recognised by DTI minister Patricia Hewitt.

Funded projects within the ISD group reflect the central theme of design of social technical systems. The EPSRC SI programme SIMP (Systems Integration for Major Projects) in collaboration with BAESystems is investigating three main themes: assessment of performance of socio-technical systems against a set of operational scenarios; modelling the communication, interaction and effectiveness of intelligent technology in socio-technical systems, and developing methods for generating scenarios for requirements validation in complex systems. The requirements engineering theme continues in the EPSRC ISRE (Immersive Scenario based Requirements Engineering) project which is investigating methods for requirements engineering with virtual prototypes, a taxonomy of scenarios to configure virtual environments and prototypes for testing, and design guidelines for development of VR toolkits for collaborative requirements analysis. The EPSRC CORK (CORporate Knowledge for requirements engineering) project is developing methods for assessing and designing inter-organisational relationships and requirements for supporting computer technology. The theories of transaction costs and value chains are being extended to create a method for analysing the complex inter-organisational relationships and business processes, transactions, and system requirements for virtual companies and alliances. A scenario-based method identifies system requirements to support business processes and enhances reuse of business knowledge and design experience held in corporate memories. The EU funded DUMAS project is developing information extraction tools and speech technology for mobile systems, so users can interact remotely with complex databases.

In the eCommerce and Internet research the EPSRC funded project Human factors in the design of Electronic Service Delivery Systems for use in complex environments aims to inform the development and assessment of interactive electronic service delivery systems for use in complex environments, where user choice exists and where the influence of social and physical factors impact on usage and usage decisions. This project is a collaboration with the Manchester Federal School of Business

and Management. The DTI/EPSRC Teaching Company Scheme—Customer relationship management through streaming media and VOIP technologies-collaborates with a local industry to support development of e-business via a more customer focused service with online help facilities and personalisation of the shopping experience. As a result of utilising VOIP and multimedia streaming technologies, the project aims to improve visit-sale conversion rates, increase repeat business, and increase the value per sale.

Links:
http://www.co.umist.ac.uk/hci_design/

Please contact:
Alistair Sutcliffe — University of Manchester Institute of Science and Technology, United Kingdom
Tel: +44 161 200 3315
E-mail:@ a.g.sutcliffe@co.umist.ac.uk

Human-Computer Interaction Activities in Italy

by Maria Francesca Costabile and Fabio Paternò

We describe the growing interest in HCI activities in Italy and the objectives of SIGCHI Italy, the Italian chapter of ACM SIGCHI, the Special Interest Group on Computer-Human Interaction.

The mission of SIGCHI is to bring together people working on the design, evaluation, and implementation of interactive computing systems for human use. This embraces work on the hardware and software engineering of interactive systems, the structure of communications between human and machine, the use of interactive systems, design methodology, and new designs themselves. SIGCHI helps to organize and coordinate information through on-line dissemination, traditional publications, conferences, workshops, and symposia relating to human-computer interaction. SIGCHI members are interested in keeping up with changes in this rapidly evolving field—making computers easier to use and extending technology into new domains for new users.

The interest in human-computer interaction is increasing rapidly in Italy. The first research groups in this area began work in the late eighties-early nineties. Every year a growing number of universities offer HCI courses in various degree programmes (eg computer science, computer engineering, psychology, communication science). A list is available at the SIGCHI-Italy website, at http://giove.cnuce.cnr.it/sigchi/HCIcourses.htm.

Since 1992, an important international conference on Advanced Visual Interfaces (AVI) has been held regularly every two years in Italy with the sponsorship of ACM. Reports have been published in the October 1994, January 1997 and October 1998 issues of the SIGCHI Bulletin and the conference proceedings are published by ACM Press.

To further support this growing interest in HCI, in 1995 a number of Italian researchers applied to set up an Italian chapter of SIGCHI. SIGCHI-Italy was officially started by ACM on April 24 1996. The purpose is to promote and increase knowledge and interest in science and technology of human-computer interaction in Italy. The Italian chapter organises meetings, conferences, discussion groups and workshops involving people from both academia and industry. The group of members has grown rapidly and now involves about 150 people. There are two levels of membership: non fee-paying non voting members and voting members, who also have to be members of and pay subscription fees to ACM ASIGCHI. Consequently, the chapter has no budget, it is managed completely by volunteers and events are organised thanks to contributions from sponsors. SIGCHI-Italy members contribute in many ways: proposing thematic meetings, participating in the organization of events, advertising the activities, creating links with other organisations. Since members of the chapter are spread throughout Italy, it is not possible to have meetings too frequently: usually meetings are held every three months often in Rome.

SIGCHI-Italy also has a strong industrial component. Italian companies that have been or are now actively involved in the chapter include SOGEI, Etnoteam , FIAT, and CSELT, the research center of Telecom, Italy. For example, the meeting of the chapter in June 2000 at CSELT provided an interesting opportunity for our members to view new prototypes for vocal interfaces and interfaces for third-generation cellular phones with UMTS support. Our most recent meeting was in March at the FIAT Research Centre, where much work is under way on the development of interfaces for car devices.

Other groups are active in Italy on HCI-related topics and SIGCHI-Italy aims to maintain close contacts with them. Examples are the Italian Society of Ergonomics, the Italian Association of Computer Science — AICA, the Italian Association of Artificial Intelligence — AIIA. AIIA includes a working group which focuses on the development of adaptive, agent-based interfaces.

To further encourage the growth of the Italian HCI community we also organised the first Italian Symposium on HCI in Rome in 1999 (information is still available at http://giove.cnuce.cnr.it/hcitaly99.html). Given the success of this event, we are now organising a second symposium which will be held in Florence, in September, in conjunction with the conference of the Italian Society of Ergonomics (information available at http://giove.cnuce.cnr.it/hcitaly01.html). We also intend to increase cooperation with other SIGCHIs, primarily those located in Europe, in order to coordinate efforts and activities.

Link:
SIGCHI-Italy: http://giove.cnuce.cnr.it/sigchi/

Please contact:
Maria Francesca Costabile — Dipartimento di Informatica, Universita' di Bari, Italy
Tel: +39 080 544 3300
E-mail: costabile@di.uniba.it

Fabio Paternò – CNUCE-CNR
Tel.: +39 050 315 3066
E-mail: fabio.paterno@cnuce.cnr.it

The Usability-Engineering Laboratory at GMD

by Wolfgang Dzida and Regine Freitag

Research at GMD involves trying to find an answer to the question: "Is a usability laboratory essential for usability testing?" Usability testing is taken as part of a comprehensive service accompanying the entire usability engineering process.

Two cases of investigations are typical for usability testing:

- Testing a hypothesis: A usability assessor inspects a user's critical work session and builds a hypothesis about what might be wrong with the product.
- Exploring task performance: The efficiency of user performance is studied for an entire work task. The usability assessor uncovers problems in system usage that make user performance unnecessarily difficult.

Investigations without a Usability Laboratory — when is this possible?

For every critical task the usability assessor must analyse the user's efficiency and effectiveness in performing this task. Having a clear understanding of the flow of work, both user and assessor are able to interpret an encountered problem. If it is immediately and indisputably clear what is causing the problem, then it suffices to observe or inspect the situation and document it appropriately. Employing a usability laboratory would deliver the same result but with much higher costs. Hence, there should be good reasons when using such a laboratory.

Investigations with the Aid of a Usability Laboratory — when is this necessary?

The previously mentioned types of investigation require specific methodological approaches while using the equipment of a usability laboratory. In the first case the investigation is governed by a hypothesis the assessor aims to verify, but it is yet unclear whether the assumption concerning the cause of the usage problem is correct. In the second case it is not even clear whether a usage problem exists at all, and this must be explored empirically.

Hypothesis-driven Investigation

Usually, the user causes the usability assessor to have a closer look at a usage problem that is hard to work around. A plausible hypothesis as to the possible causes of the trouble is generated. The first step in forming a hypothesis is to identify the 'mismatch' between the user's intention in a particular task, and the designer's intention in providing the appropriate means for that task.

Scrutinising the mismatch requires a closer look at the case. However, the case is temporary. Should one ask the user to reproduce the critical situation again and again? Since the observational situation is also a social situation between people, this could mean expecting too much from the user. The short-lived interaction with the system can be more effectively captured by means of a video camera, thus allowing the assessor to repeatedly watch the episode afterwards in order to analyse and comment on it. Interpreting a mismatch is done by analysing the intentions of the user and contrasting them with the intentions of the designer. The GMD laboratory applies a specific template for coding the intentions in a semi-formal way, so as to precisely and completely identify the sources of mismatch.

Exploring Task Performance

Exploring task performance is especially interesting during usability prototyping to support requirements development. This approach is also referred to as explorative prototyping. A typical case of mismatch when using a new system or prototype is called opportunistic use (Carroll, 1992). This may happen, if, for example, the user misses an appropriate 'affordance' (Norman, 1989). The user then does not understand the designer's intention and starts exploring some plausible opportunities.

In the case that an opportunistic strategy fails, the user runs into a usage problem. Because it can hardly be foreseen when this situation will occur, the whole session should be recorded with the aid of a usability laboratory. The laboratory primarily serves for the storage of observational data. Furthermore, the data can be exploited in order to rapidly extract and analyse the most interesting episodes out of the mass of recorded data.

The aim of exploration is to observe different users and study how they find their individual ways in conducting their tasks. Although the designer of the product has implemented an ideal way, it is unlikely that each user will intuitively follow this path. There is rarely 'one best way' of doing things. A usability laboratory is well equipped to record the variety of approaches employed by different users. The designer himself can watch the recorded situations, can listen to the user's comments and will be well prepared for a face-to-face meeting with participating users.

Usually, an investigation of interactive prototypes takes place in a usability laboratory only when the prototype has achieved a fairly advanced stage of development. Most of the preceding iterative improvement steps can be done in a more heuristic manner.

Equipment of the Semi-portable Laboratory

- Coding of behavioural data by means of 'The Observer Video Pro' (NOLDUS).
- DATech Procedures for Usability Testing based on ISO 9241-10.
- DATech Procedures for Usability Engineering Process Testing based on ISO 13407.

Link:
http://borneo.gmd.de/~freitag/

Please contact:
Wolfgang Dzida and Regine Freitag — GMD
Tel: +49 30 61 62 57 81
E-mail: wolfgang@dzida.de,
regine.freitag@gmd.de

Dependable Interactive Systems Research at York

by Michael Harrison

Much of the research at the Human Computer Interaction Group at University of York is concerned with the development of tools and techniques that support modelling and evaluation of dependable interactive systems. Their concern is that these techniques are appropriate so that systems engineers can recognise the human consequences of design decision and cost effective so that the cost of using the method is justified in terms of dependability improvement.

The analysis is applied at three levels: the artefact or device, the human system interface and the collaborative system where devices are viewed as embedded components.

Our concern is to analyse the human consequences of a design and we do this at the device level by producing precise models of interactive behaviour and exploring consequences either as prototypes or by using model-checking techniques. A number of specification approaches have been explored. For example, Degani's OFAN approach to structuring statecharts for interaction has been used as a basis for modelling interactive systems and we have implemented a translation algorithm for generating a state transition model compatible with the SMV model checker. We are currently exploring templates that match typical requirements such as mode-freedom, accessibility and visibility and can be used by engineers without the need to understand the underlying temporal logic required by the model checker. We are also exploring visual system engineer friendly forms of the traces produced as counter examples.

At both the human system interaction level and the collaborative system level it has been necessary to find ways of representing work in such a way that it can be used realistically as a basis for assessment. The problem with task based approaches (in which work is thought of as a hierarchy of goals and the plans by which these goals are achieved) is that they do not in themselves take account of the context in which the work is carried out and they presume or prescribe the way people work in practice. We have therefore focussed on typical or extreme examples of use based on scenarios or narratives about situations in which the artefact is used. We are developing two methods both based on the scenario as the unit of work. (1) THEA is concerned with assessing the vulnerability of a design to interaction failure. (2) Function allocation is aimed at helping system engineers decide which parts of the system should be automated. Both methods address key concerns of dependable interactive systems and are designed to be structured but relatively informal.

THEA is a fairly mature method of human error identification. It has been developed as a questionnaire based approach taking a representative set of scenarios and then systematically addressing components of the scenario through the questionnaire. The scenario is structured in terms of the multiple goals and corresponding tasks that are achieved. The questions are aimed at eliciting the design features that may cause or mitigate against interaction failure and are based on Norman's simple control loop model of cognition. Failures are sparked by problems of perception, interpretation, evaluation, goal formation, planning and action. The method is supported by a prototype tool (ProtoTHEA) and has been used in case studies with Matra Dynamics, BAE SYSTEMS, Airbus and NATS.

The method of function allocation is relevant to human system interaction and collaboration. It involves a method that contains two trade-offs. A set of functions and a set of roles are introduced and refined. The first trade-off concerns matching function to role and deciding how feasible it would be to automate the function. A set of functions to be partially automated are then analysed in terms of a framework (IDA-S) that is used to identify key components (in terms of information, decision, action and supervision) and which role should control these components. In the second trade-off alternative candidate functions represented in this way are then ranked in the context of the scenarios in terms of how well they do against criteria such as workload and performance.

More recent research extends interaction failure assessment to collaborative systems. We are also concerned with methods of dealing with dynamic function allocation and understanding flexible scheduling in the face of hard deadlines in collaboration with York's real time systems group. We are exploring frameworks based on distributed cognition and the role of notions such as accountability.

The research is funded by EPSRC and EU as well as BAE SYSTEMS and Defence Evaluation Research Agency.

A copy of this paper with appropriate links to fuller descriptions of the work can be found at http://www-users.cs.york.ac.uk/~mdh/ddis.htm.

Please contact:
Michael Harrison — University of York, United Kingdom
Tel: + 44 1904 432 732
E-mail: Michael.Harrison@cs.york.ac.uk

Modeling Designers' Activities in the Garment Industry

by Rémy Taillefer and Dominique Scapin

One current challenge in the garment and fashion industry is to reduce the cost and development time for new collections and to allow stylists and modelists to express all their talent, liberating them from the constraints inherent in creating physical prototypes.

To answer some aspects of this challenge, the European project COMEDIA aims at giving fashion designers 3D CAD tools that allow testing new visualization and control ideas, through dynamic simulation of clothes on dynamic virtual models. Two INRIA research teams are involved: 'Mirages' for the analysis and synthesis of images, and 'MErLIn' for ergonomics. With respect to ergonomics, several activities are being conducted in parallel:

- ergonomic evaluation of 2D software currently in use by stylists and modelists
- analysis of the state-of-the-art with respect to ergonomics in the textile industry
- collection and analysis of ergonomic recommendations concerning 3D
- analysis of the state-of-the-art of tools and devices for manipulating 3D objects, eg, eye-tracking, voice and gesture
- analysis of the work activities of stylists and modelists.

The last study is essential to better understand the activities and language of operators in order to specify the proper functions, objects, names and labels for such future systems.

Data Gathering Methodology

After initial interviews, designers were observed in different artisan and industrial businesses in France and Italy. The observations occurred during the course of their actual work, as they designed and implemented models of clothing. During the observations, we used the 'thinking aloud' technique, in which we asked subjects to describe their activities as they performed each task. This helped elucidate the procedures used and the reasoning behind the way they solved particular problems.

The data about modelling and prototyping helped elaborate the task models, which were then validated by the people who had been observed. We videotaped the styling tasks, which are currently being analyzed and will be validated prior to creating the final task model. The initial results led to the description of the tasks according to the following phases:

Styling Task Model

The steps are: (1) prepare the materials; (2) create a sketch; (3) confirm the lines; (4) specify the final details. makes a 2D sketch of the model on paper.

The objective is to make the stylists' creative ideas concrete to maximally limit the interpretations by the modelist. The stylist is autonomous, creating the model according to her personal choices, inspired by what is currently fashionable. The stylist must however consider the constraints on the modelist in her sketches, which involves the mechanical properties of the modeling fabric and the iterations between the two-dimensional sketches and patterns and the three-dimensional prototypes. The stylist can limit these iterations by creating more or less detailed technical specifications (sizes, colors, fabrics) depending on the type of collaboration with the modelist (direct collaboration versus at a distance), the type of business (workshop or industrial) and the type of design (single garment or entire collection).

Modeling Task Model

Traditionally the steps are: (1) prepare the construction work; (2) create the prototype; (3) verify and adjust.
Instrumented, the steps are: (1) find references; (2) prepare the construction work; (3) create the new model; (4) create the prototype; (5) specify the final details of the prototype.

Figure 1: Manual interaction with a wooden mannequin.

Figure 2: Modelist's interaction with Current 2D software.

Figure 3: Example of shirt dynamically simulated on a virtual mannequin (from Mirages.

The modelist's goal is to create a two-dimensional representation (based on sketches, associated descriptive elements and the interaction with the stylist) of each component of the prototype model necessary to create cardboard patterns. Modelling work is characterized by frequent iterations between the representation of 2D materials (pieces of the model laid out flat) and the three-dimensional representation (the model mounted on a mannequin). Iterations between 2D and 3D are numerous in both traditional and instrumented situations: the quality of the model depends upon them. Modelists require working in 2D on the modelling fabric in order to verify the visual properties of the model in the 3D form on the wooden mannequin. This allows them to precisely specify the cardboard patterns, which are the 2D elements that serve as the basis for manufacturing.

Prototyping

The stylist and the modelist can experiment on a wooden or human model and evaluate (on the first prototype) their work on the final fabric with all the accessories, which is not possible during the initial modeling phase. The mechanical properties of the modelling fabric (used to save money and to facilitate tracing lines and inscriptions) are different from the final fabric. The prototypes frequently reveal unanticipated errors, which requires additional modelling iterations.

Conclusion

The initial data require deeper examination, notably in the process of styling and creating prototypes. However, the following activities have been identified for the future specification of certain aspects of the tools, for example:

- mixed 2D-3D representations to limit the number of iterations between 2D-3D models
- a system of exchanging representations between stylists and modelists that permits them to share and visualize ideas
- take into account during the design of virtual models the essential visualization elements, such as the right thread, the closed position of a garment or the possibility of generating 3D versions in different sizes.

Based on further analysis and on a complementary design study focusing on the representation of objects in space, the next goal will be to conduct an evaluative study of a prototype, in collaboration with our industrial partner LECTRA Systèmes.

Please contact:
Rémy Taillefer — INRIA
Tel: +33 1 3963 5079
E-mail: Remy.M.Taillefer@inria.fr

Dominique Scapin — INRIA
Tel: +33 1 3963 5507
E-mail: Dominique.Scapin@inria.fr

Teaching Computers to spot Interesting Images

by Eric Pauwels

In the EU-funded FOUNDIT project researchers at CWI are designing search engines that can be taught what images to look for on the basis of evidence supplied by the user.

Vision is an amazingly powerful source of information about the world around us. As a consequence, a considerable part of the data we collect comes in the shape of visual material, and multimedia libraries are rapidly filling up with images and video-footage. However, while progress in electronics reduces the costs of collecting and storing all these images, an efficient database search for specific material becomes increasingly difficult. As it is virtually impossible to classify images into simple unambiguous categories, retrieving them on the basis of their visual content seems to be the

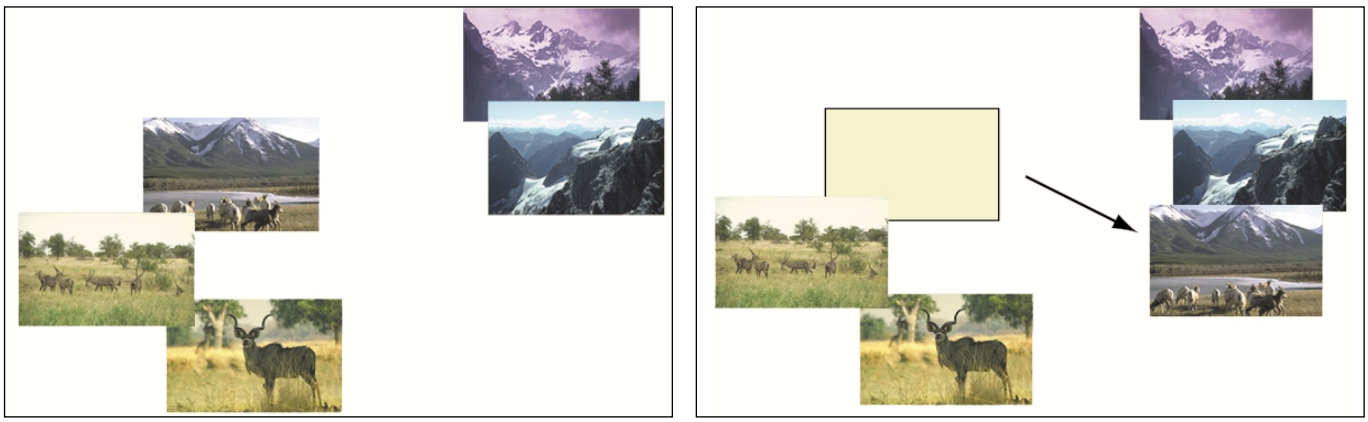

Interactive definition of image similarity. Left: The interface displays five images grouped together to reflect its current model of image similarity. Right: If this grouping does not match the user's appreciation of their similarity, he can rearrange them on screen into a more satisfactory configuration by a simple drag-and-drop operation. For example, if the user considers the top left image to be more reminiscent of mountains, he will drag it over to the image group on the right. In the next iteration step, the search engine will adapt its features to reproduce this user-defined similarity as faithfully as possible.

only viable option. Rather than asking the computer 'to display image AK-C-4355.jpg', we want to instruct it to retrieve an image showing, eg, the 3-dimensional structure of a DNA-molecule. Or, imagine an interior decorator looking for visual art to enliven his client's newly refurbished offices, requesting the system to find vividly coloured abstract paintings. These scenarios call for vision-sentient computer systems that can assist humans by sifting through vast collections of images, setting aside the potentially interesting ones for closer inspection, while discarding the large majority of irrelevant ones. Being able to do so would effectively equip the search engine with an automated pair of eyes with which it can reduce for the user the drudgery of an exhaustive search. Such systems will have to be highly adaptive. Indeed, what constitutes an 'interesting' or 'relevant' image varies widely among users, and different applications call for different criteria by which these images are to be judged.

In the FOUNDIT project search engines are designed that can be taught what images to look for on the basis of evidence supplied by the user. If the user classifies a small subset of images as similar and/or relevant, the system extracts the visual characteristics that correlate best with the user's classification, and generalises this ranking to the rest of the database. Only images marked as highly relevant are displayed for closer inspection and further relevance feedback, thus setting up an iterative cycle of fine-tuning. The interaction between user and interface is very natural and transparent: similarities between images are indicated by dragging them into piles on the screen, and relevant images are selected by simply clicking on them. Such operational transparency is of great importance for applications, such as E-commerce.

Link:
http//www.cwi.nl/ ~pauwels/PNA4.1.html

Please contact:
Eric Pauwels — CWI
Tel: +31 20 592 4225
E-mail: Eric.Pauwels@cwi.nl

Modelling SAR Images of Urban Areas

by Ercan E. Kuruoglu and Josiane Zerubia

Satellite imagery has found vast applications in a wide spectrum of areas including agriculture (eg detection of crop types), urbanization (tracking the development of urban areas), cartography (eg detection of rivers, road networks), warfare (eg detection of targets, surveying), etc. This heavy demand on satellite imagery applications lead to the development of imaging systems that are alternative to optical imagery. In particular, synthetic aperture radar (SAR) imagery in the last two decades has become increasingly popular as some of its properties are favorable to optical imagery. SAR imagery can operate regardless of weather conditions and SAR image resolution is independent of sensor height.

There are problems associated with the nature of radar imaging process. SAR images are formed by the reflections of waves from the earth surface. If the surface roughness is comparable to the wavelength, the wave reflections occur in various directions. These directions are also dictated by the incidence angle, surface structure, dielectric constant of the surface, too many factors to be successfully accounted for by a simple deterministic model. The coherent addition of these reflected waves in the receiver out-of-phase results in the granular appearance of the images which is referred to as speckle noise.

The assumptions of the classical SAR image generation model lead to the convention that the real and imaginary parts of the received wave follow a Gaussian law, which in turn means that the amplitude of the wave has a Rayleigh distribution. This classical model has been successful in some cases such as modelling particular natural scenes. However it is known that it fails to provide an efficient model for urban scenes where man-made structures provide bright reflections. The impulsiveness of these images suggests underlying heavy-tailed clearly non-Rayleigh distributions. Some alternative distributions have been suggested such as the Weibull and log-normal distributions, however, in most cases these models are empirical, and do not generalise. By relaxing some of the assumptions leading to the classical Rayleigh model, we have developed a more general and successful model for SAR images of urban areas.

The classical SAR image formation model assumes that the reflections that add up in the receiver are Gaussian due to the central limit theorem. However, this model does not consider the attenuation that is faced during the wave propagation. Consideration of this attenuation leads to a more general model for the received signal which is alpha-stable distributed. This is in accordance with a generalised version of the central limit theorem. Recent work on multiuser communication which has a similar model for wave propagation and the distribution of sources also led to the alpha-stable model. Alpha-stable distributions have received wide interest recently as efficient models for impulsive data and have found applications in radar signal processing, financial time-series analysis, communications and audio signal processing. They are generalisations of the Gaussian distribution and are governed by two additional parameters one setting the impulsiveness, the other the skewness (unsymmetry) of the distribution.

It can be shown that starting from alpha-stable assumption for the complex parts of the received wave, the amplitude is distributed with a generalisation of the Rayleigh distribution which we call heavy-tailed Rayleigh distribution. This distribution has two parameters setting the dispersion and the impulsiveness of the distribution. A group of these

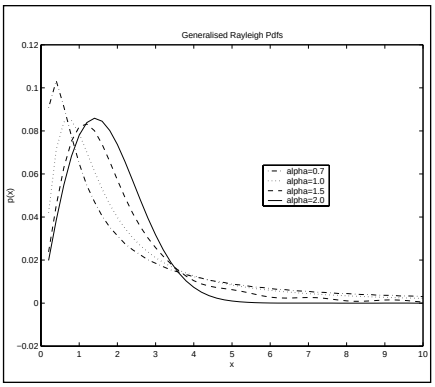

Figure 1: Generalised Rayleigh probability density functions for various values of impulsiveness parameter.

Figure 2: X-band SAR image.

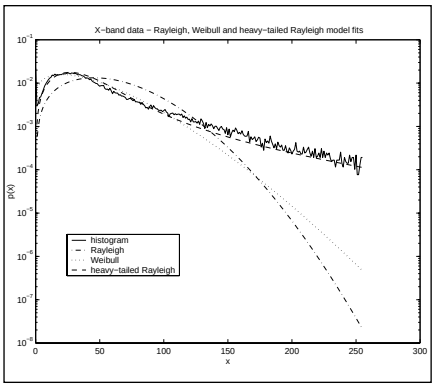

Figure 3:
Histogram of the X-band SAR image, and fittings of classical, generalised Rayleigh and Weibull models in logarithmic scales.

distributions is plotted in Figure 1 for varying values of the impulsiveness parameter.

For this model to be of any practical use, efficient algorithms for estimating the model parameters are needed. We have thus also derived novel methods for the estimation of the heavy-tailed Rayleigh distribution parameters based on negative fractional-order statistics which are of great interest in themselves since this is one of the first applications of the new concept of negative fractional order moments. Our experimental results show that the heavy-tailed Rayleigh model can describe a wide range of data which could not be described by the classical Rayleigh model.

In Figures 2 and 3, we provide the results of some simulations on SAR images of urban areas. It is clear that the new generalised (heavy-tailed) Rayleigh distribution provides a better statistical description for the SAR images of urban areas when compared to the classical Rayleigh model and other alternative models suggested in the literature such as the Weibull distribution. The estimation technique we suggested also makes it more attractive than other heavy-tailed models such as the K-distribution for which model parameter estimation is particularly difficult. A wide survey on images taken at different bands and weather conditions (such as humidity) would provide us with a relationship between model parameters and imaging conditions which would be very useful in interpreting SAR images.

Please contact:
Ercan E. Kuruoglu — IEI-CNR
Tel: +39 050 315 3128
E-mail: kuruoglu@iei.pi.cnr.it

netzspannung.org — an Experimental Architecture for interfacing Digital Art, Culture and Technology

by Monika Fleischmann, Wolfgang Strauss, Jasminko Novak and Stefan Paal

The German Federal Ministry for Education and Research regards the promotion of competence centres dealing with artistic approaches and media culture as an important boost to the development of information technology. GMD's MARS Exploratory Media Lab is developing a model for an online media laboratory as a competence centre for digital art, culture and technology.

GMD.IMK-MARS is developing an online media laboratory for interdisciplinary networking between digital art, culture and technology. As two fundamental pillars of an online media lab we consider the tools for active production and an architecture for connecting archives into a collaborative knowledge space. To this end the netzspannung.org system is conceived as a toolbox for experimentation with networked media spaces combined with structures for knowledge discovery. Such architecture is intended to serve both as a toolbox for active producers of digital culture as well as a public information interface. This requires an open and user-extendable architecture that enables the users (artists, designers, technologists) to explore, re-shape and re-design the platform to their own needs. To address these requirements we are developing a distributed community context architecture accompanied by a selected set of dynamic interfaces.

Open Platform Architecture

The basis is an open distributed system architecture. The challenge is to build a system which not only deals with different user requirements in the heterogeneous environment of the Internet, but also enables the users to modify and extend the system with new functionalities, data archives and interfaces, at different levels of expertise. Besides the ability to modify the platform itself, the users can integrate their own computers as active nodes in a distributed network of experimental media spaces. This is supported by the following system architecture:

Perceptive Interfaces

Our starting point is the design of interfaces for connecting the building

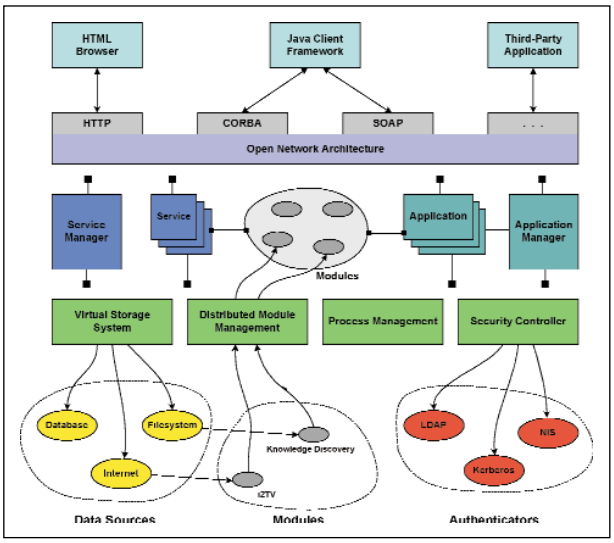

Figure 1: netzspanung.org distributed system architecture

blocks of real and virtual spaces to new modalities of perception. The basic set includes the timeline, netzkollektor, workspace, knowledge map and i2tv. The timeline is a multilayered, hypermedia structure that shows the development and mutual interaction of different works and discourses from art, science and technology in the form of contextualised information and a network of relationships. i2tv (interactive Internet television) is a mixed reality communication interface for live media productions and networked scenarios. It combines 3D Internet environments with digital television and mobile communication interfaces.

Knowledge Discovery

Methods for accessing, structuring and visualising existing information flows and databases on the basis of semantic relationships and self-organising maps enable the extension of isolated archives to connected knowledge spaces.

Distributed Module Management

Transparent integration of additional functionalities is enabled by dynamic location and loading of modules, and resolving of their dependencies. Modules can be located anywhere on the Internet and accessed as transparently as if they were stored locally on the machine.

Virtual Storage System

An open storage interface and data abstraction layer enable transparent access to heterogeneous data sources, without regard for where and how the data is stored. Different applications can share the same data archives, even when they span different storage types such as object-oriented and relational databases, XML repositories or network connections.

Open Network Architecture

A protocol abstraction layer supports the implementation of networked systems not only in Client-Server-Architecture, but also in a domain-based architecture or within a P2P network. This enables concurrent use of different network protocols (eg CORBA, SOAP) as well as individual selection based on the needs of a given application.

Distributed Server Framework

Though netzspannung.org does not want to specify how the implementation of server or client nodes must be designed, it provides a ready-to-go framework for distributed server implementations. The framework is a basis for extension by custom modules and provides both runtime and developer support. The C++ release is accompanied with a Java version, that comes with Java Servlet, Java Server Pages (JSP) and Extended Server Pages (XSP) support.

Java Client Framework

A community platform, which requires only a native HTML browser, can be accessed by most Internet users, without installing additional software portions. However, for more dynamic interfaces, active clients are needed. For this purpose, we provide a Java Client Framework that is able to run in a browser window and needs only a Java Plugin for execution. Additional software packages can be loaded from the server and extend the client with the needed functionality.

Pilot Projects

An important part of the development of netzspannung.org are pilot projects. The production potential of the platform is demonstrated and evaluated in exemplary applications in co-operation with individual artists and cultural institutions. The projects are selected thematically and involve research into concepts for collaborative knowledge spaces and networked media spaces on the basis of artistic models and strategies.

Links:
http://netzspannung.org/
http://imk.gmd.de/mars

Please contact:
Jasminko Novak and Stefan Paal — GMD
Tel: +49 2241 14 2809
E-mail: info@netzspannung.org

Figure 2: i2tv interface for 'Ottos Mops' distributed play.

City-Traffic: An Artificial Reality System of Traffic in an Urban Area

by Heinz Mühlenbein, Ulrich Bartling and Uwe Beyer

The computing power of large PC clusters has made possible the implementation of computer systems which process hundreds of thousands of individual agents in real-time. This is designated an Artificial Reality System (ARS) if the system is in continuous interaction with the real system via sensor data. An ARS is a detailed image of the real system on the computer, and can be used for analysis, planning, and control. Even prediction is possible if part of the ARS can be cloned and made to run faster than real-time. Researchers at the Institute of Autonomous Intelligent Systems of GMD have developed a new technology for implementing AR systems. This technology has been used to develop a large ARS, called City-Traffic. The demonstrator version simulates the traffic of the city centre of Bonn.

New computer models are essential in understanding today's world. Research and development that relates to mapping real-world systems into computer systems has so far concentrated on the visual presentation of objects (virtual reality). It is now moving to augmented reality and artificial reality. Augmented reality shows real objects in more detail than can be perceived in reality. Artificial reality models the dynamic behaviour of a set of real objects. Thus virtual, augmented and artificial reality systems show true reality from different perspectives.

Artificial reality means reproducing the real world in detail on computers. Furthermore, the artificial and real systems can be made to interact. The state of the real system is continuously monitored by thousands of sensors (in the case of City-Traffic by induction loops, floating car data etc). At intervals of, say, three minutes, the states of the real system and the artificial system are compared. In a process similar to relaxation in mathematics, the state of the artificial system is kept almost identical to the state of the real system.

Short-term prediction is done as follows. Every six minutes, for example, the ARS creates a child which then runs faster than real-time. Thus, if the child runs ten times faster than real-time, in two minutes it can make a prediction for the traffic situation twenty minutes in the future.

For the implementation of artificial reality systems, the potential of multi-agent systems has steadily grown. Here, every agent represents a sub-unit of the system. Depending on the context considered, these sub-units could be local processes, components, individuals, species, cars, firms etc. The agents may interact directly or indirectly on different time scales and/or spatial scales.

Research in artificial reality at the GMD started with a theoretical investigation of networks of agents. We then designed an architecture called the 'Flip-Tick Architecture' (FTA) to facilitate the implementation of artificial reality systems. This research has been funded since 1994 by the Real World Computing Partnership financed by METI from Japan.

Flip\,-Tick Architecture

The origins of the Flip-Tick Architecture (FTA) are based on several well-known organisational principles, software architectures and hardware structures. Major contributing factors to the development of FTA include:

- blackboard architectures
- distributed computing
- multi-agent systems
- object-oriented design
- Petri nets.

The central processing units of FTA are fine-grained software agents called 'actors' (up to 800 000 per host on a 256 MB PC). Actors read typed data containers (called 'tags') from a special memory called a 'tagboard', and write tags to tagboards for later use. An FTA implementation can have many tagboards (up to 30 000 per host) and each actor has access to any tagboard. A tagboard has two faces - the face with data accessible for reading (read-only) and the face with the new data (write-only). The FTA architecture is cycle-oriented. After a cycle, the tagboards are flipped. The data read in the previous cycle is erased. The data written is now available for reading. The efficient implementation of FTA on a PC cluster was a considerable technological software challenge.

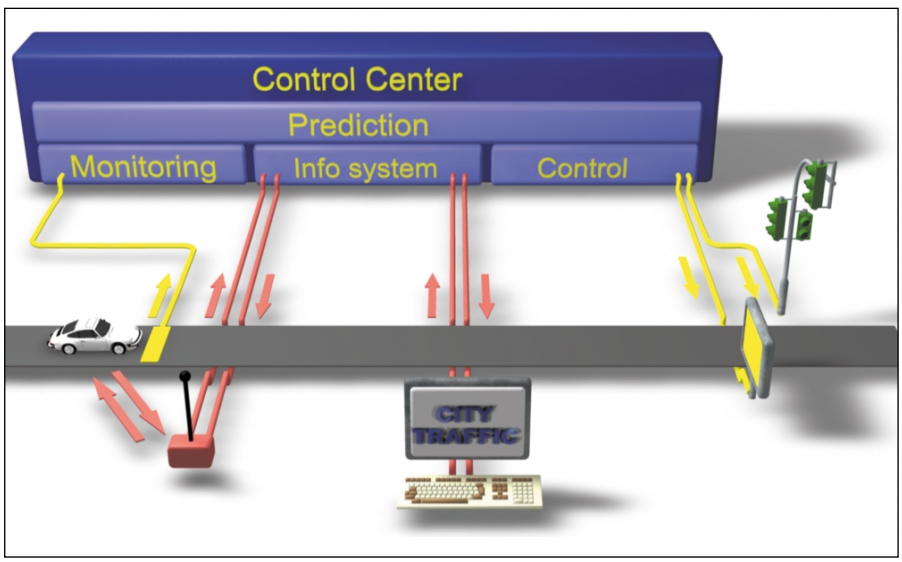

City-Traffic consists of a monitoring system, an information system, a control system, and a prediction system. The system is intended to run continuously in a traffic control center.

City Traffic

Previous investigations on traffic dynamics have mainly concentrated on freeway traffic, and therefore deal with phenomena like traffic jams, stop-and-go traffic, etc. Quite a different situation exists in cities, where we do not have a large number of cars heading in one direction, but rather a large number of cars with very different destinations and driver schedules. A suitable way to implement such a real-world situation on computers uses agents representing individual cars, which move on a detailed map of the real city road system.

For the implementation of a commercial artificial reality system we started a project with the city of Bonn. The following information is used as an input for the system: road maps with intersections, traffic lights on intersections, average number of cars per home in a certain area, and time-dependent traffic density.

The traffic light system of the city of Bonn produces online information about the current flow of traffic. The system is not able to identify or distinguish individual cars. There is no information about individual starting or destination points of traffic members. Thus City-Traffic simulates the same amount of traffic, but not individual routes. The online information is used to continuously update the state of the artificial reality system.

We have recently completed a simulation of the city ring of Bonn. For the traffic of the whole urban area we will use a PC Cluster of 24 processors. With this configuration we will achieve an execution time of two hours real-time per 24 hours of real traffic in Bonn. A further function of City-Traffic is an information system where citizens can access actual traffic information via the internet. The information system also makes routing proposals. The total system is shown in the figure.

In addition to the particular situation of Bonn, our computer program is set up to handle even larger cities, with up to 20,000 intersections, up to 40,000 road sections, and up to 200,000 cars moving at the same time.

Acknowledgment

The research has been funded by the Real World Computing Partnership financed by METI from Japan.

Links:
http://borneo.gmd.de/AS/art/index.html
http://borneo.gmd.de/AS/ct/ct_engl.html

Please contact:
Heinz Mühlenbein — GMD
Tel: +49 2241 142405
E-mail: muehlenbein@gmd.de

The Possible Effect of a Turbine Testing Plant on the Quality of the Air — A Case-study for Air Pollution Transmissions

by Sándor Márton and Tamás Rapcsák

The amelioration of the state of environment is a key-issue for the countries belonging to the European Union (EU) and for those wishing to join EU. One of the most serious problems is air pollution, caused mostly by industrial plants. Based on the Hungarian National Standards, scientists at SZTAKI have been studying the transmission of pollutants emitted by point-sources (stacks), linear sources and surface sources, for several years, elaborating numerous case-studies in the field.

The expected effect of a future turbine testing plant on the air quality has been studied in the framework of a case-study ordered by General Electric Ltd. This project was carried out with the help of the ATP program-system developed in course of our research.

The future plant is to be based on the model of a US plant working in a desert belt. The location in Hungary, however, is surrounded by inhabited settlements, therefore, in the interest of a safer forecast, the transmission calculations were prepared in circle representation as well, besides the most frequent wind direction. By circle representation modelling the concentration at every point of the map can be obtained in case the wind blows towards the examined point from the point of emission. Following the stipulations of the Hungarian Standards, also the level-differences in the surface were considered, by the help of the relevant relief digital map.

The composition and the volume of the emitted gas in course of the planned turbine testing will change, depending on the type of the turbine and the test loading. Testing would run exclusively during day-time, and the turbines will be in operation for 2 hours as an average, as per loading levels. On the effect of the high speed of the exhaust output and its high temperature, an effective stack-height of 160-300 meters develops. Consequently, in the case of a stack designed as 6-10 m high, on the basis of the calculations, the difference of some meters in the height would result a practically negligible (1-2%) difference in the effective stack-height and in the developing concentrations. According to the agreement with the Consigner, calculations regarding short averaging time were done based on the most unfavorable emission- and meteorological parameters, and regarding long averaging time, we calculated yearly dispersion. To determine the sphere of effect to be expected, we considered two methods. More than 50 different cases were examined, and the results were visualized in the form of tables, diagrams and on 30 digital maps annexed to the study.

Computations were done regarding three pollutants: nitrogen oxides (NO_X), sulphur dioxide (SO2) and carbon

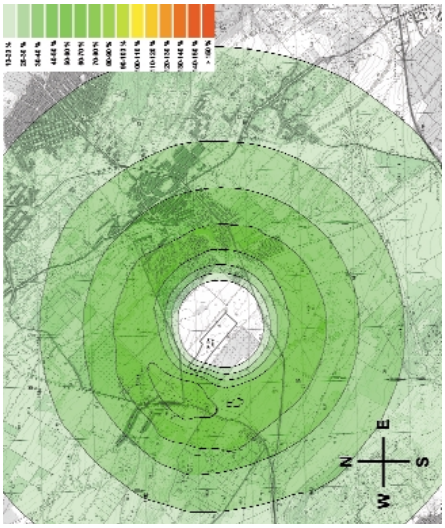

Figure 1: Isoconcentration zones in circle representation for the planned turbine testing plant, for the most frequent meteorological condition (9,7^{o} C air temperature, 2 m/s wind speed, B stability indicator), NO_X, emission: 362,2 kg/h, imission limit value: 150 µg/m3.

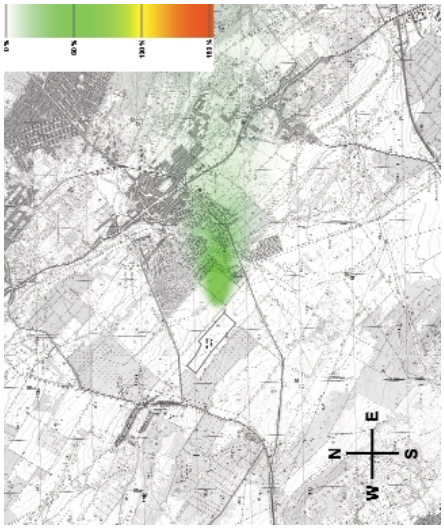

Figure 2: Concentration tail in W wind-direction for the planned turbine testing plant, for the critical meteorological condition (35^{o} C air temperature, 5 m/s wind speed, B stability indicator), NO_X, emission: 362,2 kg/h, imission limit value: 150 µg/m3.

monoxide (CO). Out of these, NO_X is the most significant. However, the ground-surface concentration to be expected, regarding this pollutant didn't attain or got at the limit value for short averaging time (30 min) stipulated by the Hungarian Standard, either in the most unfavorable meteorological state. The maximum value to be expected at populated areas is the concentration corresponding to the 65% of the limit value, but this concentration zone almost ends at the settlement, and is gradually reduced under 10% onsite. The developing maximum concentration, related to the limit value of 24 hours is the 1,5% of the limit value, and the yearly dispersion is almost zero, consequently, practically insignificant in both cases. The effects of SO_2 and CO emissions calculated for short averaging time are insignificant, and in daily and yearly dispersion cannot be detected.

It derived from the results that:

- the increase in wind-speed results in the increase of the developing maximal concentration, and at the same time, the significant attenuation of the smoke-tail departing from the stack. This is due to the fact that in case of rather intense wind blow, the effective stack height is lower, ie, the smoke-tail precipitates sooner
- a smoke-tail of bigger size (wider and longer) belongs to smaller wind-speeds
- an increase of 15^{o}C in the temperature results in a concentration increase of 1-2% in relation to the limit value, thus air-temperature makes influence on the concentrations developing on the ground-surface, only to a small extent
- by the increase of the stability indicator of the meteorological state, the developing maximum concentration decreases
- a wind-speed increase of 3 m/s has a bigger influence on the concentrations (intensifying them) than the increase of the Pasquill stability indicator by a unit (decreases the concentrations).

Figure 1 shows NO_X in circle representation, developing in the most frequent meteorological condition. The deviation of the isoconcentration zones from the regular concentric circles derives from differences in the altitude. The isoconcentration zones increase with concentrations corresponding to the 10% of limit value regarding air quality, defined by the Hungarian Standard. Consequently, the value of the concentration in the innermost and the outermost zones is between the 10-20% of limit value, and 50-60% in the maximum concentration zone, in W-NW direction from the stack. The stack position is denoted (indicated) by the point marked inside the concentration zones, on the site of the planned plant.

Figure 2 shows the concentration tail developing in case of W wind-direction and in critical meteorological condition. The modelling calculations were made for maximum NOx emission, in both figures. Consequently, the concentrations in the figures show the highest values to be expected, based on the calculations.

This research was supported in part by the Hungarian Research Foundation.

Links:
Department and Laboratory of Operations Research and Decision Systems, SZTAKI: http://www.oplab.sztaki.hu

Please contact:
Tamás Rapcsák — SZTAKI
Tel: +36 1 209 5266
E-mail: rapcsak@oplab.sztaki.hu

ERCIM NEWS

ERCIM News is the magazine of ERCIM. Published quarterly, the newsletter reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. ERCIM News has a circulation of 7000 copies.

ERCIM News online edition is available at http://www.ercim.org/publication/ERCIM_News/

ERCIM News is published by ERCIM EEIG, BP 93, F-06902 Sophia-Antipolis Cedex
Tel: +33 4 9238 5010, E-mail: office@ercim.org
ISSN 0926-4981

Director: Bernard Larrouturou
Central Editor:
Peter Kunz
E-mail: peter.kunz@ercim.org
Tel: +33 1 3963 5040
Local Editors:
CLRC: Martin Prime
E-mail: M.J.Prime@rl.ac.uk
Tel: +44 1235 44 6555
CRCIM: Michal Haindl
E-mail: haindl@utia.cas.cz
Tel: +420 2 6605 2350
CWI: Henk Nieland
E-mail: henkn@cwi.nl
Tel: +31 20 592 4092
CNR: Carol Peters
E-mail: carol@iei.pi.cnr.it
Tel: +39 050 315 2897
FORTH: Constantine Stephanidis
E-mail: cs@csi.forth.gr
Tel: +30 81 39 17 41
GMD: Dietrich Stobik
E-mail: stobik@gmd.de
Tel: +49 2241 14 2509
INRIA: Bernard Hidoine
E-mail: bernard.hidoine@inria.fr
Tel: +33 1 3963 5484
SARIT: Harry Rudin
Tel: +41 1 720 3707
E-mail: hrudin@smile.ch
SICS: Kersti Hedman
E-mail: kersti@sics.se
Tel: +46 8633 1508
SRCIM: Gabriela Andrejkova
E-mail: andrejk@kosice.upjs.sk
Tel: +421 95 622 1128
SZTAKI: Erzsébet Csuhaj-Varjú
E-mail: csuhaj@sztaki.hu
Tel: +36 1 209 6990
TCD: Carol O'Sullivan
E-mail: Carol.OSullivan@cs.tcd.ie
Tel: +353 1 60 812 20
VTT: Pia-Maria Linden-Linna
E-mail: pia-maria.linden-linna@vtt.fi
Tel: +358 9 456 4501

Free subscription

You can subscribe to ERCIM News free of charge by:

- sending e-mail to your local editor
- contacting the ERCIM office (see address above)
- filling out the form at the ERCIM website at http://www.ercim.org/

Updating Abramowitz & Stegun

by Nico Temme

A project is underway at the US National Institute of Standards and Technology (NIST) to develop a replacement for the Handbook of Mathematical Functions, commonly known as 'Abramowitz & Stegun', after its principal editors. This will be a major new mathematical reference source on the World Wide Web for special functions and their applications.

The Handbook, published in 1964 by the National Bureau of Standards, is probably the most widely-used book on special functions, and perhaps even in mathematics at large. This book, whose sales number in the hundreds of thousands, may well be the most frequently cited source of all mathematical reference works. A paperback version has been photo-offset and sold by Dover Publications since 1965. Although Dover does not reveal sales data, it undoubtedly outsells the government edition many times over.

The Handbook has never been revised, although there have been numerous advances in basic theory, computational methods, and domains of application since its publication. Its structure is that of a static reference volume; though still very useful, it does not meet the needs of modern users for information that can be conveniently exploited in a highly computer-oriented technical environment. Such needs have been communicated regularly to NIST (formerly, National Bureau of Standards). About five years ago it was decided at NIST that a successor to the NBS Handbook should be designed in the form of a knowledge base, called the Digital Library of Mathematical Functions (DLMF). Via free access over the World Wide Web, using standard Web browsers, individuals will be able to obtain validated mathematical information in a semantic-based representation that incorporates metadata, interactive features, and extensive linkages. The Web server will be constructed, maintained and operated by NIST. The many tables in the existing version of the Handbook will not be present in the new version. They are inadequate to current needs, particularly when used to validate numerical software. Second, software packages exist that can compute vastly extended numerical tables. New chapters will be:

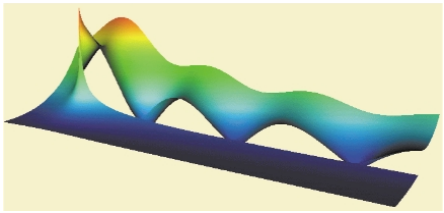

***$H_0^{(1)}(z)$:* Hankel function of the first kind and of order zero.**

- Asymptotic Approximations
- Computer Algebra
- q-Hypergeometric Functions
- Painlevé Transcendents
- Asymptotics of Integrals with Coalescing Saddles
- Wavelets
- 3j, 6j, 9j Symbols

CWI is involved in this project. Nico Temme, a specialist on Asymptotics and Special Functions, is a member of the Editorial Board, and will write several chapters (on exponential integrals, error functions and parabolic cylinder functions). The Executive Committee for the whole project consists of the four Principal Editors: Daniel W. Lozier, Frank W. J. Olver, Charles W. Clark, and Ronald F. Boisvert. The Board of Associate Editors: R. A. Askey, M. V. Berry, W.Gautschi, L. C. Maximon, M. Newman, I. Olkin, P. Paule, Johannes Kepler, W. P. Reinhardt, N. M. Temme, and J. Wimp. The project is expected to be completed by 2003.

Link:
http://dlmf.nist.gov/

Please contact:
Nico Temme — CWI
Tel: +31 20 592 4240
E-mail: nicot@cwi.nl

CALL FOR PARTICIPATION

First IFIP Conference on e-commerce, e-business, e-government

Zurich, Switzerland, 3-5 October 2001

This conference provides a forum for users, engineers, and scientists in academia, industry, and government to hear and discuss the latest findings in these e-application areas and the underlying technology.

From the 130 manuscripts submitted, the program committee has selected 33 papers for the conference's main track. Additional 'minitracks' will cover the topics of e-government, e-democracy, models for e-business, re-engineering for e-commerce, and mobile-commerce. These presentations will take place on Thursday and Friday. On Wednesday of the same week, a number of tutorials will be presented. Please see the URL below for details.

Further information:
http://www.ifi.unizh.ch/I3E-conference

CALL FOR PARTICIPATION

DELOS Workshop on Interoperability in Digital Libraries

Darmstadt, Germany, 8-9 September 2001

The workshop is intended to bring together researchers and developers working on digital libraries and related areas for in-depth analysis and discussion of new models, theories, frameworks, and solutions to interoperability in digital libraries. Interoperability and mediation in distributed, heterogeneous digital libraries require a middleware that provides transparent access to inherently heterogeneous digital collections. Among others, such a middleware must support data translation between different data types, representations, detection of same-objects, and data propagation for ensuring global consistency. There are research and industrial efforts under way to develop techniques for the modeling, creation, and management of metadata and ontologies for integrating and exchanging content. Transparent search (and metasearch) and consolidation of information from multiple digital libraries pose many challenges beyond the typical problems encountered in Web search and distributed heterogeneous databases. This workshop will provide a forum for discussing these problems and the approaches that are being taken to address them.

The two-day workshop will have a mix of invited and contributed presentations from academia and industry. In addition to presentations, there will be ample time for moderated discussion and panel discussions.

Further information:
http://www.ercim.org/delos/

CALL FOR PARTICIPATION

cast01—Conference on Communication of Art, Science and Technology: Living in Mixed Realities

Sankt Augustin, Germany, 21-22 September 2001

cast01 is the first in a series of international conferences on communication of art, science and technology. The conference theme 'Living in Mixed Realities' refers to topics like: Semantic Web, Mixed Reality, Advanced Interfaces and Future Media Spaces that symbolise the influence of information technology on patterns of life and work in a networked society. cast01 aims at establishing an ongoing forum for interdisciplinary exchange by confronting ground breaking media art with inspiring research projects and cuting-edge technological development. The conference topics include:

- Agents and Narrative Intelligence
- Artistic Productions of Mixed Reality Environments
- Awareness, Memory Space and Knowledge Discovery
- Cultural Archives
- Distributed Systems and Parallel Architectures for the Web
- Hypermedia Formats (XML, VRML, MPEG-4, MPEG-7)
- Interactive TV
- Mixed Reality Architectures
- Performative Interfaces
- Tracking, Tracing, Vision Systems.

Accepted papers and blueprints will be published in the Conference proceedings. A special issue of netzspannung.org journal of Art, Design and Innovation Research will be published with cast01 conference best papers. cast01 is organised by netzspannung.org and by GMD. It is supported by the German Federal Ministry for Education and Research and by the European Commission.

Further information:
http://netzspannung.org/cast01/

CALL FOR PARTICIPATION

ECDL 2001—5th European Conference on Research and Advanced Technology for Digital Libraries

Darmstadt, Germany, 4-8 September 2001

ECDL is the fifth in a series of annual European digital library conferences on current technological developments and innovative applications in Digital Libraries and their impact on science, technology and society.

This year, the conference will explore the on-going convergence of libraries, archives, and museums into integrated digital information spaces and collaborative work environments. The program will consist of 38 paper presentations, three invited talks, six full-day workshops, four half-day tutorials, and several panels. Sponsors of this year's conference are Technische Universität Darmstadt, Die Deutsche Bibliothek, ERCIM, DELOS Network of Excellence, the European Commission's Human Potential Programme, and GMD.

Further information:
http://www.ecdl2001.org/

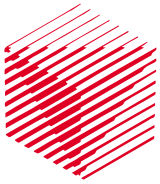

ERCIM – The European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development, in information technology and applied mathematics. Its national member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.

Central Laboratory of the Research Councils
Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, United Kingdom
Tel: +44 1235 82 1900, Fax: +44 1235 44 5385
http://www.cclrc.ac.uk/

Consiglio Nazionale delle Ricerche, IEI-CNR
Area della Ricerca CNR di Pisa, Via G. Moruzzi 1, 56124 Pisa, Italy
Tel: +39 050 315 2878, Fax: +39 050 315 2810
http://www.iei.pi.cnr.it/

Czech Research Consortium for Informatics and Mathematics
FI MU, Botanicka 68a, CZ-602 00 Brno, Czech Republic
Tel: +420 2 688 4669, Fax: +420 2 688 4903
http://www.utia.cas.cz/CRCIM/home.html

Centrum voor Wiskunde en Informatica
Kruislaan 413, NL-1098 SJ Amsterdam, The Netherlands
Tel: +31 20 592 9333, Fax: +31 20 592 4199
http://www.cwi.nl/

Foundation for Research and Technology – Hellas
Institute of Computer Science, P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
Tel: +30 81 39 16 00, Fax: +30 81 39 16 01
http://www.ics.forth.gr/

GMD – Forschungszentrum Informationstechnik GmbH
Schloß Birlinghoven, D-53754 Sankt Augustin, Germany
Tel: +49 2241 14 0, Fax: +49 2241 14 2889
http://www.gmd.de/

Institut National de Recherche en Informatique et en Automatique
B.P. 105, F-78153 Le Chesnay, France
Tel: +33 1 3963 5511, Fax: +33 1 3963 5330
http://www.inria.fr/

Swedish Institute of Computer Science
Box 1263, SE-164 29 Kista, Sweden
Tel: +46 8 633 1500, Fax: +46 8 751 72 30
http://www.sics.se/

Swiss Association for Research in Information Technology
Dept. Informatik, ETH-Zentrum, CH-8092 Zürich, Switzerland
Tel: +41 1 632 72 41, Fax: +41 1 632 11 72
http://www.sarit.ch/

Slovak Research Consortium for Informatics and Mathematics
Dept.of Computer Science, Comenius University, Mlynska Dolina M SK-84248 Bratislava, Slovakia
Tel: +421 7 726635, Fax: +421 7 727041
http://www.srcim.sk

Magyar Tudományos Akadémia – Számítástechnikai és Automatizálási Kutató Intézete
P.O. Box 63, H-1518 Budapest, Hungary
Tel: +36 1 4665644, Fax: + 36 1 466 7503
http://www.sztaki.hu/

Trinity College, Department of Computer Science,
Dublin 2, Ireland
Tel: +353 1 608 1765, Fax: 353 1 677 2204
http://www.cs.tcd.ie/ERCIM

Technical Research Centre of Finland
VTT Information Technology, P.O. Box 1200, FIN-02044 VTT, Finland
Tel:+358 9 456 6041, Fax :+358 9 456 6027
http://www.vtt.fi/