

European Commission, DG Information Society

Future and Emerging Technologies (FET)

A report on

The FET Global Computing Initiative

July 2005

Table of Contents

Acknowledgements

This document is based on the report on the third year reviews produced by Kurt Geihs (Universität Kassel, Germany, kurt@geihs.de), previous reports by Jan van Leeuwen (Universiteit Utrecht, The Netherlands) and David Pearce (Universidad Rey Juan Carlos, Madrid, Spain), and specific project documents, including review reports. We also wish to acknowledge the work of the project officer who launched this initiative, Leonardo Flores Añover.

Fabrizio Sestini fabrizio.sestini@cec.eu.int
Wide Hogenhout wide.hogenhout@cec.eu.int
European Commission, DG Information Society, Future and Emerging Technologies

Executive Summary

The Global Computing (GC) initiative was launched in 2001 by the Future and Emerging Technologies (FET) unit within the Information Society Technologies (IST) programme of the European Commission.

The mission of this initiative was to support high-quality research that develops models, frameworks, methods and algorithms for constructing large-scale open distributed systems. Projects have addressed the fundamental challenges posed by rapidly evolving systems composed of very large numbers of autonomous, interacting and mobile computational entities.

The initiative funded 13 RTD projects for a total contribution of 18.2 M€. Other FET projects contributed to the same topic, for an additional 6 M€ funding. The projects addressed a wide spectrum of subjects, such as formal calculi and correctness proofs for distributed computing, system models and model checking, peer-to-peer and overlay computers, security and trust management, as well as mechanism design for resource allocation.

The GC initiative brought together most of the best European theoretical computer science groups, complemented by strong teams in agents, databases, and systems and network engineering. Scientific output in terms of publications, PhD theses, workshops etc. was found to be very high. Results were at the forefront of the state of the art, and in certain areas major breakthroughs were achieved (e.g. using type theory for security proofs). Overall there is a high awareness of activities and results within other projects. This has led to an active, open and cooperative exchange between projects. On the whole, the GC initiative has succeeded in building a strong and visible European research community focussed on theoretical aspects of Global Computing, in the realm of theoretical computer science.

While the impact of GC was remarkable in several areas of theoretical computer science, less impact could be noted in respect to systems building and practical exploitation. This is one focus of the follow-on Global Computing II initiative in FP6, whose projects have started in September 2005. Overall, GC has been an incubator for many follow-on research activities.

Introduction

FET (Future and Emerging Technologies) is the IST Programme nursery of novel and emerging scientific ideas. Its mission is to promote research that is of a long-term nature or involves particularly high risks, compensated by the potential of a high societal or industrial impact. FET goals are achieved via an "open" scheme and a "proactive" scheme. While the open scheme is constantly open to proposals for bold and visionary research on any subject related to IST, the proactive scheme is focused on a small number of proactive initiatives. These are strategic areas identified by the EC (in consultation with the main stakeholders) as holding particular promise for the future, in order to open new possibilities and set new trends for future research programmes in Information Society Technologies.

FET launched the Global Computing (GC) proactive initiative in 2001, within the Information Society Technologies (IST) priority of the 5th Framework Programme of Community Research (FP5) of the European Commission. In order to illustrate the thematic context of Global Computing, Figure 1 shows the pro-active initiatives of FET in FP5.

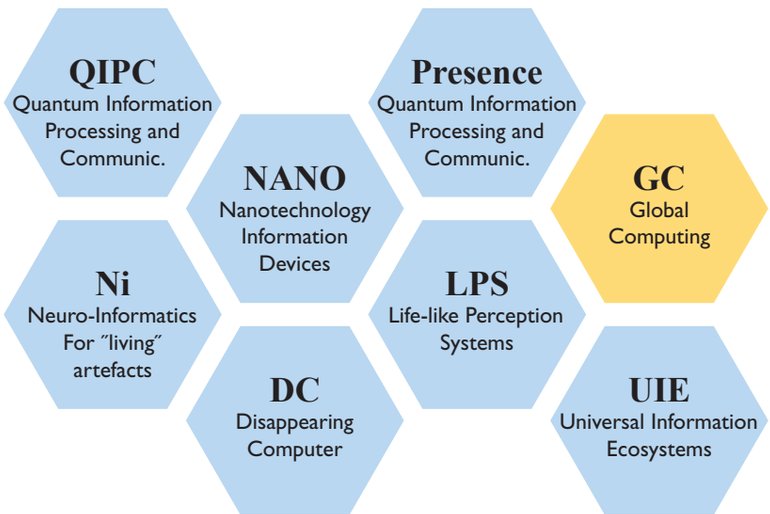

Fig.1: FET pro-active initiatives in FP5

In FP6 additional proactive initiatives were created, such as *Beyond Robotics*, *Complex Systems*, *Disappearing Computer II*, *Advanced Computing Architectures, Situated and Autonomic Communications*, and the follow-up *Global Computing II*, whose call was published in June 2004, for projects starting in the second half of 2005.

1.1. Objectives

The overall mission of the Global Computing initiative in FP5 was to support high-quality research that develops models, frameworks, methods, and algorithms for constructing large-scale open distributed systems. Its detailed objectives were based on the recommendations of a consultation workshop which was held in September 2000. Projects were requested to address the fundamental challenges posed by rapidly evolving systems composed of very large numbers of autonomous, interacting and mobile computational entities. In these global systems, activity

is not centrally controlled, the configuration varies over time, and the information about the environment is incomplete. The ultimate goal was to provide a solid scientific foundation for understanding, analysing and designing such systems so that they are flexible, dependable, secure, robust and efficient.

1.2. Project Set-up

The call for projects was published in February 2001. The initiative funded 13 RTD projects for a total funding of 18.2 M€, representing around 6 % of the programme's total effort in FET in FP5. The GC projects began on January 1st, 2002, and lasted nearly three years and a half.

1.3. Annual reviews and jamborees

All the projects in this initiative underwent clustered annual project reviews, grouped in a period of a few days and joined with public events dedicated to the community, so to favour collaboration and cross-fertilisation between different projects. The first two main GC reviews were held in March 2003 and March 2004, both in Rovereto (Trento, Italy); in 2003 some projects were reviewed in Cyprus. The third year review took place at the University of Edinburgh (UK), in April 2005, co-located with the ETAPS'05 conference.

Most of the GC project representatives also attended the public Workshop on *Trustworthy Global Computing*, organized as part of the ETAPS'05 conference programme. The workshop presented and discussed further insights in Global Computing, in particular introducing the three projects which had been selected in the call on Global Computing II in 2005.

2. Overview of Progress in Projects and Clusters

The 13 projects of the GC initiative addressed a broad spectrum of subjects in line with the general GC objectives. They were grouped initially in three thematic clusters: *Foundations of Networks and large Distributed Environments*, *Analysis of Systems* and *Languages and Programming Environments*. In this section, we give a brief summary of the projects' objectives and their major achievements in year 3, grouped according to the clusters. We also point to general project success stories and highlights. Further details and the projects' web sites can be found in the annexed project fiches.

2.1. Foundations of Networks and large Distributed Environments

This cluster was meant to lay the foundations of global computing infrastructures. Its projects were dealing with issues such as scheduling algorithms and mechanism design for resource sharing between autonomously interacting agents, use of type-theoretic approaches for the specification and verification of distributed computing systems, and other paradigms related to the reasoning about properties of such systems. The projects in this cluster were:

CRESSCO (Critical Resource Sharing for Cooperation in Complex Systems), conducting research on foundational aspects of managing critical resources (e.g. bandwidth, frequency, energy, processor time) in GC infrastructures connecting very large numbers of independent, possibly mobile and selfish agent entities.

CRESSCO has been one of the first research projects (together with FLAGS) that applied economic game and auction theories to the question of resource sharing in distributed computing systems, which has led to fundamental insights into resource management with monetary incentives.

DBGLOBE (A Data-Centric Approach to Global Computing), aiming at extending current database technology in order to address data management requirements in large-scale networks of mobile entities using a data-centric and service-oriented approach to GC.

The project developed a new powerful language for integrating documents and queries, called Active XML, a major success according to the project reviewers. There are links from the AXML developers to the W3C XQUERY group which might lead to a take-up of AXML features in the XQUERY standard.

FLAGS (Foundational Aspects of Global Computing Systems), whose objective is to provide a general set of design principles and mechanisms in order to facilitate the construction of global computing systems with autonomous and selfish agents competing for resources.

As CRESSCO, FLAGS also applied economic game and auction theories to the question of resource sharing in distributed computing systems. By doing so, it has achieved an impressive set of foundational results for a wide range of design aspects of GC systems.

SOCS (Societies of Computees), investigating computational and logical models for describing, analyzing, and verifying properties of individual agents as well as of societies of agents. The main focus is on formal models that identify core functionalities that the entities need to have in a GC environment.

The project has shown that logic programming techniques simplify considerably the proofs of properties of global computing systems.

2.2. Analysis of Systems

The focus of this cluster has been on the use of formal techniques such as type theory, proof-carrying code, and formal models for trust management in order to validate system properties such as correctness and security.

DART (Dynamic Assembly, Reconfiguration and Type Theory), developing formalisms and techniques for the modeling of the temporal dimension of GC systems in order to support arbitrary interleaving of type-checking, meta-programming, and normal computational activities while retaining safety.

This consortium has produced notable software prototypes that demonstrate practical applications of the formal methods, especially with respect to "smart recompilation". These may be of interest to tool vendors.

MRG (Mobile Resource Guarantees), extending the concept of Proof-Carrying Code (PCC) to include guarantees of resource usage, both in high-level source code, i.e. in a user-oriented programming language, and in low-level target code, e.g. a virtual machine's byte code.

Although small in size, MRG has made a relatively large contribution in the realm of PCC-based system analysis and verification. A theoretically well founded, comprehensive PCC infrastructure as well as an innovative framework for formal certificates of resource consumption were developed. The project has increased substantially the insights into the power and usefulness of this technology.

MYTHS (Models and Types for Security in Mobile Distributed Systems), exploring and developing type-based theories of security for mobile and distributed systems.

This has been a small project with outstanding results; according to the reviewers it achieved a breakthrough in the development of type-based theories enforcing the security of GC systems and applications.

PROFUNDIS (Proofs of Functionality for Mobile Distributed Systems), conducting research on formal modelling and verification techniques to explore key issues in mobile distributed systems, such as security protocols, authentication, access rights and resource management, as well as automatic support for the design and implementation of such features.

The project has been very successful in pushing the state of the art in the formal analysis of mobile distributed systems. The tools for the verification and testing of security protocols are considered valuable assets that can be used in other environments.

SECURE (Secure Environment for Collaboration among Ubiquitous Roaming Entities), developing a formal model of security based on the notion of trust as well as algorithms for the dynamic and self-configuring management of trust.

Project results include formal models of trust and trust evolution, which have been the basis for extensive practical design, implementation and testing activities. It is quite likely that the SECURE trust engine can be transferred into middleware systems for mobile and ad-hoc communication environments.

2.3. Languages and Programming Environments

While the above two clusters have concentrated primarily on formal techniques and proofs, this third cluster has used and extended these theoretical foundations to come up with software tools and frameworks for building global computing infrastructures.

AGILE (Architectures for Mobility), exploring an UML-based architectural approach to software engineering for mobile systems based on a uniform mathematical framework. This includes theoretical foundations as well as pragmatic techniques for designing mobile computing systems by focusing on software architecture and corresponding UML models.

This project demonstrated the power and elegance of UML-based specification techniques for mobile computing that are based on sound algebraic foundations. Model-checkers and design tools can be further exploited in future research and practice.

DEGAS (Design Environments for Global Applications), deriving specific models such as process algebras and performance models from annotated high-level application models in

UML in order to support the design and implementation of global computing applications. It has also implemented a set of tools to mechanise the process, including the environment "Choreographer", which has been impressively demonstrated during the last review.

Commercial exploitation seems feasible for the Choreographer tools. The industrial partner (i.e. Motorola Spain) will continue to use the tools for commercial products where applicable.

One of the researchers in this project, Dr. Jane Hillston (University of Edinburgh), received the Roger Needham Award for her research achievements. In her award lecture and in other invited lectures she has presented results of DEGAS to the general public.

MIKADO (Mobile Calculi Based on Domains), developing new formal programming models for global computing based on the concept of process mobility and the notion of "domain". The formal models are used to build new prototype programming languages and runtime environments.

Several calculi for distributed and mobile computing systems have been obtained, whereby these calculi were implemented using a new generic implementation framework.

PEPITO (Peer-to-Peer Implementation and Theory), investigating the foundations of scalable distributed computing based on the peer-to-peer (P2P) computing paradigm. PEPITO explored P2P algorithms, programming language features and a generic, language-independent distribution platform for P2P computing.

The project provided a practical proof that server-less distributed systems are feasible and can provide very similar services as conventional, client/server systems. Some of the demo applications have raised interest for commercial take-up.

3. Progress and Impact of the Initiative

Three years of research on GC have produced an enormous amount of high-class scientific results. Their implications as for impact and community building are discussed in detail in this Section.

3.1. Project Interactions and Community Building

After a gradual start, in the last year collaborations between GC projects have become intensive. Joint research leading to joint publications, exchange of PhD students between project partners, visits of researchers to other project sites, joint workshops, exchange of knowledge, sharing of application scenarios etc. have all been stepped up continuously during the course of the initiative.

Figure 2 shows the links created between the GC projects in the initiative. A link means that collaboration had been established, as reported in the third year review meetings and by members of the project teams. The diagram does not reflect simple overlaps in technical goals or technical approaches, because this is assumed to be normal for such a thematic initiative.

The clustered reviews and joint activities permitted to create mutual awareness of activities and results among all GC projects. This has led to an active and open cooperation between projects which was found of considerable benefit to the whole initiative (also resulting in saving of resources).

As the diagram illustrates, there were as many inter- as intra-cluster collaborations. This suggests that eventually collaborations did not follow the initial project clustering patterns.

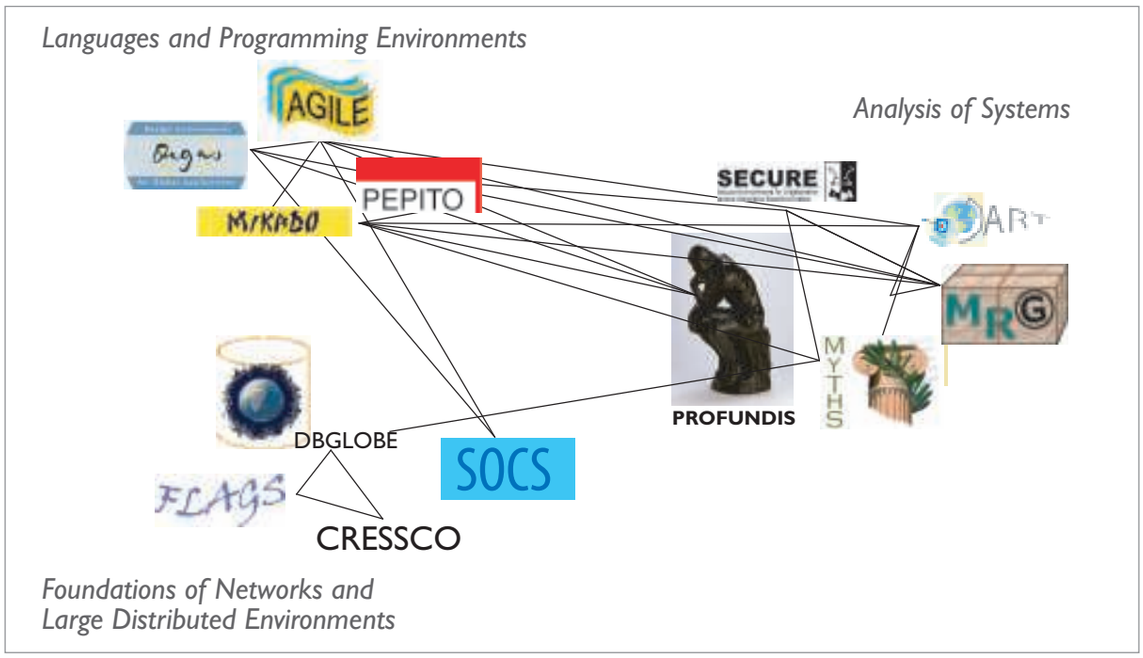

Fig.2: Project interactions in GC (as reported at the 3rd review)

Some collaboration was clearly due to pre-existing personal working relationships between participants in the initiative. It is also interesting to note that the number of links is not proportional to project size in terms of partners or subject spectrum. There are small projects, e.g. MRG with 2 partners and MYTHS with 3 partners that both have many links to other projects. This appears in line with the view that large projects can progress well even if "self-contained" whereas smaller initiatives need more frequent contacts with external partners and activities.

Overall, GC has led to a network of excellent research groups working on the foundations of distributed computing systems.

3.2. Community building

The interaction diagram in Figure 2 above demonstrates nicely that at the end of the initiative there is a rather strong overall sense of community in GC, which was not present in the beginning. Within this GC community there seem to be two or three "inner circles" of participants that collaborate very intensively, and will continue to do so in future research projects, e.g. in the GC 2 initiative.

The GC initiative was defined with the aim to create a *"solid scientific foundation for the design of such [GC] systems, and to lay the groundwork for achieving effective principles for building and analysing such systems."* The thirteen GC projects have collectively been extremely successful in achieving these goals, as demonstrated by the number of publications, books, workshops and PhD theses (see 3.3 and the individual projects' fiches in annex). According to the reviewer's report, "quantity and quality of the theoretical results are extraordinary."

The original GC Call for Proposals mentioned some particular subjects that projects should explore i.e.

1. *the design of GC systems,*
2. *analysing and reasoning about their behaviour,*
3. *avoiding and/or detecting undesirable behaviour through control of the system and/or its environment, and*
4. *understanding their limits and potential.*

Virtually all the 13 GC projects have addressed these issues. The main difference between the projects lies in (a) the underlying application or architectural focus (e.g. mobile systems, Peer-to-Peer, wireless communications etc.) and (b) the adopted methodology. A predominant method to analyse, reason about and control the behaviour of GC systems has been *type theory*. DART, MIKADO, MRG, and MYTHS have explicitly used this approach. *Algebraic models* of distributed systems have also been an issue of major interest in most projects. A third major concern was the design and analysis of *algorithms* for e.g. resource sharing among selfish agents, information management in peer-to-peer systems, trust management and propagation, etc. In these three areas the GC projects have collectively contributed to establish and extend the leading role of European researchers.

Orthogonal to the methodological viewpoint, *security* of global computing systems has been a recurring subject in several projects, i.e. in DEGAS, MRG, MYTHS, PROFUNDIS, and of course SECURE. The importance of this subject and the contributions of GC projects to this research area have been acknowledged through many publications in highly respected conferences and journals, and most notably through the Roger Needham Award that was awarded to a researcher in the DEGAS project, as reported above. Also notable are the contributions to the international state of the art from MRG related to Proof-Carrying Code (PCC) (applied to Java bytecode programs and resource bounds).

Building on these results of GC1, security has become a central topic in the GC2 project MOBIUS. This project aims at building a security architecture for global computers (based on ideas from PCC), providing innovative trust management, static enforcement mechanisms and support for system component downloading.

3.1. Dissemination

Dissemination of results through academic publication channels has been excellent. Publishing more than 1,000 peer-reviewed papers in journals and conferences (including many very prestigious titles), as well as about 15 books and book chapters, about 50 open workshops, and more than 50 PhD theses is truly outstanding.

Project web sites have substantially improved over the course of the initiative's lifetime, mainly due to reviewers' recommendations. Generally, web pages are informative and contain links to experimental results (code, graphics) and source code. They are a useful reference for other researchers, and they are quite easy to find using appropriate keywords.

By outsiders of GC the results of the initiative typically are not directly attributed to the term *Global Computing*. This seems to be a kind of general "marketing problem" for the GC initiative. Future activities should pay more attention to this issue. For example, joint public workshops under the umbrella of popular conferences and overview papers in well-respected journals specifically addressing the objectives and results of the initiative should be part of the dissemination activities. The initiative is missing a separate Specific Support Action (SSA) performing such kind of marketing, organizing and managing the various public activities.

3.2. Impact and Exploitation

Although the GC projects had a significant impact on the state of the art in theoretical research, in terms of commercial exploitation and take-up they have not yet provided very visible results. Contacts with industry have been limited, as is the number of directly exploitable results. A notable exception is project DEGAS, where a take-up of tools by an industrial partner is under negotiation.

Indeed GC projects have followed a basic research approach, whose real significance and value will surface later when the theoretical results will be taken up by other researchers as a foundation for new system designs and verification techniques. In this respect the choice of many projects to make infrastructure support software available on a non-commercial basis is deliberate in order to speed up the acceptance and development of these technologies, and not uncommon for this kind of effort. It is unfortunate that the European contribution to the world-wide efforts on non-commercial software is so often hidden, since the focus is on commercial uptake in a very concrete sense.

3.3. Future Research

The GC initiative has opened up many new research challenges. Besides the continuation and extension of the GC research subjects, other threads of future activities seem possible:

1. Experimentation with prototypes and building systems

   The GC projects were meant to address mainly foundational research issues. It would now be a logical step to harvest the excellent theoretical results achieved in the first three years of GC and to shift the attention to a more practical type of research with prototype building, experimentation, and practical evaluation.

2. Application of GC results to new application scenarios

   Within three years of GC, technological progress has generated new research areas of general interest, e.g. sensor networks, self-* systems, service-oriented computing, identity computing etc. Without any doubt many of the GC results can be applied or extended to these areas.

4. Global Computing in FET and Links with Other Initiatives

Similar work to GC has not been carried out elsewhere in FET or more broadly in IST. The *GRID* and *Disappearing Computer* initiatives also address distributed computing but in different contexts and with different aims: the GRID projects focus on large scale resource sharing, while DC projects emphasize pervasive computing and ambient intelligence techniques and applications.

4.1. Global Computing II

The *Global Computing II* (GC II) initiative is a follow-up of GC. It addresses many of the research challenges listed above in Section 4.5, building on the successes of GC projects. GC II projects are not a mere continuation of GC topics, but try to position themselves towards new trends and research subjects such as service-oriented architectures and sensor networks.

There are three Integrated Projects selected in GC II, having the following objectives:

AEOLUS
• Develop algorithmic principles to enable transparent and efficient access to the resources of an internet-based global computer. • Address also communication between wireless and mobile nodes. • Build sample services and applications as proof-of-concept.
MOBIUS
Investigate trust and security for small devices as parts of global computers.
Use the Proof Carrying Code paradigm.
Perform Java-enabled global computer case studies.
SENSORIA
Develop an engineering methodology for service-oriented overlay computers.
Build context-adaptive, personalisable global services.
Apply results to e-business, automotive systems, telecommunications.

GC II funds 3 Integrated Projects involving a total of 53 partners from 18 European countries with a total funding of approximately 20 M€. As a comparison, the 13 GC projects involved 68 partners from 15 countries (including US).

The figure is a reflection of the relations between the three projects.

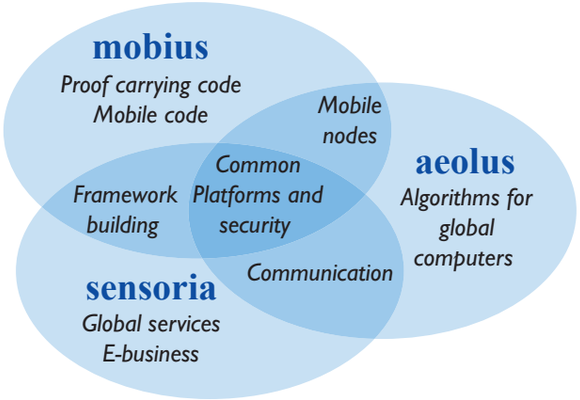

Country	Institution	AGILE	CRESCCO	DART	DBGLOBE	DEGAS	FLAGS	MIKADO	MRG	MYTHS	PEPITO	PROFUNDIS	SECURE	SOCS	SECSAFE	POP	ASAP	ITRUST	APPSEM II	TYPES	CATNETS	MOBIUS	AEOLUS	SENSORIA
Belgium	Katholieke Universiteit Leuven																						■	
Belgium	Universite Catholique de Louvain										■													
Cyprus	University of Cyprus				■		■							■									■	
Czech Republic	Universita Karlova v Praze																						■	
Denmark	Aarhus Universitet												■						■					
Denmark	Danmarks Tekniske Universitet					■									■									■
Denmark	Koebenhavns Universitet																		■					
Denmark	Roskilde Universitetscenter																■							
Estonia	Cybernetica AS																						■	
Estonia	Tallinn University of Technology																		■			■		
France	Centre National de la Recherche Scientifique		■																				■	
France	Ecole Normale Superieure Paris									■														
France	France Telecom							■														■		
France	Institut National de Recherche en Informatique et Automatique				■						■	■			■			■	■			■	■	
France	Trusted Logic S.A														■							■		
France	Universite de Nice - Sophia Antipolis		■																					
France	Universite de Paris XI																			■				
France	Universite Paris 7 Denis Diderot																		■	■				
Germany	Albert-Ludwigs-Universitaet Freiburg																		■					
Germany	Christian-Albrechts-Universitaet zu Kiel		■																				■	
Germany	FAST Gesellschaft für angewandte Softwaretechnologie mbH																							■
Germany	Ludwig-Maximilians-Universitaet Muenchen	■							■										■	■		■		
Germany	Max-Planck-Ges. z.Fd.W. represented by MPI für Informatik																						■	
Germany	RWTH Aachen																					■		
Germany	S&N AG																							■
Germany	SAP Aktiengesellschaft																					■		
Germany	Technische Universitaet Darmstadt																		■					
Germany	Technische Universitaet Muenchen																			■				
Germany	Universitaet Bayreuth																				■			
Germany	Universitaet Dortmund																	■						
Germany	Universitaet Karlsruhe																				■			
Germany	Universitaet Paderborn						■																■	
Greece	Athens University of Economics and Business				■																			
Greece	Intracom S.A.																	■						
Greece	National and Kapodistrian University of Athens						■																■	
Greece	Plefsis Information Systems S.A.																	■						
Greece	Research Academic Computer Technology Institute		■		■		■																■	
Greece	University of Crete																	■						
Greece	University of Ioannina				■																		■	
Greece	University of Patras		■													■							■	
Greece	Virtual Trip Ltd																	■						
Hungary	Budapest University of Technology and Economics																							■
Ireland	National University of Ireland, Dublin																					■		
Ireland	Trinity College Dublin												■					■						
Italy	Consiglio Nazionale delle Ricerche	■														■		■						■
Italy	Istituto Trentino di Cultura																				■			
Italy	Motorola Electronics S.P.A.					■																		
Italy	Omnys S.R.L.																							
Italy	Telecom Italia Learning Services S.p.A.																						■	
Italy	Telecom Italia SpA																							■
Italy	Universita Ca' Foscari di Venezia									■														
Italy	Università di Bologna													■										■
Italy	Università di Ferrara																							
Italy	Università di Firenze	■						■																■
Italy	Università di Genova			■															■					
Italy	Università di Padova																						■	
Italy	Università di Pisa	■				■						■		■					■					■
Italy	Università di Roma Tor Vergata		■																				■	
Italy	Università di Roma La Sapienza																							
Italy	Università di Salerno		■																				■	
Italy	Università di Torino			■																■				
Italy	Università di Trento																							
Italy	Università di Udine					■														■				■
Italy	Università Politecnica delle Marche																				■			
Netherlands	Stichting Katholieke Universiteit Nijmegen																			■		■		
Norway	Sintef																	■						
Norway	Universitet I Oslo																							
Poland	TLS-Technologie Spo ka z graniczon odpowiedzialno ci																					■		
Poland	Uniwersytet w Bialymstoku																			■				
Poland	Uniwersytet Warszawski	■																				■		■
Portugal	Atx Software, S.A.	■																						■
Portugal	Universidade de Lisboa							■																
Portugal	Universidade do Minho																		■					
Portugal	Universidade Nova de Lisboa											■												
Spain	Universidad Autonoma de Barcelona																	■						
Spain	Universidad Politecnica de Madrid																■					■		
Spain	Universitat Politecnica de Catalunya						■									■					■	■	■	
Sweden	Chalmers University of Technology																		■	■		■		
Sweden	Kungliga Tekniska Hoegskolan										■													
Sweden	Swedish Institute of Computer Science											■												
Sweden	Uppsala Universitet																							
Switzerland	Ecole Polytechnique Federale de Lausanne										■								■					
Switzerland	Eidgenössische Technische Hochschule Zürich																					■	■	
Switzerland	Universite de Geneve		■										■										■	
United Kingdom	City University London													■										
United Kingdom	Council for The Central Laboratory of The Research Councils																	■						
United Kingdom	Graham, Amanda Klyne, Nine by Nine																							
United Kingdom	Heriot-Watt University			■																				
United Kingdom	Hewlett-Packard Ltd														■			■						
United Kingdom	Imperial College of Science, Technology and Medicine			■										■								■		
United Kingdom	King`s College London																	■						
United Kingdom	London Software Systems																							■
United Kingdom	The University of Birmingham																		■					
United Kingdom	The University of Edinburgh					■			■										■			■		■
United Kingdom	The University of London, Queen Mary and Westfield College																	■	■					
United Kingdom	The University of Nottingham																		■					
United Kingdom	The University of Sussex							■		■														
United Kingdom	The Victoria University of Manchester																			■				
United Kingdom	University of Bristol																■							
United Kingdom	University of Cambridge										■		■						■	■				
United Kingdom	University of Leicester	■																						■
United Kingdom	University of Oxford																		■					
United Kingdom	University of Southampton																■							
United Kingdom	University of Strathclyde												■					■						
United Kingdom	University of Wales, Cardiff																				■			
United States	The Regents of the University of California				■																			

Table 1 – Participation in the GC related projects

4.2. Related IST FET OPEN projects

There are a number of links between GC projects and other EC projects that will be further exploited in future work. For example, the ongoing IST FET OPEN projects *CATNETS* (evaluation of the free-market paradigm for decentralized operation of dynamic application networks) and *ASAP* (advanced analysis and specialization for pervasive systems) are directly related to the general objectives of GC, and will partially build on GC results.

Other (completed or near to completion) FET OPEN research activities that are thematically or methodologically close to the projects of the GC initiative have nicely complemented the objectives of the GC initiative. The most notable examples are: APPSEM II (Thematic network on application-oriented semantics of programming languages), TYPES (Working group focusing on the development of languages and tools for formal reasoning based on type theory), SECSAFE (Research project developing methods and tools for the validation of the safety and security of Java-based smart card software), POP (Research project developing platform-independent extensions to the OpenMP programming model, especially for non-SMP architectures), and ITRUST (Working Group focusing on computational trust models in dynamic open systems integrating inter-disciplinary perspectives).

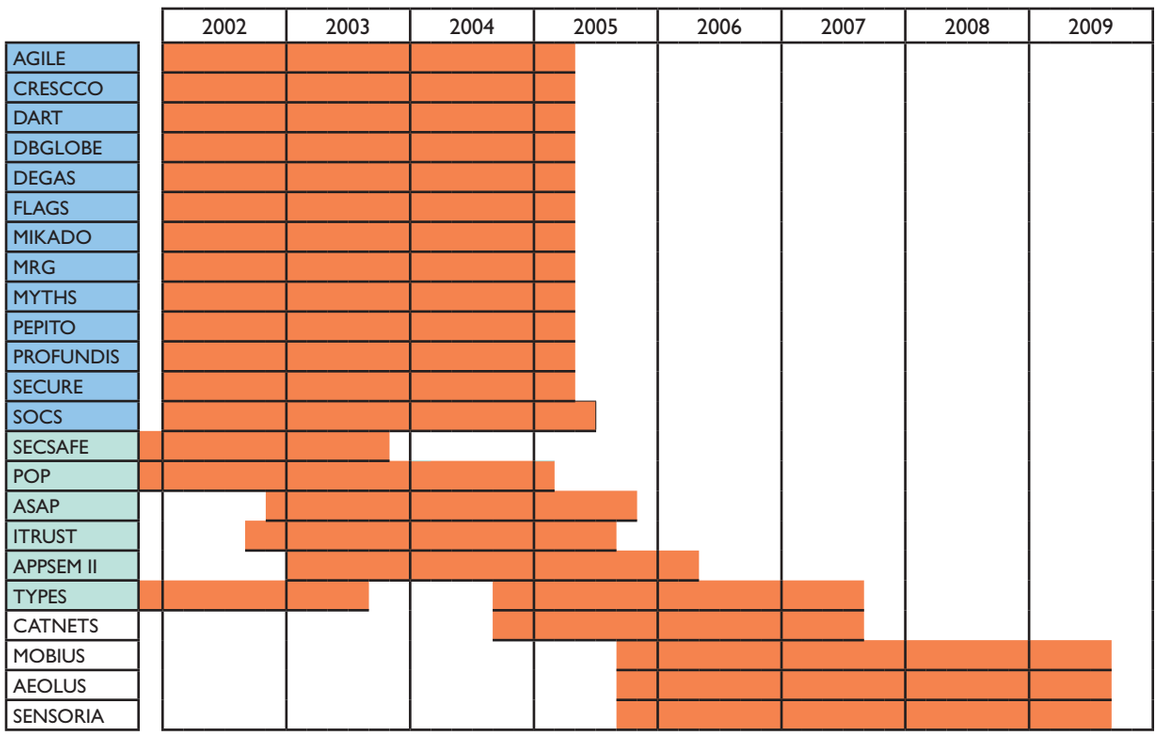

Duration of the different projects related to the GC initiative

All together, these additional FET OPEN projects received a funding of 6 million euro. The diagram above shows the duration of the different project related to the Global Computing initiative, whereas Table 1 lists the partners in all the GC related projects.

4.3. Other related national activities

GC projects have established a variety of direct links to national research activities. Representative examples include:

- In the UK, the EPSRC research agency has a number of programmes that are related to GC, e.g. the *e-science* initiative. The programme *Research in the Fundamental Computer Science for e-Science* has many similarities with GC but a slightly smaller funding. Furthermore, the *Grand Challenges for Computing Research* programme, sponsored by the UK Computing Research Committee with support from EPSRC and NeSC, addresses some of the GC topics.

- In Ireland, the Science Foundation of Ireland (SFI) is supporting some related research in its Information and Communications Technology division.

- In Italy, the MIUR_COFIN programme has funded / funds four projects that have direct links to GC, i.e. TOSCA (Concurrency Theory, Higher Order Languages and Types), NAPOLI (Network Aware Programming: Objects, Languages and Implementations), MEFISTO (Formal Methods for Security), and COMETA (Computational Meta-Models).

- In France, project partner CNRS of GC project MYTHS collaborated very successfully with the French INRIA project CRISTAL; they jointly produced a new programming language which has attracted industrial interest.

- In Switzerland, the university EPFL, partner in GC project PEPITO, has set off a *Global Computing Center* in which about 90 members of the school's institute of computer science are involved.

- In Sweden, planning for a *Center of Excellence on Global Computing* is underway. The Swedish research agency VINNOVA granted planning funds for a proposal by S. Haridi (GC project PEPITO) and some others.

4.4. Other related international activities

Viewed from an international perspective, GC was relevant and timely launched. Other funding agencies like the CISE division of the National Science Foundation (NSF) have also launched similar programmes:

- *Theory of Computing,* supporting fundamental research in core theory, fundamental algorithms, and application-specific theory (http://www.nsf.gov/pubs/progdesc/1998/cise/2860.htm)

- *Trusted Computing,* seeking to establish a scientific foundation for managing privacy and security (http://www.nsf.gov/pubs/2001/nsf01160/nsf01160.html)

- *Distributed Systems* focusing on novel concepts for the design of distributed systems (http://www.nsf.gov/pubs/progdesc/2003/cise/2876.html)

While both GC and these other international programmes have set new challenging research directions in the international arena, there were no formal links established between these initiatives. The GC project participants however seemed to be well aware of these activities, and had often the opportunity to cross-fertilise ideas and knowledge through their participation to international conferences and workshops.

5. Conclusions and recommendations

The Global Computing initiative has been very successful in exploring the theoretical foundations of global computing systems. It has changed substantially the research landscape. As noted above, major scientific breakthroughs and success stories were achieved, like in DEGAS, MYTHS and CRESSCO. However, for the type of research performed in GC it is fair to say that the real value and impact can only be judged after some time of exploring and applying the theoretical results. Research results take time to mature and to turn into real innovations.

The scientific output in terms of publications and PhD theses has been very high. A strongly networked community of researchers has emerged that collectively define the state of the art in the GC research areas. The initiative has opened up as many challenging new research questions as it was able to solve. More research is necessary. Overall, GC has been an incubator for many follow-on research activities, including GC II.

The enormous momentum existing at the end of GC should be preserved. In this respect, the GC II research partners are encouraged to create a Network of Excellence supporting their future collaboration, and to find means for informing the interested public about the research work, in order to facilitate the dissemination, transfer and exploitation of the results of GC II.

Although the GC projects have published widely, efforts are also undergoing in order to present the main results of the initiative to a broader audience. This document is a first step in that direction.

ANNEX — GLOBAL COMPUTING PROJECTS

Global Computing II	
Algorithmic Principles for Building Efficient Overlay Computers (AEOLUS)	AEOLUS

The recent explosive growth of the Internet gives rise to the possibility of a global computer of grand-scale consisting of Internet-connected computing entities (possibly mobile, with varying computational capabilities, connected among them with different communication media), globally available and able to provide to its users a rich menu of high-level integrated services that make use of its aggregated computational power, storage space, and information resources. Achieving this efficiently and transparently is a major challenge that can be overcome by introducing an intermediate layer, the overlay computer.

The goal of this project is to investigate the principles and develop the algorithmic methods for building such an overlay computer that enables this efficient and transparent access to the resources of an Internet-based global computer.

In particular, the main objectives of this project are:

- To identify and study the important fundamental problems and investigate the corresponding algorithmic principles related to overlay computers running on global computers.
- To identify the important functionalities such an overlay computer should provide as tools to the programmer, and to develop, rigorously analyze, and experimentally validate algorithmic methods that can make these functionalities efficient, scalable, fault-tolerant, and transparent to heterogeneity.
- To provide improved methods for communication and computing among wireless and possibly mobile nodes so that they can transparently become part of larger Internet-based overlay computers.
- To implement a set of functionalities, integrate them under a common software platform in order to provide the basic primitives of an overlay computer, as well as build sample services on this overlay computer, thus providing a proof-of-concept for our theoretical results.

Results

This project has not started yet.

AEOLUS Integrated Project ISTF FP6-15964-– total budget 6.85 M€
Coordinated by University of Patras (Greece)

http://www.

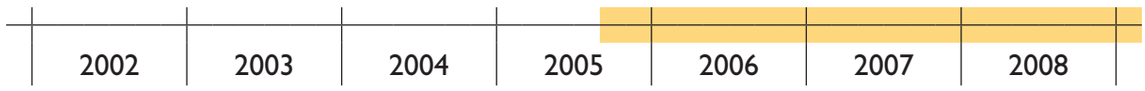

Global Computing

Architectures for Mobility (AGILE)

Architecture-based approaches have been promoted as a means of controlling the complexity of system construction and evolution, namely for providing systems with the agility required to operate in turbulent environments and adapt very quickly to changes in the enterprise world. Recent technological advances in communication and distribution have made mobility an additional factor of complexity, one for which current architectural concepts and techniques are not prepared for.

AGILE aimed to provide means for addressing this new level of complexity by developing an architectural approach in which mobility aspects are a key factor which can be modelled explicitly and mapped on the distribution and communication topology made available at physical levels. The whole approach is to be developed over a uniform mathematical framework based on graph-oriented techniques that will support sound methodological principles, formal analysis, and refinement.

Results

The project has pursued in a comprehensive manner the "architectural approach" to modelling, building and analysing mobile software systems. It has developed logical frameworks that clearly separate the three orthogonal issues of computation, coordination, and distribution/mobility from each other and that emphasize by this separation of concerns modularity in design and reasoning.

Key results include:

- primitives for explicitly addressing mobility within architectural models,
- algebraic models of the evolution processes that result from system reconfiguration caused by mobility of components,
- extensions to modelling languages like the UML that make the architectural primitives available to practitioners,
- analysis techniques for supporting compositional verification of properties,
- refinement techniques for relating logical modelling levels with the distribution and communication topology available at physical levels.

The project has produced 145 refereed publications.

AGILE Research Project IST-2001-32747 – total budget 1.86M€
Coordinated by Ludwig-Maximilians-Universitaet Muenchen (Germany)

http://www.pst.informatik.uni-muenchen.de/projekte/agile/

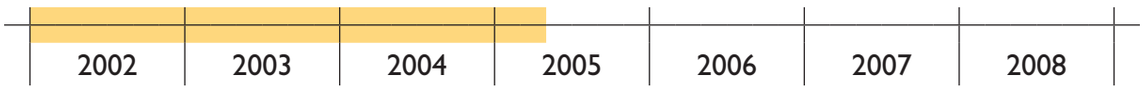

FET OPEN

Applied Semantics II

APPSEM II

The objective of this accompanying measure (Thematic Network) is to maintain and further develop an existing European network for research and technology transfer in the field of application-oriented semantics of programming languages. Programming languages are one of the essential ingredients for successful software development. This project is a follow-up of earlier EU-funded Working Groups; it funds joint workshops and mutual visits between its 20 sites; it does not fund research directly.

Theoretical results and expertise that are available in the consortium will be used to solve problems relating to programming with proofs, improving specification and verification, type systems, distributed computing, programming with continuous data, programming the web. At the same time and enabled by these practical problems the theoretical toolbox will be further developed, in particular the semantic understanding of sequentiality, distribution, object-orientation and of resource-bounded computation. This will happen through mutual research visits, workshops and one summer school.

Results

The project is not concluded yet. Insofar, it has been successful in organising a general workshop, including an Industrial panel session, and several thematic workshops (on Logic, Domains and Programming Languages, and on Normalisation by Evaluation). It also co-sponsored some relevant conferences and summer schools, and maintained a website which describes the research topics of APPSEM, provides access to all publications and can thus be used as a starting point for someone in search of information on programming language semantics.

This project also supported over 20 individual visits in 2004 resulting in 20 scientific publications (15 in refereed journals or competitive international conferences). In general, more than 200 scientific publications by were produced by the members relative to the APPSEM-II themes.

APPSEM-II Thematic Network IST-2001-38957 – total budget 0.4 M€
Coordinated by Ludwig-Maximilians-Universitaet Muenchen (Germany)

http://www.tcs.informatik.uni-muenchen.de/~mhofmann/appsem2/

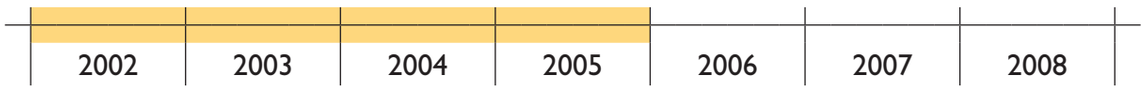

FET OPEN

Advanced Specialization And Analysis For Pervasive Systems (ASAP)

The overall aim of this project is to develop techniques which enable the development of sophisticated and reliable software systems that are easy to maintain and can be deployed on new generation, pervasive computing platforms.

Such platforms are becoming ever smaller and more powerful, and are embedded everywhere, even in living organisms. They can contain sophisticated models of our personal environment that help us to make everyday decisions; they have the power to do mathematical and logical reasoning in order to perform intelligent tasks. As a result, system development and validation techniques have to keep pace with the huge requirements for intelligent, user-oriented applications that must run on devices with reduced computing resources.

The development tools available for such pervasive computing platforms are lagging behind the state-of-the-art and there is a strong need by industry to optimise development cost, reliability, and time to market. Unfortunately, the tendency of current software engineering techniques is to produce more and more bloated, general-purpose systems, built from large libraries of pre-written modules. While this approach can increase productivity and reliability, it is wasteful of computing resources. Pervasive system developers are always forced to operate on the boundaries of available computing resources, even though processors, storage and communication networks are getting cheaper. One wants to pack as much functionality as possible into a given device. The main approaches adopted are either to produce by hand "stripped-down" versions of full systems in an ad-hoc manner, or to write applications from scratch taking limited resources into account. This project proposes a different approach: the automated generation and validation of specialized systems from general ones.

Results

This project has not yet finished. Some key achievements are:

- An off-line specialization system with an automatic binding time analysis.
- An integrated analysis and specialization system for CLP which incorporates predefined assertions and backwards analysis.
- A self-tuning specialization system, using a genetic algorithm to perform resource and architecture aware specialization
- A new program slicing tool, which can be used to reduce the code size of programs.
- A set of case studies selected from wearable computing projects, where CLP specialization and analysis tools have been applied.
- A new approach to resource-aware certification of mobile code named abstraction-carrying code.

ASAP Research Project IST-2001-38059 – total budget 1.44 M€
Coordinated by Universidad Politecnica De Madrid (Spain)

http://clip.dia.fi.upm.es/Projects/ASAP/

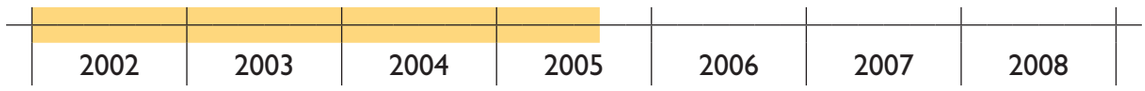

FET OPEN

Evaluation of the Catallaxy paradigm for decentralized operation of dynamic application networks (CATNETS)

Future Grid network technology will face the problem of the efficient provisioning of services to clients by a scalable and dynamic resource allocation (matching) mechanism. The objective of this project is to study free-market principles (the ***Catallaxy*** paradigm) as a means of guiding resource allocation in application layer networks, where client requests need to be matched with services and the latter with resources.

In particular ***CATNETS*** aims to determine the applicability of a decentralized economic self-organization mechanism for resource allocation in application layer networks (ALN), by
(1) producing a 'proof-of-concept' prototype in a real ALN, and
(2) by evaluating its performance against existing resource brokerage approaches in a simulated ALN.

The term ALN integrates different Internet overlay network approaches, like Grid and P2P systems. The allocation of resources in these networks (e.g. matching of demand and supply, deployment of service instances, and service discovery) can principally be conducted in a centralized (e.g. using resource brokers) or in a decentralized fashion (e.g. using self-organizing mechanisms). Centralized approaches reach their limits with increasing network size and growing numbers of elements; self-organizing approaches thus gain attention, e.g. in IBM's Autonomic Computing initiative. The ***CATNETS*** project investigates a 'free market' economic self-organization approach, the 'Catallaxy' by Friedrich A. von Hayek, as the basis for self-organizing resource allocation in ALNs. A preliminary evaluation of 'Catallactic' mechanisms in the FET assessment project CATNET (IST-2001-34030) by simulation has shown positive results, upon which the ***CATNETS*** project builds. The performance measurements of the 'proof-of-concept' implementation will be compared against the simulation results, with the goal of being able to make a substantiated statement on the applicability of economic self-organization as a major component of ALN networks. A positive evaluation of the Catallactic approach would have a high potential impact, with new possibilities for resource brokering in future ALN, and maybe for self-organization in computing in general.

Results

This project has only recently started.

CATNETS Research Project FP6-003769 – total budget 1.78 M€
Coordinated by Universitaet Bayreuth (Germany)

http://

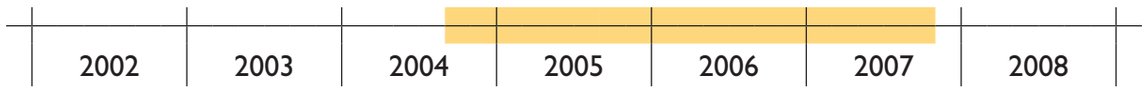

Global Computing

Critical Resource Sharing for Cooperation in Complex Systems (CRESCCO)

CRESCCO

The CRESCCO project aims at the integration in modern Information and Communication Technologies of:
i) heterogeneous communication infrastructures (optical, ATM networks);
ii) mobile users accessing these backbone networks and the Internet, and;
iii) dynamic selfish agent entities, results in highly dynamic, complex, global systems.

To design and implement high-speed, cost-effective, and reliable communication and computing solutions for such environments, the CRESCCO project investigated the bottleneck and critical issues involved at the fundamental algorithmic level. In particular, the focus was on the efficient management of scarce and critical resources such as frequency spectrum and energy in wireless networks, bandwidth in optical and ATM networks, as well as CPU time, space and communication time in dynamic environments of selfish agents. A multidisciplinary approach was pursued that combines and extends techniques from computer science, economics and operations research.

This project investigated at a foundational level the critical aspects and bottlenecks arising in highly dynamic and extremely complex global systems emerging from the integration of heterogeneous communication infrastructures, highly dynamic mobile user populations and dynamic selfish agent entities. The focus was on the efficient management of limited and scarce communication and computing resources to ensure co-operation and efficient access to advanced services in such environments by investigating the design and implementation of high-speed, cost-effective and reliable communication and computing solutions.

Results

New algorithms and techniques for bandwidth sharing in wireless networks, distributed computing with selfish agents, scheduling mechanisms, and algorithms for optical networking were developed.

The consequent application of economic game and auction theories to the question of resource sharing in distributed computing systems has led to fundamental insights into resource management with monetary incentives.

CRESCCO Research Project IST-2001-33135 – total budget 2.12 M€
Coordinated by University of Patras (Greece)

http://www.ceid.upatras.gr/faculty/kakl/crescco/

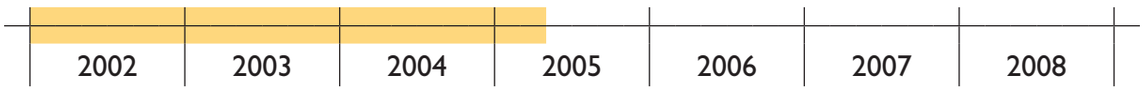

Global Computing

Dynamic Assembly, Reconfiguration and Type-checking (DART)

The project will develop formalisms for dynamic assembly, reconfiguration and type-checking of complex distributed software systems, such as telephone and banking systems, that should be kept running as they evolve through patches or upgrades, and should be able to adapt to changes in the environment. Such formalisms will advance the state of the art in modelling the "temporal" dimension of Global Computing (GC), where the ability to interleave meta-programming activities, like assembly and reconfiguration, with computational activities is a must. The development of these calculi will rely on decisive progress in three areas: calculi for dynamic assembly, calculi for object evolution and adaptation, flexible and compositional type systems.

The project objective is to advance the state of the art in modelling and programming software evolution while retaining safety. The main goal will be to carry out the foundational work, which will take the form of frameworks for dynamic assembly and reasoning about properties of different assembly strategies, calculi for object evolution and adaptation, type systems with properties that will make them particularly suitable for use in a dynamic context.

Results

The main contributions by the DART project are in the development of

- foundational calculi for dynamic assembly and reconfiguration, which are able to describe separate compilation, dynamic linking and loading
- foundational calculi supporting objects capable of changing their behaviour, as well as calculi that are environment adaptable
- type systems that support compositional analysis through the existence of principal typings. This includes the use of such type systems for separate compilation, incremental type inference and dynamic type-checking.

The consortium has produced software prototypes that demonstrate practical applications of the project results, especially with respect to "smart recompilation", which may be of interest to tool vendors. Also the innovative idea of "polymorphic byte code" may have long-term influence on the middleware industry.

The project has produced 10 publications in journals and 38 in conferences and workshops.

DART Research Project IST-2001-33477 – total budget 1.38 M€
Coordinated by Universitá degli Studi di Genova (Italy)

http://www.disi.unige.it/project/DART/

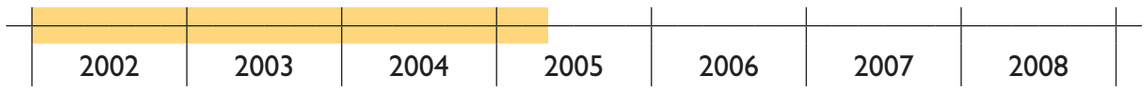

Global Computing

A Data-Centric Approach to Global Computing (DBGLOBE)

On one view, global computing can be seen as a database problem: how to design, build and analyse systems that manage large amount of data. However, the traditional database approach of storing data of interest in monolithic database management systems becomes obsolete in such environments. In current database research, data are relatively homogeneous, exhibit a small degree of distribution (just a few network sites) are passive in that they remain unchanged unless explicitly updated. All these assumptions do not hold in the global computing world. This creates the need for new theoretical foundations in all aspects of data management: modelling, storage and querying.

DBGlobe aims at revolutionizing the way we think of databases, as the mundane task of processing static data in monolithic database management systems. It broadens database management research focus to attack the issues of mobility, autonomy, incomplete information, scale, and adaptability that arise in dynamic environments.

Results

A key result was the development of a new powerful language for integrating documents and queries, called Active XML. The new language was cited as a major success by the project reviewers. There are links from the AXML developers to the W3C XQUERY group which might lead to a take-up of AXML features in the XQUERY standard.

DBGLOBE Research Project IST-2001-32645 – total budget 1.11 M€
Coordinated by University of Ioannina (Greece)

http://softsys.cs.uoi.gr/dbglobe/

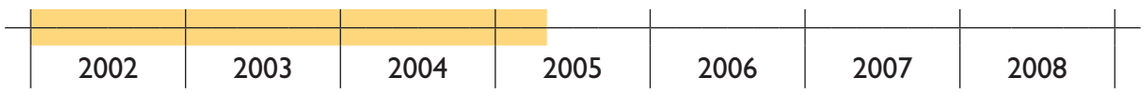

Global Computing

Design Environments for Global ApplicationS (DEGAS)

This project aims to combine structured (semi-formal) graphical methods for specification by picture and animation of global applications with formal methods for their analysis and verification. The project investigated to what extent UML is already suitable to model global applications and proposed extensions. Formal models of these applications based on the operational semantics of foundational process calculi for mobility were also proposed. Static and dynamic analysis concentrate on two key-features of global computing: performance prediction and security. The DEGAS project assessed the foundational studies in a prototypical proof of concept environment that hides to the user as much as possible of the formal treatment. The development was supported with case-studies on wireless telecommunication applications..

DEGAS addresses foundational aspects for the design of global applications by enhancing the state of the art in scientific as well as engineering principles. The main concerns are the specification in UML and qualitative and quantitative analysis of global applications.

Results

The project has a relatively narrow perspective in the sense that the main goal has been the verification of protocols using algebraic means. This is, however, a very important part of any software development leading to global applications. In that respect the scientific results of the consortium are extremely valuable. The protocol verification as proposed in this project permits to answer the questions of security and scalability that have been identified in the different GC calls as corner stones of global applications. The project has proposed a modular architecture based on extractors and reflectors that permit the combination of several formal techniques in a comfortable way. The tools have been successfully used to detect breaches in some classical communication protocols, and to model different case studies. The activity in proving correctness of protocols is receiving an increased interest by the scientific community of GC with techniques like theorem proving (project SOCS) or model checking.

Dr. Jane Hillston from University of Edinburgh received the Roger Needham Award for her research achievements, the last three years of which she has worked on the DEGAS project. In her award lecture and in other invited lectures she has presented results of DEGAS to the general public.

DEGAS Research Project IST-2001-32072 – total budget 2.28 M€
Coordinated by Universitá degli Studi di Trento (Italy)

http://www.omnys.it/degas/

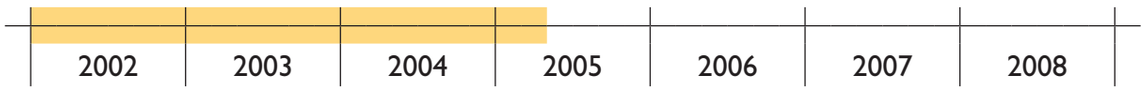

Global Computing

Foundational Aspects of Global Computing Systems (FLAGS)

FLAGS

New global computing and communication environments are emerging that integrate:

(a) autonomous, interacting, selfish entities,

(b) highly dynamic multi-agent environments and

(c) ad-hoc mobile networks.

For the efficient and robust implementation of global computing scenarios in such systems, we aim to provide a unifying foundational framework and a coherent set of design rules, focusing on the issues of:

(i) co-operation and antagonism of autonomous entities;

(ii) stability and fault-tolerance in multi-agent environments and

(iii) communication and motion in ad-hoc mobile networks.

The project aims to provide a unifying scientific framework and a coherent set of design rules for global systems resulting from the integration of autonomous interacting entities, dynamic multi-agent environments and ad-hoc mobile networks. The focus will be on the issues of co-operation and antagonism, stability and fault-tolerance, as well as communication and motion in such global systems.

Results

By applying economic game and auction theories to the question of resource sharing in distributed computing systems, the project has achieved an impressive set of important foundational results for a wide range of design aspects of GC systems.

The project has created a comprehensive set of design rules, proofs and applications for various aspects on the three architectural layers that have been defined.

FLAGS Research Project IST-2001-33116 – total budget 1.53 M€
Coordinated by Research Academic Computer Technology Institute (Greece)

http://em1server.cti.gr:8080/FLAGS/

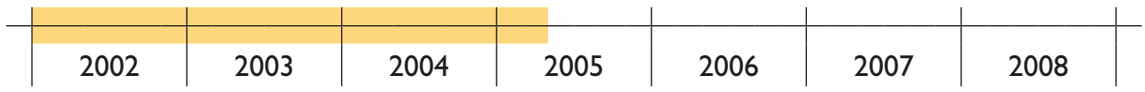

Global Computing

Working Group on Trust Management in Dynamic Open Systems (*ITRUST*)

Merely having a certified identity in a dynamic and open environment is not enough for informed decisions on access restrictions and controls, selection among potential candidates for interaction, and even less adequate for reasoning about the expected performance and dependability of entities for which no prior knowledge is available. The aim of the ***iTrust*** Working Group is to provide a forum for cross-disciplinary investigation of the application of trust as a means of establishing confidence in the global computing infrastructure, recognizing trust as a crucial enabler for meaningfull and mutually beneficial interactions. ***iTrust*** will strive to integrate perspectives of trust from the social sciences, philosophy and law with trust-related research in formal reasoning frameworks, security technologies, and distributed software systems. ***iTrust*** will seek to establish synergies with ongoing RTD activities.

ITRUST Research Project IST- 2001-34910 – total budget 0.4 M€
Coordinated by University Of Crete (Greece)

http://www.itrust.uoc.gr/2002

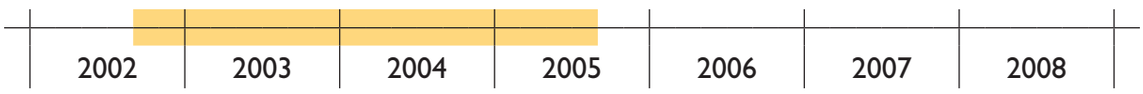

Global Computing

Mobile Calculi Based on Domains (MIKADO)

Current middleware and programming language technologies are inadequate to meet the challenges posed by a global computing environment. In particular, they tend to support only a limited range of interactions, have a limited view of components and objects, fail to properly and uniformly support properties such as mobility, predictability, security, fault-tolerance, and they are not amenable to rigorous investigation for verification, validation and test purposes. The Mikado project intends to overcome these limitations by defining and prototyping new formal models for both the specification and programming of highly distributed and mobile systems, and to develop specification and analysis techniques which can be used to build safer and trustworthy systems, to demonstrate their conformance to specifications and to analyse their behaviour.

The goal of the Mikado project is to construct a new formal programming model, based upon the notion of domain as a computing concept, which supports reliable, distributed, mobile computation, and provides the mathematical basis for a secure standard for distributed computing in open systems.

Results

The most significant contributions have been in the area of parametric models of migration and mobility, proof techniques for distributed models and programming languages which support such models. MIKADO has been at the heart of recent developments of bisimulation based proof techniques for mobile and distributed computing. Mobile distributed computing is inherently higher-order, and the project has created foundations for further research in this area.

The research of MIKADO has led to approximately 70 refereed journal and conference publications and 5 book chapters. The project has led to new collaborations also outside the project, for example on IMC.

MIKADO Research Project IST-2001-32222 – total budget 3.15 M€
Coordinated by Institut National de Recherche en Informatique et en Automatique (France)

http://mikado.di.fc.ul.pt/

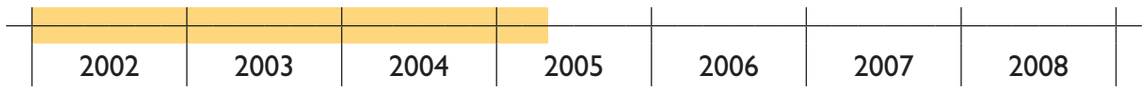

Global Computing II

Mobility, Ubiquity and Security (MOBIUS)

While global computers have the potential to profoundly enhance our daily life by providing a global and uniform provision to services, they will not realize their full potential, unless the necessary levels of trust and security can be guaranteed.

The objective of the MOBIUS project is to develop the technology for establishing trust and security for the next generation of global computers based on Java-enabled devices.

To realise its goal, MOBIUS shall pioneer a Proof Carrying Code (PCC) framework applicable to global computers, and allowing enforcement of functional properties as well as advanced security properties of software.

Results

This project has not started yet.

MOBIUS Integrated Project IST FP6 15905– total budget 8.11 M€
Coordinated by INRIA Sophia-Antipolis (France)

http://mobius.inria.fr/

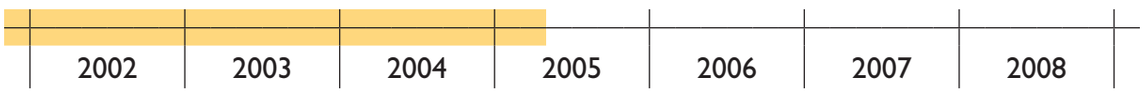

Global Computing

Mobile Resources Guarantees (MRG)

The use of mobile code in a global environment aggravates existing security problems and presents altogether new ones, one of which is the maintenance of bounds on quantitative resources. Without some technological foundations for providing such guarantees, global computing will be confined to applications where malfunction due to resource bound violation is accepted as normal and has little consequence. With more serious applications, resource awareness will be a crucial asset. This project aims at developing the infrastructure needed to endow mobile code with independently verifiable certificates describing resource behaviour. These certificates will be condensed and formalised mathematical proofs of a resource-related property, which are by their very nature self-evident and un-forgeable. Arbitrarily complex methods may be used to construct these certificates, but their verification will always be a simple computation.

The overall aim of MRG project is to develop proof-carrying code for resource-related properties where proofs are generated from typing derivations in a resource-aware type system.

Results

Successful completion of a PCC infrastructure including foundations (operational semantics, formalized soundness and relative completeness proofs), type systems, and web demonstration systems with downloadable software for the infrastructure components.

Although a rather small project in size, the project has made a relatively large contribution in the realm of PCC-based system analysis and verification. A theoretically well founded, comprehensive PCC infrastructure as well as an innovative framework for formal certificates of resource consumption were developed. The project has increased substantially the insights into the power and usefulness of this technology.

MRG Research Project IST-2001-33149 – total budget 1.25 M€
Coordinated by The University of Edinburgh (United Kingdom)

http://groups.inf.ed.ac.uk/mrg/

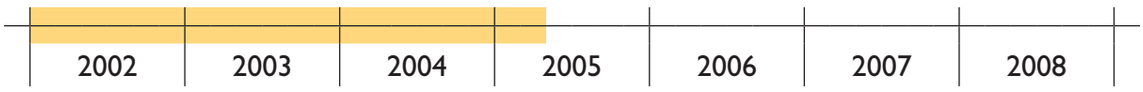

Global Computing

Models and Types for Security in Mobile Distributed Systems (MYTHS)

Peer-to-peer computing on the Internet, mobile code downloading, and e-Commerce are already ubiquitous aspects of our computing environments. Casting a view to the immediate future, we see a global computational infrastructure emerging that will rely on the sharing of an open-ended distributed network of mobile resources among mobile computing agents. This infrastructure can only be successful insomuch as it provides adequate security guarantees of secrecy, integrity, availability, and more. This project seeks to develop type-based foundational theories of security for mobile and distributed systems. By relying on strong typing as the basic principle, ***MyThS*** addresses the foundations of programming languages and paradigms that allow static detection of security violations, and aims at developing type theoretic methods and tools that enable formal analyses of security guarantees appropriate for systems and applications on the global computing platform.

MyThS aims to deliver an integrated set of type systems for innovative resource management and security policies, information flow analyses, and formal validation of distributed crypto-protocols, and establish rigorous design principles for secure and provably flawless systems.

Results

This project contributed to the development of type-based theories enforcing the security of GC systems and applications. In particular issues such as access control, information flow analysis, and verification of cryptographic protocols were addressed. These issues are of paramount importance in preventing attacks, in avoiding leaks and in guaranteeing (partial) correctness.

Key results of ***MyThS*** include prototypes which concern resource control in embedded systems, checks of incoming agents, and absences of information flow leaks that may be used by system designers for developing compliant software. In particular prototypes concerning information flow leaks of XML document transformations by middleware may be amenable to commercial exploitation.

MyThS Research Project IST-2001-32617 – total budget 0.99 M€
Coordinated by The University of Sussex (United Kingdom)

http://www.cogs.susx.ac.uk/projects/myths/

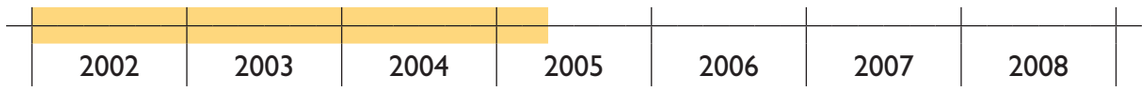

Global Computing

Peer-to-Peer: Implementation and Theory

PEPITO

Traditional centralised system architectures are ever more inadequate. A good understanding is lacking of future decentralised peer-to-peer (P2P) models for collaboration and computing, both of how to build them robustly and of what can be built. The ***PEPITO*** project will investigate completely decentralised models of P2P computing.

The project will:

- study the use-models of P2P systems, that is how they are perceived by users and what new applications are possible;
- develop the foundations of P2P computing, including formal foundations (calculi, proof techniques, security and resource models) and new distributed algorithms (for diffusing information and coping with multi-consistent views);
- provide a language-independent distribution subsystem tailored for P2P computing and;
- provide programming languages and platforms using this, showing that they are useful by implementing convincing demonstrator applications.

Peer-to-peer computing (P2P) is a paradigm in which applications are connected to a shared network as peers, that is with the same capabilities and responsibilities. Current P2P applications are limited to information exchange.

The objectives are to remove this limitation by:

- developing formal models to understand P2P computing;
- developing the distributed algorithms required for implementation;
- implementing a language-independent set of basic services;
- implementing languages and devise programming techniques and convincing demonstrator applications.

Results

Extension of the formalization of Internet protocols, which led to the discovery of several errors in popular implementations; completion of a generic P2P middleware, that was used – among other demo applications – to build a distributed file system.

The project provided a practical proof that server-less distributed systems are feasible and can provide very similar services as conventional, client/server systems. Some of the demo applications have raised interest for commercial take-up.

PEPITO Research Project IST-2001-33234 – total budget 2.78 M€
Coordinated by Swedish Institute of Computer Science (Sweden)

http://www.sics.se/pepito/

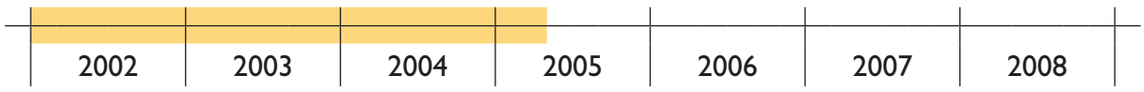

Global Computing

Performance Portability of OpenMP (POP)

In the last years, the industry is adopting OpenMP as a common shared-memory programming standard. Using this standard eases the task of implementing parallel programs for SMPs and makes parallelism more popular in industrial environments. On the other hand, complete OpenMP is currently only available for SMP machines, but no efficient implementation exists for other architectures. Our objective is to build an environment able to generate efficient OpenMP code for any architecture, which will avoid the necessity of using different programming models for different architectures. These architectures will range from shared-memory architectures and multi-threaded architectures to clusters of workstations/ SMPs with software distributed shared memory (SDSM).

POP Research Project IST-2001-33071 – total budget 2.17 M€
Coordinated by Universitat Politecnica de Catalunya

http://www.cepba.upc.edu/pop/2002

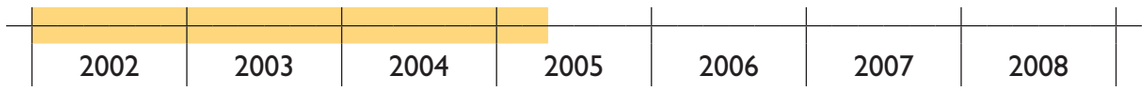

Global Computing

Proofs of Functionality for Mobile Distributed Systems (PROFUNDIS)

PROFUNDIS aims at developing methods to analyse the behaviour of distributed mobile systems, in order to ascertain that they function correctly. This involves modelling the systems in an abstract way and formulating rigorous correctness properties; it will be necessary to consider open and extensible systems with unknowable parts. For this purpose operational models (based on automata), algebras, logical languages, and associated type systems have been developed. Analysis was conducted through computer tools, both fully automatic and interactive. The novelty of the project lies in integrating several theoretical strands into one framework and one set of tools geared towards mobile distributed systems. Security properties and systems used in electronic commerce received specific attention.

The objective of ***PROFUNDIS*** is to advance the state of the art of formal modelling and verification techniques to the point where key issues in mobile distributed systems, such as security protocols, authentication, access rights and resource management can be treated rigorously and with considerable automatic support. A specific objective is to verify properties typical in so called open systems, where the behaviour of some parts (like intruders or adversaries) is unknowable, in extensible systems, where parts may be added or removed as the system executes, and in mobile systems where physical and logical connectivity between parts may change. This includes implementing automatic and partly automatic analysis methods for ascertaining correct behaviour of such systems.

Results

Important results of profundis have been in the area of models, such as the formulation of History-Dependent (HD) automata, toolkits, new logics such as spatial logics and type systems.

HD automata have been shown to provide an effective model suitable for finite state verification of finite control processes.

Specific toolkits were created provide for example minimisation of HD automata, checking of properties of cryptographic protocols, and implementation of symbolic execution of cryptographic protocols.

Type systems were proposed as solutions for specific problems in distributed mobile systems, such as interference, control of the migration of processes and of the access to local resources that a migrated process is granted, resource allocation, the specification of guarantees on the services that a process can offer message-deliverability, termination and access control.

The project has resulted in 82 articles in reviewed journals and conferences.

PROFUNDIS Research Project IST-2001-33100 – total budget 1.59 M€
Coordinated by Uppsala Universitet (Sweden)

http://www.it.uu.se/profundis/

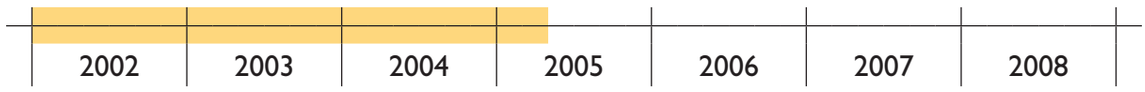

Global Computing

Secure Environments for Collaboration among Ubiquitous Roaming Entities (SECURE)

It is clear that new approaches to security are needed for the infrastructure envisaged by the global computing initiative, which is characterized by decentralised control. The SECURE project investigated the design of dynamic and self-configuring security mechanisms that are appropriate for global-scale and pervasive computing based on the human notion of trust. In particular, the project addressed how entities that encounter each other in unfamiliar computing environments, such as the Internet, can overcome initial suspicion to allow collaboration to take place.

The project's aim was to develop a model in which trust relationships are established from the record of interaction between entities, and a security mechanism expressed in terms of such trust. SECURE also investigated how to specify access control policy based on trust.

The objectives of SECURE are the definition of a computational trust model, a collaboration model, means to specify and to enforce security policies based on trust, means to evaluate security policies based on trust, algorithms for trust management and the validation of the approach in the context of the formal model.

Results

SECURE has achieved high visibility in the growing community of *trust* and *trust management* research, which clearly demonstrates that the project is working at the forefront of the state-of-the-art. It has also established a prominent presence in the more traditional security and distributed systems communities.

The project has shown how entities can build up and compute trust in each other, and can take risk factors into account before proceeding with an interaction. The project undertook the design of a security framework, which is based on a formal computational model of trust, and applied it to applications such as SPAM filtering and reputation management for participants in electronic auctions.

The project also formally defined a computational trust model and a collaboration model capturing the dynamic aspects of the trust model; means to specify and to enforce security policies based on trust; means to evaluate security policies and implementations based on trust; and algorithms for trust management.

The consortium is planning to publish its major results and contributions in the form of a book in the Springer-Verlag LNCS series.

SECURE Research Project IST-2001-32486 – total budget 1.73 M€
Coordinated by Trinity College Dublin (Ireland)

http://secure.dsg.cs.tcd.ie

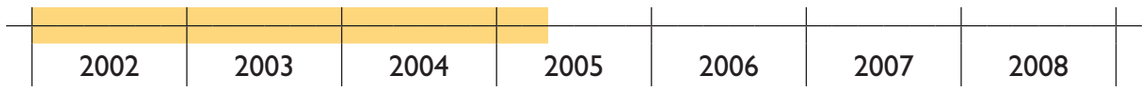

Global Computing

Secure and Safe Systems based on Static Analysis (*SECSAFE*)

Secsafe
Secure and Safe Systems based on Program Analysis

Static analysis of programs is a proven technology in the implementation of compilers and interpreters. Recent years have begun to see application of static analysis techniques in novel areas such as software validation (for example Ariane V) and software re-engineering (for example the Y2K problem). This project will demonstrate that static analysis technology facilitates the validation of systems based on the Internet and on smart cards.

The objective of the project is to assess the scalability of static analysis technology to the validation of security and safety aspects of realistic languages and applications. Two domains where security is all-important have been identified: smart cards and Internet programming. The project intends to develop methods that apply to both domains by focussing a substantial part of our efforts on the Java programming language and its dialect Java Card, treating source-level as well as bytecode-level applications.

SECSAFE Research Project IST-1999-29075 – total budget 1.64 M€
Coordinated By Imperial College of Science, Technology and Medicine (United Kingdom)

http://www.doc.ic.ac.uk/~siveroni/secsafe/2002

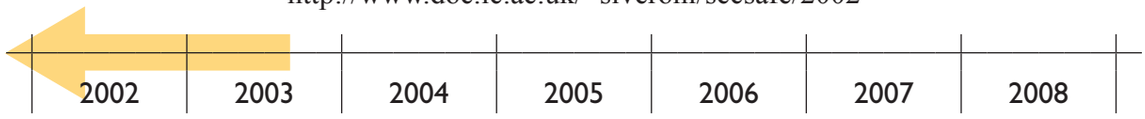

Global Computing II

Software Engineering for Service-Oriented Overlay Computers (SENSORIA)

Service-oriented computing is an emerging paradigm where services are understood as autonomous, platform-independent computational entities that can be described, discovered, and dynamically assembled for developing distributed, interoperable systems and applications.

The aim of SENSORIA is to develop a novel comprehensive approach to the engineering of service-oriented software systems where foundational theories and techniques are fully integrated in a pragmatic software engineering approach.

The results of SENSORIA will include new modelling and programming primitives for services, new mathematical analysis and verification techniques and tools for system behaviour and quality of service properties, and novel model-based transformation techniques.

Results

This project has not started yet.

SENSORIA Integrated Project IST-FP6 16004 – total budget 10.37 M€
Coordinated by Ludwig-Maximilians-Universitaet (Germany)

http://www.pst.ifi.lmu.de/projekte/sensoria/

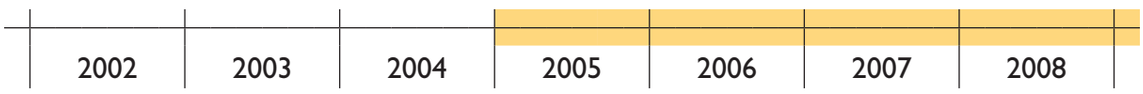

Global Computing

A computational logic model for the description, analysis and verification of global and open societies of heterogeneous computees (SOCS)

The aim of this project is to investigate computational and logical models for describing, analysing, and verifying individual and aggregates of computational entities – referred to as computees - interacting in the context of global, open and dynamic environments. Computees can form complex organisations, which, for the purposes of this project, we call societies of computees. We propose computational logic-based techniques for specifying and verifying formal properties of computees and their societies. We further validate the framework by a series of grounded controlled experiments using a prototype demonstrator embodying the formal model. The results of the project will also provide a practical basis for the design of classes of systems and applications which require aggregate behaviour of computational entities.

The projects aims to:

(1) deliver novel computational logic-based descriptions of computational entities - referred to as computees - with heterogeneous knowledge, objectives, roles and patterns of behaviour and interaction;

(2) describe societies of computees that are capable of interacting in a global, open, and dynamically changing environment;

(3) provide tools for the specification, analysis and verification of properties of computees and societies of computees emerging from their interactions.

Results

SOCS developed a prototype platform supporting the implementation and deployment of computees and societies of computees. In this platform, computees may exhibit heterogeneous profiles of behaviour, may interact with one another by communication and by observation, and are highly responsive to the dynamic changes in their environment. The prototype allows for expectations to be generated about the ideal social behaviour of the computees that are part of it.

SOCS Research Project IST-2001-32530 – total budget 1.91 M€
Coordinated by Imperial College of Science, Technology and Medicine (United Kingdom)

http://www-lia.deis.unibo.it/Research/Projects/SOCS/

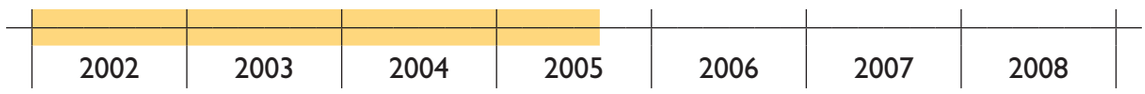

Global Computing	
Types for Proofs and Programs (TYPES)	TYPES

The project is a coordination action and the aim is to communicate and disseminate developments in the technology of formal reasoning and computer programming based on Type Theory. These developments are mainly in improving the languages and computerised tools for reasoning, and applying the technology in several domains such as analysis of programming languages, certified software, formalisation of mathematics and mathematics education.

In an earlier phase, TYPES has led to logics based on various type theories, proof tools based on these logics, and mathematics libraries which are essential for practical use of the proof tools. The aim was to develop systems that will be useful to mathematicians and other scientists in their proof-oriented work, much as computer algebra systems have been useful for symbolic computations.

The project plans to focus on the following four key araes:

- *Correctness of Computer Systems*: tools and techniques aimed specifically at application of formal methods to system correctness, e.g. programming language specific tools and problem-specific automation of proof search.
- *Formal Mathematics and Mathematics Education*: this is the prototype example for proof in the large, including very high level mathematical vernacular languages, the construction and use of necessarily large libraries of previous work, and distributed work on long-term projects.
- *Proof Technology:* the details of proof, including unification, resolution, rewriting, general proof search, tactic languages and declarative proof languages.
- *Foundational Research:* underlying the previous three areas must be research on the expressiveness and relative correctness of the foundational logics, including syntax, semantics, definitional mechanisms, allowed computation and subtyping.

TYPES will organise three annual meetings to communicate recent work throughout the consortium, six smaller thematic workshops on designated research themes, one summer school, short courses and short visits between sites.

TYPES Research Project FP6-510996 – total budget 0.6 M€
Coordinated by Chalmers Tekniska Hoegskola Aktiebolag (Sweden)

http://www.cs.chalmers.se/Cs/Research/Logic/Types/