# The FET Global Computing Initiative

**Carles Sierra**
**30 November 2006**

# 1 Introduction

The Global Computing (GC) initiative was launched in 2001 as part of the IST FET Programme. The vision of the call, also contained in the current Global Computing II initiative, was to focus research on large-scale open distributed systems. This was a timely vision given the exponential growth of Internet and the turmoil generated in the media and in scientific fora of some international initiatives such as the Semantic Web, or the IBM autonomic computing concepts, and the peak of Napster usage in 2001 with more than 25 million users. The first GC call (CGI) received 18.2 M€ for projects that had a life span of more than three years on average. Most projects had a high interdisciplinary nature, and a large number of groups from theoretical computer science, agents, networks and databases worked well together.

Global computing refers to computation over "global computers", i.e., computational infrastructures available globally and able to provide uniform services with variable guarantees for communication, co-operation and mobility, resource usage, security policies and mechanisms, etc., with particular regard to exploiting their universal scale and the programmability of their services. As the scope and computational power of global infrastructures continues to grow, in order to harvest their potential benefits, and ultimately improve our quality of life, a vision needs to be realised which goes well beyond incremental and disconnected improvements of diverse (and often incompatible) implementations.

The key aim of this initiative is to define innovative theories, computational paradigms, linguistic mechanisms and implementation techniques for the design, realisation and deployment of global computational environments and their application and management. The expected result in the long term is to achieve real, integrated global computing in a wide range of application scenarios by providing foundational advances on suitably large classes of global computers, together with the integration of methods and concepts necessary to advance global computing as a whole.

Through the abstractions of common characteristics representing global computers, referred to as "overlay computers", Global Computing tackles the four issues of security, resource usage and management, scalability and distribution transparency. There is growing overlap with the area of Grids, especially as the area of Grids has started to consider issues such as mobility, while researchers in the area of Global Computing have increasingly considered issues such as dependability and monitoring.

The focus of GCI was on three main topics: analysis of systems and security, languages and programming environments, and foundations of networks and large distributed systems. Along these lines, GCI projects dealt with formal techniques, mobility, distribution, security, trust, algorithms, and dynamics. The focus was ambitious and foundational, with an abstract view of computation at global level having as particular examples the GRID or the telephone network. Both functional and non-functional (e.g. Quality of Service) properties had to be studied.

The analysis done during the reviews of GCI projects revealed a stupendous production of scientific results but rather few immediate applications. The ‘Napster’-like big systems were not being generated as a direct consequence of the research. During 2005 a second call was launched, Global Computing II (GCII), with a shift in focus towards issues that would help in the actual deployment of such big applications, namely, security, resource management, scalability, and distribution transparency. Also, there was need for a more even balance between theory and building experimental systems in GCII.

Both calls have contributed to the creation of a European community on global computing that has organised several meetings where the main research ideas have been discussed. The first meetings within GC I took place in Rovereto in 2003 and 2004 and as specialised workshops on foundation of

Global Computing collocated with the conferences ICALP and Concur. Within the context of GC II two workshops with the appropriate name of Trustworthy Global Computing have been organised: TGC05 (April 7-9, 2005, Edinburgh) and TGC06 (November 7-9, 2006, Lucca). Formal proceedings as special issues of LNCS and ENTCS have been produced. The community aims at continuing the meetings in the future as a stable event in the scientific calendar. Contacts with the Grid community have been established and a large meeting took place on 28-29 November 2006 in Sophia-Antipolis to discuss the current research roadmap on the area of open and distributed systems (http://www.coregrid.net/mambo/content/view/270/288/). There are obvious research questions in common (e.g. security) that permit sharing ideas and solutions. Overall, Global Computing is becoming mainstream in computer science, and companies think most solutions on a distributed way with openness in mind. The number of applications using ad-hoc sensor networks, or adopting P2P solutions, is increasing and therefore the self-sustainability of the community is closer.

# 2 Overview of Global Computing II Projects

The projects of GCII are built on top of the results of GCI projects as mentioned in the introduction. Many techniques are followed up, and the consortia are to a large extent similar to those in GCI. The MOBIUS project shows a high potential for having an impact on practical systems. SENSORIA and AEOLUS are creating foundations that will become a powerful basis as the integration with applied parts of case studies takes off during the next phase of the projects.

## *AEOLUS*

*http://aeolus.ceid.upatras.gr/*

The focus of AEOLUS is on the algorithmic methods for global computing. Several issues require new algorithmic solutions relating to the topological properties of the networks of computing nodes, the selfish nature of the computing entities, the stability of the solutions given the dynamic nature of global computers, and the partial and uncertain knowledge that nodes have. A distinctive feature of AEOLUS is that it aims at improving the communication methods among wireless and mobile nodes in such networks. Testing the results on actual systems and networks is also a goal of the project. The research direction is centred on efficient and scalable resource usage (e.g., load and data management), although other aspects such as security are also considered.

## *MOBIUS*

*http://mobius.inria.fr/*

MOBIUS focuses on security for global computing, setting out a vision of program components transmitted between nodes on a distributed network. When using such a network one of the nodes might dynamically download and compose components to build a complex application with sophisticated functionality. A central concern in this vision is how to verify properties of the components so as to guarantee the behaviour of the composed system, with the flexibility to handle both qualitative and quantitative properties. The main outcome of the project will be a security framework that can verify and enforce such requirements on individual mobile components. To enable this vision, MOBIUS proposes a radical approach based on proof-carrying code, a novel technology that takes advantage of foundational work on logic and type systems being developed within the project. The key breakthrough in proof-carrying code (PCC) is that programs carry mathematical proofs of their own safety which a recipient can independently verify before execution. Unlike conventional cryptographic signatures, these proofs provide direct evidence about the program itself, rather than simply some claim about its author, and do not depend on any centralised trust infrastructure.

## *SENSORIA*

*http://www.sensoria-ist.eu/*

The SENSORIA project aims at helping engineers of GC applications along the whole design process. It builds upon existing standards, such as Web Services Description Language (WSDL) and the Business Process Execution Language (BPEL), to define a specification language for services. The verification of properties of the services expressed in that language is specially challenging. To help on this process the project proposes a (semi-) automated (via graph transformations) generation of specifications in different formalisms adapted to different types of property verification. For that purpose a rich variety of ideas has been explored in the definition of several process calculi into which the automatic transformations are made. Concepts like session, conversation, composition, events, synchronous, and asynchronous communication are going to be studied. SENSORIA is intensively studying graph transformation techniques that preserve the semantics of models written in different formalisms. Interestingly, they can be applied to behavioural models in order to prove properties over them.

## *CATNETS*

*http://www.iw.uni-karlsruhe.de/catnets/*

CATNETS is a STREP project funded in the FET Open Programme. Because of its close similarity to the GC Proactive Initiative projects, it has been included in this report.

The CATNETS project aims at finding scalable solutions for the matching of needs and offers in application networks, like the Grid or P2P systems. The solution explored in this project is the use of economic mechanisms. Through the use of catallactics, an economic theory that analyses all actions from a perspective based on individual utility functions, the market prices are pursued as well as the individual choices that give rise to these prices. This paradigm is going to be used to model the self-organisation of highly dynamic networks of nodes providing services to each other. A service cost is assumed to be included in WSDL documents. The project will develop a simulator to try and obtain scalable solutions and to obtain good resource allocation mechanisms. Basic negotiation mechanisms are explored, e.g. English auctions.

## *Complementarity of the projects and coverage*

All projects cover to different extent the four goals of the GCII call. However, the emphasis is different depending on the project. AEOLUS concentrates on resource management, MOBIUS on security, SENSORIA on design and engineering, and CATNETS on scalability. AEOLUS and SENSORIA aim at building platforms, MOBIUS aims at developing a framework for security that could be used in all sorts of global computers, and CATNETS aims at building a simulator that could be used to study scalability properties of highly dynamic networks. Interoperability is a concern of all projects, as it couldn't be otherwise, and different perspectives are taken on this: AEOLUS aims at mixing different devices, mobile, wireless, and fixed; MOBIUS at mixing different components developed by different people; and SENSORIA at permitting the specification of components using different languages and paradigms. These relations are reflected in Figure 1.

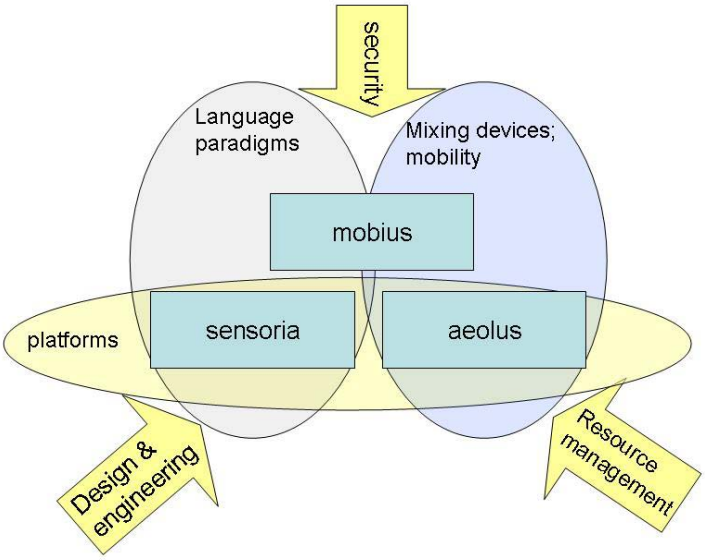


FIG 1. Schema showing the complementary nature of the IP projects in GC II.

### *Links between GCI and GCII*

In the GCI initiative 13 projects were funded, all of which have now been completed. A description of these projects and links to their web pages can be found at http://cordis.europa.eu/ist/fet/gc-sy.htm.

Many techniques developed and used within the context of GCI projects have been pursued further. Examples are proof-carrying code (MRG), game theory (FLAGS) and type-theory (MYTHS, TYPES, DART). Also, state-of-the-art surveys on a large number of techniques which are foundational for global computers have been produced within the project AEOLUS (e.g. algorithms, resource management, mechanism design, load management, scheduling, reputation). Mobility, which was a subject of study in the GCI project AGILE, is continued in the project AEOLUS. The scalability of algorithms, which is the main concern of AEOLUS, was a topic in GCI project CRESCCO. Scalability was also the concern of the assessment project CATNET on top of which the CATNETS STREP project builds up.

Static verification has continued as a driving force, especially in MOBIUS, as was the case in the GCI projects SECSAFE and PROFUNDIS  The development of platforms to deploy global computers, started in the GCI project SOCS, is an inspiration of all three projects currently funded under GCII.

## 3 Progress and main achievements

During the first year of development of the GCII projects, a strong progress has been made and many initial results have been published.

AEOLUS has delivered original contributions in the study of structural properties of the topology of computing nodes, the study of the trade-offs between fairness and efficiency and on secure distributed computation. Its subproject on wireless networks gives the project a unique dimension within the GCII initiative.

"Compositionality" as a basic architectural concept is being analysed in SENSORIA from several perspectives. It studies orchestration primitives for the composition of complex applications. And also proposes a new language, JAVA/A, to support architectural programming by permitting reasoning about software components. AEOLUS also has a compositional view on the combination of components to build complex applications. SENSORIA has proposed service-oriented extensions of UML and their corresponding transformations into different process calculi. Also, spatio-temporal logics are used to reason about complex choreographies: spatial and epistemic operators configure a new family of logics that permit model checking of distributed systems.

Mobius has already produced a comprehensive set of security requirements and proof carrying code scenarios for global computers. The project has also obtained significant results on type systems for confidentiality and resource control, and on program verification techniques that support practical reasoning about bytecode programs. The consortium is now focusing on the further development of these enabling technologies, in order to warrant the successful deployment of the security infrastructure in later stages of the project.

## 3.1 Security

AELOUS has developed an approach for coping with nodes in a global computer showing malicious and disruptive behaviour. This complements the classical approach based on cryptography, as used in SENSORIA, with an approach based on Economics. In particular, AELOUS has developed a new mechanism that permits anonymity in the interaction, which has great potential in areas like automated negotiation. Bounds on the number of malicious nodes in a network have been established for which a computing model proposed in AEOLUS can still guarantee the reliability of the computation. Also, a purely distributed protocol that guarantees security without the need of a third party has been developed in AEOLUS. MOBIUS has the most radical departure in terms of security as it aims at having components for which the behaviour can be certified. The foundational research on the supporting technologies for PCC is underway (e.g. its logic basis), and initial work has been carried out on performing the enforcement of the information flow between components with type systems.

The accepted importance of the Web-Services Business Process Execution Language (WS-BPEL) has inspired some developments within SENSORIA, for instance a process calculus that permits to model web services and orchestration languages. Services have also inspired the definition of a process calculus that models the definition and invocation of services as well as the crucial concept of session handling. The concept of session has also inspired the definition of type systems that specify interaction protocols as types. Thus, security can be assessed by simple static type checking. In fact, static security mechanisms are those privileged in SENSORIA (as in MOBIUS) due to their simplicity. An interesting approach explored in SENSORIA is a call-by-property mechanism where client requests impose constraints on the service behaviours. In SENSORIA security concerns have materialised in orthogonal extensions of process calculi that support the exchange of data. The extensions are currently for symmetric cryptography but easily extendable to other mechanisms. Finally, confidentiality on Publish/Subscribe systems can be guaranteed by some provably secure protocols.

## 3.2 Resource management

Resource management has been a central concern of AEOLUS, for instance in managing communication bandwidth of wireless networks and in managing communication, and through responding to the new challenges that mobility creates in communication because of the dynamicity of the network topology. The project has also proposed new methods for data propagation that minimise energy consumption. It builds up on the success of CRESCCO where similar resource management issues were investigated. As scheduling of scarce resources to guarantee Quality of Service (QoS) is a

major concern in overlay computers, several (centralised) models have been analysed in AEOLUS. A promising distributed method, based on an economical model, is currently being studied.

MOBIUS also focuses on resources, in particular for small portable devices like mobile phones where computing power is tightly constrained. This has a security impact for global computing: both because some resources are intrinsically valuable, such as phone calls or text messages, and because overuse of limited resources on a device can compromise availability, leading to denial of service. To address these MOBIUS proposes techniques for checking in advance the potential resource usage of components, so that security policies can be enforced before code is even executed.

The CATNETS project explicitly researches similarly into a free market (Catallactic) resource management model.

## 3.3 Scalability

The analysis of scalability requires tools that permit the analysis of the structure of the network supporting an overlay computer. AEOLUS has produced methods that permit to analyse the efficiency of a network given its structural properties, for instance for load balancing analysis. Game theoretical approaches in AEOLUS permit the scalability of ad-hoc and sensor networks by minimizing power consumption and distributing the communication over the structure of the network. Also, an algorithm that partitions a network into a number of clusters is proposed that permit efficient broadcastings and limited energy consumption; again, a useful proposal for wireless networks. CATNETS is using catallactic methods to analyse the resource allocation in very large and highly dynamic networks. A simulator has been programmed that implements several economic models to be used in the negotiations between service providers and service consumers.

Non-functional aspects are essential for global computing and SENSORIA approaches the issues of performance and dependability by means of stochastic extensions of process calculi. Very promising results have been demonstrated in SENSORIA that permit the direct analysis of performance from specifications written in process calculus, enabled by a transformation into differential equations. The scalability and performance of the system can then be analyzed by direct simulation on the differential equations without the need to actually build the system.

Achieving scalability of the PCC security framework requires substantial advances in the enabling technologies of type systems and program logics on which this framework relies. To ensure scalability, the second and third research and technological objectives of MOBIUS are the development of type systems and program logics

## 3.4 Distribution transparency

Some techniques that permit a transparent load balancing that maximises the efficiency of the overall network of computing nodes have been developed within AEOLUS. This project has also developed new techniques of replication and cache that are based on the topology of the network and that are the brick-and-mortar for distributed transparency. Distribution transparency is moreover accomplished in AEOLUS for heterogeneous components by a seamless integration of wireless networks with fixed networks. To cope with distribution transparency SENSORIA has performed interesting developments using the architectural shape concept that applies local modifications to automatically adapt distributed programs.

## 3.5 Service Oriented architecture

The project SENSORIA has a clear view on what a Service Oriented Architecture (SOA) has to be. SENSORIA is proposing a language for service modelling and has also proposed a method to integrate

legacy systems into a SOA. Its formal methods and model-based transformations are assumed to become extensions of UML so that industry can quickly take on the development of services. A Computer Aided Software Engineering (CASE) tool for the definition of services is being developed. The graphical language contains interesting proposals for service composition and expands on current standards like BPEL. The industrial interest in SOA is an indication of the expectations that Global Computing is creating.

# 4 Potential and impact

The main impact of this research is in the area of geographically distributed, open applications that require a collaborative infrastructure, development techniques and tools for supporting distributed computing and integration of distributed information. The potential for GCII in realising this impact is great. Some of the best groups in theoretical computer science are present in the consortia, good results are being obtained, and a large set of quality publications is being generated. The results of GCII will certainly influence the development of the distributed computing field.

One route to achieving significant real-world impact is through addressing challenges such as scale and heterogeneity, through, for example, efficient, composable and reusable service components. It may be possible, through fundamentally different architectures, to address the challenges posed by environments with millions or billions of nodes. The SENSORIA project, through services, and AEOLUS from the networking side, have the most direct potential to have an impact in this area.
In such environments a substantial impact would be achieved through management and coordination of resources, allowing for the real leverage of the advantage of open infrastructures. This can only be achieved through information, monitoring and exchange frameworks beyond what is available today, which would allow for management and co-ordination of resources. In terms of platforms that will be available at the end of GCII, AEOLUS aims at building a platform for all the algorithms for resource management and security, enabling the building of scalable and transparent applications. This will be constructed on top of the open source peer-to-peer platform JXTA. The functionalities are defined and the implementation is underway.

Security is a key factor in collaborative environments, and MOBIUS proposes a radical new security infrastructure based around proof-carrying code, where distributed devices can reliably check safety properties of mobile programs. This complements existing cryptographic mechanisms, to support truly open global computing.

The projects complement one another: AEOLUS is tackling the novel algorithms required to drive the global computing network, SENSORIA takes a foundational view of software engineering for global computers, and MOBIUS sets out a radical approach to guaranteeing security. Together, they promise to deliver a comprehensive framework for the future of global computing.

There has been a substantial scientific impact. In terms of dissemination the projects have developed a very intense academic activity with more than 130 papers published in journals and conferences during their first year. There are plans to present Global Computing projects to the industry at large (e.g. at SYSTEMS 2007 in Munich) and also some dissemination in the form of schools has been made in SENSORIA and AEOLUS. The projects are well positioned to propose a Masters-level curriculum for the study of service-oriented software engineering.

# 5 Future perspectives

In general, the projects have explored mostly ‘syntactic’ approaches to service oriented modelling. They have an opportunity to explore semantic approaches by interacting with research projects on the area of Semantic Web. The problems of semantics (touched upon in the area of ‘Resource Discovery’ of AEOLUS, for instance) are one of their basic concerns, and tools and methods are being developed (e.g. ontology alignment, semantic routing) that could be reused in GCII projects. Examples of such projects are DIP (http://dip.semanticweb.org/) and SEKT (http://www.sekt-project.com/).

There are opportunities for synergies among the projects. For instance, the advances of MOBIUS in security could make AEOLUS and SENSORIA consider the integration of the PCC infrastructure into their tools. Also, the use of logic-based techniques for verification in MOBIUS and SENSORIA calls for an increased interaction between the projects. Economic modelling is a concern in AEOLUS and CATNETS, and the sharing of results could be useful in the generation of scalable solutions for global computers. Finally, there is a clear opportunity for the projects to collectively determine the contents of a master course on Service Oriented Software Engineering.
As a start, SENSORIA and MOBIUS will soon organize a joint summer school. Other initiatives could follow.

The work on service-oriented modelling could explore interaction with research projects about the Semantic Web. For example, the issue of resource discovery in AEOLUS is a relevant concern, and tools and methods are being developed (ontology alignment, semantic routing) that could be reused in GCII projects.

The projects’ excellent results on security issues in the form of guaranteeing behaviour or getting bounds for non-functional properties opens the possibility for interchange of ideas with the agents community that are currently exploring trust and reputation methods inspired by sociological studies.

The four projects have an opportunity to deepen into the study of the dynamic aspects of Global Computers. As the node topology is dynamic, the algorithms and methods can be transformed and made incremental and adaptive, considering that static solutions might not work or scale. Some of the challenging questions for the future are how algorithms compute new solutions when minimal changes in the topology are introduced, and how long it takes to compute a new solution when the topology has changed locally