Global Computing II

Terms of Reference

---

1. FOREWORD

Global information and communication infrastructures support an ever-growing share of our economic and social activities. The extent of our dependence on such infrastructures creates an urgent, overwhelming demand for safety, soundness, robustness, effectiveness, and efficiency. 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.

2. VISION AND AIMS OF THE CALL

A global computer is a programmable computational infrastructure distributed at worldwide scale and available globally. It provides uniform services with variable guarantees for communication, cooperation and mobility, modalities and disciplines for resource usage, security policies and mechanisms, and more. Global computing refers to computation over such infrastructures, with particular regard to exploiting their universal scale and the programmability of their services.

Obvious (non exhaustive) examples of global computers are the Internet, the (world-wide) virtual private networks, the Web, the telephone network, the GRID, the infrastructures envisioned in the Ubiquitous and in the Disappearing Computer frameworks. Their respective characteristic services are: reliable stream transport (Internet); privacy and confidentiality (VPN); client/server extended handshake interaction (Web); guaranteed quality of service (telephone); sharing of computing power (GRID); seamless ubiquitous mobility (UC); minutely diffused computational capacity (DC). These are already (or will become) critical infrastructures for the economic and social activities of our society.

It is highly unlikely that further evolutions of current global computers will be replaced by convergence to a single infrastructure. Instead, a future of multiple, diverse infrastructures each realising a different global computer is envisaged. Such multiplicity has a utopic aspect, namely a global computer optimised for each and every need, and a dystopic aspect, namely a disarray of incompatible infrastructures.

The key aim of the Global Computing 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.

We aim at realising the "utopia" by calling for research to advance and surpass the current global computers. At the same time, we aim at avoiding the "dystopia" by calling for projects which focus on multiple (though not necessarily all) specific global computers.

Specific challenges for the programme are illustrated later in this document. The next section discusses in more detail the overall aims introduced above, and proposes a conceptual framework which may be used to focus on and address the double challenge they present (i.e., realising the utopia and reducing the dystopia).

3. THE CONCEPTUAL FRAMEWORK (Global and Overlay Computers)

To work effectively on both the utopic and the dystopic aspects in a cohesive manner, conceptual abstractions are required. One suitable instrument is provided by the notion of overlay computer. An overlay computer is an abstraction that can be implemented on top of a global computer, e.g. the Internet, to yield another (global) computer. It is overlaid because it lays upon existing global infrastructures, is implementable, programmable and, normally, computationally complete in its application domain.

By abstracting over common characteristics, an overlay computer may represent a family of potential or actual global computers, or possibly, in the limit case, a single one. The nature of the abstraction is not prescribed by this call, and it can range from a calculus to a set of computational primitives, from a set of characteristic axioms to a (layered) architecture, and more.

As the fundamental aspect of global computers is the actual infrastructure they provide to programmers, the focus in overlay computers is on the abstraction they build over the underlying global computer(s). Examples of overlay computers include: resource discovery services, which abstract a notion of resource sharing on top of global, distributed networks; search engines, which provide the abstraction of information repository over the Internet; systems of trusted mobile agents, which endow processes with a notion of autonomic, exploratory behaviour; IPSEC, which provides the abstraction of secrecy on top of IP; and several global infrastructures currently deployed over the Internet, such as the Web or the peer-to-peer file sharing protocols. The fundamental hypotheses underlying the work in the "Global Computing I" initiative (2001) give rise to an overlay computer too, characterised by the complete absence of centralisation and trustworthy entities, and by extreme structural variability. The GRID itself, seen as a layered architecture overlaid on the Internet and providing, e.g., resource abstraction, is an overlay computer. The provision of secure computing infrastructures, of business processes languages, and P2P mechanisms of various sorts, can all be phrased in terms of overlay computers. And since little abstraction is abstraction too, global computers can also be thought as (particularly concrete) overlay computers (as it was occasionally done above).

Less obvious, but equally valid notions of overlay computers are, e.g., programming languages for mobile agents bundled together with mechanisms for code safety certification; and languages for entity certification, when coupled with cryptographic mechanisms for authentication and certificate chain discovery.

In contrast, abstractions relying on synchrony, low latency, trustworthiness, limitless bandwidth are not overlay computers, as they cannot be realistically implemented in a global scenario. For the same reasons, **approaches neglecting scalability and security are not suitable for this call.** A specific non-example is the Web without server-side programming, as it likely fails to be computationally complete.

The dualism of global and overlay computer reflects and reveals a bi-dimensional structure of this pro-active initiative's workprogramme. According to it, global computers can be seen as vertical entities, self-contained computational models, that can be improved or newly invented (the utopic aspect). Dually, overlay computers are horizontal abstractions, embracing several notions at once, so that problems can be solved in generality, i.e., with a degree of independence from the actual underlying global computer (targetting the dystopic aspect). The purpose of overlay computers is therefore twofold. Firstly, they allow to experiment as openly as possible with infrastructures of potentially high impact before actually having to implement them. Secondly, they allow to address issues common to a wide range of specific global computers, which can be represented as overlay computers. This paves the way for general, unifying approaches which, therefore, decrease the dystopia and provide the depth of understanding necessary to actually broaden our future horizons.

4. THE CHALLENGES OF THE PROGRAMME

Project proposals are invited which focus on novel or existing real global computers or envisioned overlay computers, implemented or formalised. Projects should seek to advance beyond the boundaries of what currently exists, in order to realise the vision of this call, and aim at addressing problems at as general a level as possible. Projects may elect to confine their efforts to specific global computers (for instance as particular instances of overlay computers), if that is fully justified, e.g., by the nature or the difficulty of the investigation. For instance, some aspects of security, say e.g. “information flow,” appear to be so system-sensitive as to make little sense in a less than specific framework. Proposals of such kind should however carry a promise to discover general principles and learn lessons of wider applicability, to be suitable for this call. Ideally, projects should have the potential to radically change the way we do computing and, in due time, to influence key socio-economic developments that could improve our quality of life.

Research proposals should tackle all the four issues of *security*, *resource usage and management*, *scalability,* and *distribution transparency*, intended in a broad sense and in the context of global computing, as they are pivotal to realising the aim of the initiative. Here, security can be understood to encompass issues of trust and privacy as well as language-based security/safety, and topics such as authentication, authorisation, secure communication and non-repudiation. Distribution transparency should be understood to include topics such as location and access transparency as well as failure, migration and scaling transparency.

The proposed research should aim to devise theories and techniques concerned with these four issues that are either applicable to global computing in general, and can therefore form the conceptual backbone for the whole initiative, or that are applicable to specific classes of global/overlay computers, where this is fully justified. In doing so, the research is encouraged to find connections between overlay computers, address common characteristics, and provide common solutions.

Research themes of interest to the initiative include, but are not limited to:

- Methods and infrastructures for trust;
- Models of interaction and cooperation;
- Abstraction mechanisms;
- Components and modularity;
- Programming languages concepts and support;
- Validation and verification;
- Algorithmic principles;
- Autonomy, adaptivity, and self-organisation;
- Interoperability;
- Design support and software techniques.

5. THE APPROACH TO THE RESEARCH

Proposals must show a substantial integration between theory and systems building, following a foundational approach typical of Computer Science research. In the past, through instruments such as abstraction, virtualisation, modularity, encapsulation, software architectures, and abstract machines, Computer Science has succeeded in getting to the core of issues and separating the essential from the incidental. Such tools have brought fundamental advances to date, and will be at the heart of the technologies enabling our vision. Combined with strength in system building and experimentation, network technologies and foundational research, these tools will constitute the core of "Global Computing" projects research approach.

In practical terms, theories unrelated to the core issues of global and overlay computers are not suitable for this call, and neither is implementation work not supported by solid theories, nor advancing the general understanding of global computers.

Experiments and practical evaluation are expected to be embedded in the projects' workplans from the very start. The use of theory must be focused, of clear relevance to the practice of global computing, and of justified feasibility. The favourite vehicle to show such relevance is to illustrate it, e.g. by system building, as part of the project workplan. Systems must implement, or experiment, or build upon, or at very least be informed by theories and innovations arising from work on global computers. The favourite way to show the connections with foundational work is via an integrated workplan. In general, however, the spectrum of eligible projects ranges from projects in which theory leads and system work serves as proof-of-concept, to projects in which system building leads and theory serves a specific supporting role. The topic of the project together with its deliverables and workplan must justify the approach taken and the specific interplay between theoretical and system work.

6. CHARACTERISTICS OF A SUCCESSFUL PROPOSAL

The Global Computing II proactive initiative will be implemented through Integrated Projects (IPs), and possibly a Network of Excellence (NoE). Proposers should read carefully the documents related to FET proactive initiatives (http://www.cordis.lu/ist/fet/int-p.htm) and the description of the IP and NoE instruments, both in general terms (http://europa.eu.int/comm/research/fp6/instruments_en.html) as well as in the frame of FET (http://www.cordis.lu/ist/fet/int-n.htm for NoE).

With respect to IPs, and as common to the FET research activities, proposals will be evaluated on the basis of the innovativeness of the ideas they pursue, the risk and adventure in research they present, and the potential to give rise to new computing paradigms and achieve major, rather than incremental, advances.

In summary, the ideal IP proposals should:

- have the potential to radically change the way computing in global environments is done;
- aim at providing both foundational and practical advances on suitably large classes of global computers, and the integration of methods and concepts needed to lead to results that teach us general principles (i.e., with at least a perspective of applicability to multiple global computers);
- present a substantial integration of theory and practice, following a foundational approach;
- identify clearly the class of global or overlay computers they focus on, and make sure that such a class is not unnecessarily restrictive;
- have clear objectives related to the list of core issues and challenges of global computing mentioned above and expressed in terms of solutions realistically implementable in a global scenario;
- identify clearly the expected impact and concrete results, and have a set of clear methods and criteria to evaluate success