Available online at www.sciencedirect.com
ScienceDirect
Theoretical Computer Science 389 (2007) 337–340
Theoretical Computer Science
www.elsevier.com/locate/tcs

Semantic and logical foundations of global computing: Papers from the EU-FET global computing initiative (2001–2005)

Preface

Global computing refers to systems of interacting computational agents exhibiting the following characteristics:

- they are autonomous, in that their activity is not centrally coordinated;
- they are mobile, and their number, connectivity, and communication bandwidth may change during computation;
- they have a limited knowledge of their working environment, and no global information about the state of the computation is available, nor is a globally trusted authority.

These themes have been addressed in a series of international meetings, including Foundations of Wide Area Networks (FWAN 2002), Foundations of Global Computing (FGC 2003, FGUC 2004), and Trusted Global Computing (TGC 2005–2007). Notably, global computing has been the topic of EU funding initiatives (FET Global Computing 2001–2005 and FET Global Computing II 2005–2009), and formed the basis of one of the UK Grand Challenges in Computing, Global Ubiquitous Computing (http://www-dse.doc.ic.ac.uk/Projects/UbiNet/GC). This volume collects a representative selection of the best papers from the first EU-FET initiative, as illustrated below.

FET Global Computing (2001–2005) consisted of 13 projects clustered together in three main, essentially disjoint groups (cf. Fig. 1).

Foundations of Networks and Large Distributed Systems. This cluster focused on global computing infrastructures, dealing with issues such as scheduling algorithms and mechanism design for resource sharing between interacting agents, specification, verification and reasoning about such systems. The cluster's projects were: CRESSCO, Critical Resource Sharing for Cooperation in Complex Systems; DBGLOBE, A Data-Centric Approach to Global Computing; FLAGS, Foundational Aspects of Global Computing Systems; and SOCS, Societies of Computees.

CRESSCO researched foundational aspects of managing critical resources (such as bandwidth, frequency, energy, processor time) in GC infrastructures connecting very large numbers of independent, possibly mobile and selfish agent entities. FLAGS aimed at a general set of design principles and mechanisms for global computing systems of such agents competing for resources. CRESSCO and FLAGS pioneered the application of economic game and auction theories to resource sharing in distributed computing systems. DBGLOBE focused on extending current database technology to address data management requirements in large-scale networks of mobile entities. Their approach is data-centric, and developed a new language for integrating documents and queries, called Active XML. Finally, SOCS investigated computational and logical models for analysis and verification of societies of agents, with emphasis on techniques and models adapted from logic programming.

Analysis of Systems and Security. The focus here is on formal techniques such as type theory, proof-carrying code, and formal models for trust management for validating system properties of safety and security. The cluster was composed of: DART, Dynamic Assembly, Reconfiguration and Type Theory; MRG, Mobile Resource Guarantees; MYTHS, Models and Types for Security in Mobile Distributed Systems; PROFUNDIS, Proofs of Functionality for Mobile Distributed Systems; and SECURE, Secure Environment for Collaboration among Ubiquitous Roaming Entities.

0304-3975/$ - see front matter
doi:10.1016/j.tcs.2007.09.017

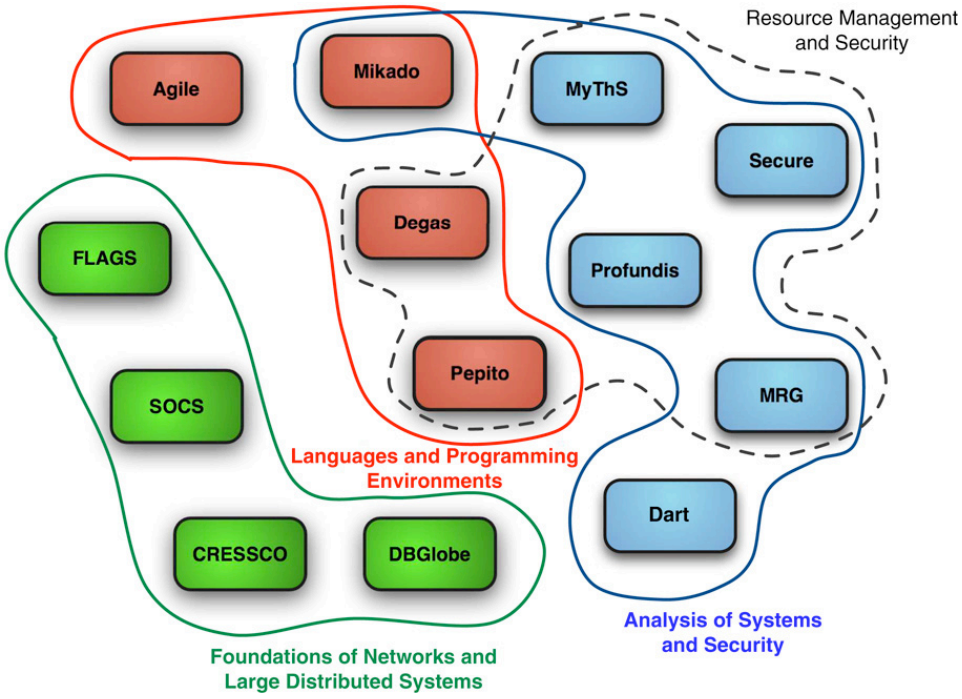

Fig. 1. EU-FET Global Computing projects and their clustering.

DART developed formalisms and techniques for supporting arbitrary interleaving of type-checking, meta-programming, and normal computational activities while retaining safety; a notable outcome is their work on "smart recompilation." MRG set out to extend the concept of proof-carrying code to include guarantees of resource usage, both in programmer-level languages and in low-level assembly and byte code. MYTHS explored type-based theories of security for mobile and distributed systems, for the analysis of security protocols, and for the type-oriented manipulation of XML data. Both very small sized, MRG and MYTHS achieved significant breakthroughs in their respective fields. PROFUNDIS exploited formal models and verification to explore key issues in mobile distributed systems, such as security and authentication, access rights control and resource management, as well as automatic support for the design and implementation of such features. The verification and testing tools produced by PROFUNDIS are regarded as valuable assets for the community. Finally, SECURE proposed a novel formal model of security based on the notion of trust and of algorithms for the dynamic management of trust, forming the so-called SECURE trust engine.

Languages and Programming Environments. This third cluster has used and extended theoretical foundations such as those developed in the previous two clusters for designing suitable software tools and frameworks. The cluster's projects were: AGILE, Architectures for Mobility; DEGAS, Design Environments for Global Applications; MIKADO, Mobile Calculi Based on Domains; and PEPITO, Peer-to-Peer Implementation and Theory.

AGILE explored a UML-based architectural approach to software engineering for mobile systems, including theoretical foundations as well as pragmatic techniques for deriving software architectural components corresponding to UML models. DEGAS derived specific process-algebra-based performance models from annotated UML application models both as a design principle and as an implementation technique. They delivered a set of tools for mechanising the process, notably the environment "Choreographer", which was very well received by the community. MIKADO – which can also be seen as belonging to the 'Analysis' cluster – researched formal programming models for global computing based on the notion of "domain", and delivered new prototype programming languages and runtime environments. Finally, PEPITO investigated the foundations of scalable distributed computing based on the peer-to-peer computing paradigm, exploring algorithms and programming languages. They developed a generic, language-independent distribution platform for P2P computing which provides a practical proof for server-less systems.

As can be observed in Fig. 1, a theme that cuts across clusters was security and resource management.

An invitation to read

The following papers are included in this special issue. Some are extended versions of papers that were published in the proceedings of TGC 2005, which was published by Springer as volume 3705 in the Lecture Notes in Computer Science series, edited by Rocco de Nicola and Davide Sangiorgi, and other conferences. All have been refereed to journal standard.

Formalising Java RMI with Explicit Code Mobility, *by Alexander Ahern and Nobuko Yoshida.* This paper formalises the key components of Java Remote Method Invocation (RMI) as a distributed and object-oriented process algebra, with primitives for distributed programming and explicit code mobility. The formalisation captures the crucial processes of serialising and unserialising data, passing data, code and class definitions across the network, and the semantics of synchronisation operations. The formalisation is shown to be type-safe, and it is used to prove the correctness of several optimisations for distributed programs. This work contributes to the *Analysis of Systems* activity.

A Program Logic for Resources, *by David Aspinall, Lennart Beringer, Martin Hofmann, Hans-Wolfgang Loidl, and Alberto Momigliano.* This paper defines a sound and complete logic for proving statements about resource consumption in a subset of the Java Virtual Machine Language, designed to be used in a proof-carrying code scenario in which mobile programs are equipped with formal evidence that they have good resource behaviour. A second logic, used to express termination, is also shown to be sound and complete. Both logics and all meta-theoretical results have been formalised in the Isabelle/HOL theorem prover as part of a prototype implementation of proof-carrying code applied to resource bounds in the MRG project.

A Semantic Framework for Open Processes, *by Paolo Baldan, Andrea Bracciali, and Roberto Bruni.* Taking the lead from the idea that Global Computing systems – and more generally open networks – can be formalised as coordinators, i.e. terms with holes to be plugged in by components and further coordinators, the authors set out to develop a general theory of such terms. They first use symbolic transition systems to describe coordinators operationally, and then present a novel bisimulation semantics that preserves the coordinators' openness, and fully accounts for it. This work sits quite well with the research approaches of both the clusters *Languages and Programming Environments* and *Analysis of Systems and Security*.

A Formal Semantics for Protocol Narrations, *by Sébastien Briais and Uwe Nestmann.* Protocol narrations are intuitive descriptions of crypto-protocols as sequences of encrypted exchanges among participants. The authors propose a formal operational semantics for protocol narrations and – on the basis of such a semantics – explain and justify a natural and precise translation of narrations into spi-calculus which led to the realisation of automated tools and the analysis of case studies, also presented in the paper. This work exemplifies quite significantly the *Languages and Programming Environments* cluster.

A Framework for Analyzing Probabilistic Protocols and its Application to the Partial Secrets Exchange, *by Konstantinos Chatzikokolakis and Catuscia Palamidessi.* This paper presents a new probabilistic process algebra which extends the pi-calculus with probabilistic constructs in order to specify randomised security protocols. Its semantics is given in terms of Segala–Lynch probabilistic automata, and a may-testing pre-order between processes is defined. This is used to specify and verify the correctness of the Partial Secrets Exchange protocol. This is another contribution to the *Analysis of Systems* activity.

The Complexity of Fixed Point Models of Trust in Distributed Networks, *by Karl Krukow and Andrew Twigg.* Global trust in the presence of delegation of authority can be modelled as a fixed point of local policies taken ideally over the entire global network. This paper explores distributed algorithms for *computing* or *approximating* global trust efficiently, and determines their lower bounds and computational complexity. The key notion adopted by the authors is that of *distributed proof-carrying request*. This work was produced in the context of the SECURE project in the *Analysis of Systems and Security* cluster.

Acknowledgements

We wish to thank the colleagues who acted as *referees* for this volume. They are: Chiara Bodei, Frank de Boer, Roberto Bruni, Michele Bugliesi, Derek Dreyer, Sergio Maffeis, Alberto Momigliano, Uwe Nestmann, Joël Ouaknine, Paula Quaglia, Julian Rathke, Bernard Reus, Peter Sewell, Martin Steffen, David Walker, and Nobuko Yoshida. Our description of the EU-FET GC clusters is inspired by the EU final report for FET GC (cf. http://cordis.europa.eu/ist/fet/gc.htm).

Donald Sannella*
University of Edinburgh,
Edinburgh, United Kingdom
E-mail address: dts@inf.ed.ac.uk.

Vladimiro Sassone
University of Southampton,
Southampton, United Kingdom
E-mail address: vs@ecs.soton.ac.uk.

* Corresponding editor.