



European Cooperation  
in the field of Scientific  
and Technical Research  
- COST -

Brussels, 30 January 2012

-----  
Secretariat  
-----

**Full proposal reference oc-2011-2-10054 for a COST new Action**

---

Subject: Full proposal for a new COST Action:  
BETTY: Behavioural Types for Reliable Large-Scale Software Systems

---

**Proposer:** Dr. Simon GAY  
University of Glasgow  
School of Computing Science  
Sir Alwyn Williams Building  
18 Lilybank Gardens  
Glasgow - UK  
Simon.Gay@glasgow.ac.uk

**National Coordinator:** [\*]

**Domain Committee:** Information and Communication Technologies

---

*[\*] Will be completed by the COST Office*

**DRAFT**  
**MEMORANDUM OF UNDERSTANDING**  
**For the implementation of a European Concerted Research Action**  
**designated as**

**COST Action**

**BETTY: Behavioural Types for Reliable Large-Scale Software Systems**

The signatories to this "Memorandum of Understanding", declaring their common intention to participate in the concerted Action referred to above and described in the "Technical Annex to the Memorandum", have reached the following understanding:

1. The Action will be carried out in accordance with the provisions of document COST 299/06 "Rules and Procedures for Implementing COST Actions", or in any new document amending or replacing it, the contents of which the Signatories are fully aware of.
2. The main objective of the Action is [\*]
3. The economic dimension of the activities carried out under the Action has been estimated, on the basis of information available during the planning of the Action, at [\*] Euro [\*] million in [\*] prices.
4. The Memorandum of Understanding will take effect on being signed by at least five Signatories.
5. The Memorandum of Understanding will remain in force for a period of years, calculated from the date of the first meeting of the Management Committee, unless the duration of the Action is modified according to the provisions of Chapter V of the document referred to in Point 1 above.

*[\*] Will be completed by the COST Office*

## **A. ABSTRACT & KEYWORDS**

### **A.1 ABSTRACT**

Modern society is increasingly dependent on large-scale software systems that are distributed, collaborative and communication-centred. Correctness and reliability of such systems depend on compatibility between components and services that are newly developed or may already exist. The consequences of failure are severe, including security breaches and unavailability of essential services. Current software development technology is not well suited to producing these large-scale systems, because of the lack of high-level structuring abstractions for complex communication behaviour.

This Action will use behavioural type theory as the basis for new foundations, programming languages, and software development methods for communication-intensive distributed systems. Behavioural type theory encompasses concepts such as interfaces, communication protocols, contracts, and choreography. As a unifying structural principle it will transform the theory and practice of distributed software development.

The significance of behavioural types has been recognised world-wide during the last five years. European researchers are internationally leading. There is an urgent need for European co-ordination to avoid duplication of effort, facilitate interactions among research groups, and ensure that the field proceeds efficiently from academic research to industrial practice. This Action will provide the co-ordination layer and leverage the efforts of European researchers, to increase the competitiveness of the European software industry.

### **A.2 Keywords**

distributed software infrastructure, software development methodology, foundations of programming languages, service-oriented computing, behavioural types

## **B. BACKGROUND**

### **B.1 General background**

Modern software systems operate on a large scale. They are distributed, collaborative and communication-centred, and are an essential part of the technological infrastructure of society. Many critical services are offered via the Internet and depend on other network technologies such as Clouds. These services include commercial operations such as banking, e-commerce, social networking and document sharing applications, as well as public-sector systems such as infrastructures for local, national and trans-national e-governments, e-healthcare and e-science. They are long-lived and run continuously, constantly collecting

information and dynamically acquiring new functionalities, while network-aware devices that access these services, including mobile phones and tablets, adapt dynamically to their environment.

Recent years have seen the development of a new paradigm for the development of distributed software systems, based on the internalisation of the concept of a distributed service, often called service-oriented computing. In this paradigm, the computational behaviour of a system is organised into a collection of services, which run concurrently and interact with each other, instead of being organised into a collection of functions which run sequentially. This collection of services implements another service. For example, a login service for a popular web portal is usually implemented by multiple distributed components realising more specialised functions, such as an external identity provider, a browser, an internal identity registry, and an internal authentication service.

How can improper, undesired behaviour be prevented in complex and socially critical distributed applications? Protocol incompatibilities, malicious resource usage, security breaches and deadlock/livelock can have wide-ranging consequences, from temporary service outage to information leakage to exploitation of security vulnerability by criminal organisations, affecting a large number of users.

In traditional sequential programming, the systematic and rigorous application of data types and type checking has yielded huge improvements in software reliability and programmers' productivity. The pinnacles, so far, of type theory in mainstream programming languages are the sophisticated generic type systems of Java and C#, where programming abstractions are built in at a fundamental level. Data types abstract and specify the behaviour of permitted *operations* on data. Even if types are not used for static checking, they often play a fundamental role in runtime checking (as in Python, a language widely used for server-side applications). Because data types offer basic abstraction of programs' behaviour, cutting edge logical verification tools such as Java Modelling Language (based on Java) and Spec# (based on C#) articulate their logical assertions for programs by building on the type abstractions of the respective programming languages.

Thus data types offer an effective basis of programming and verification for sequential data processing. For building distributed, service-oriented computing software, a corresponding structural view requires *behavioural types*, which specify the permitted *interaction* within a distributed system. The behavioural type of a service specifies its expected patterns of interaction using a simple, expressive type language, in a way that can be used to determine automatically whether a service interacts correctly with other applications. At present, although there is a significant body of theoretical work on behavioural types and a growing collection of prototype programming language designs, it is not yet possible for industrial software developers to exploit behavioural types in the design, analysis and implementation of service-oriented computing systems.

The goal of this Action is to transform the theory and practice of service-oriented software development by the systematic application and unification of behavioural type systems. The result will be new foundations for service-oriented systems, new programming languages whose designs are informed by behavioural types, new program analysis tools based on behavioural types, and new software development methodologies based on these foundations, languages and tools.

In order to achieve this goal, significant research challenges need be tackled.

- How can mathematical theories of behavioural types be unified, capturing different aspects of software behaviour at levels ranging from bare protocols to more refined properties such as security policies, and together offering a consistent, comprehensive foundation for the large-scale distributed infrastructure which constitutes global service-oriented computing?
- How can behavioural types be integrated into fully-fledged language designs, compilers and tools, both as new production-level programming languages and as enhancements of widely used programming languages, so that software developers can effectively structure their code around behavioural types?
- How can behavioural types be used in practice as a basis for specifying and verifying the correctness of software components in large-scale, long-lived distributed systems, both at development time and at runtime, so that effective validation of components and services is possible?

Recently there has been a rapid growth of activity in the field, and an accelerated expansion of the research community. There is a danger that a proliferation of alternative and incompatible approaches will lead to duplication of effort, hindering the transition from foundational theories and prototype programming languages to industrial-strength tools and methodologies. There is therefore a need for direction and co-ordination. On the basis of a large number of research groups and many national research projects, these activities will reach a tipping point towards their practical world-wide adoption by having a supporting structure within which to organise the community, with funding for strategic meetings and scientific exchanges. A COST Action is an excellent mechanism for this purpose.

## B.2 Current state of knowledge

### B.2.1 Current best-practice in industry

Some of the better known existing large-scale distributed infrastructures include the backends of Amazon, Google and Microsoft Windows Live. Widely-practised methods for building distributed services in such large-scale infrastructures include:

- **Remote procedural call (RPC)**, or its object-oriented variant, remote method invocation (RMI). For example, a major part of Google's backend is implemented in C++ with RPC, together with the serialisation formats of Google Protocol Buffer.
- **Web service technologies through HTTP/SOAP** (Simple Object Access Protocol), which often use WSDL (Web Service Description Language, a simple interface language). While the shape of interactions is essentially the same as RPC, all resources are accessible through HTTP, a single application-level protocol.
- **Messaging**, which often uses messaging brokers, centring on a publish-subscribe model. This is based on asynchronous message passing and is especially useful when rapid and reliable dissemination of data is necessary.

Another traditional method for integrating distributed components is to have implicit data flow through shared state. However, the explicit communication-based frameworks listed above have the merit that the infrastructure is clearly divided into distributed components with explicit interfaces; this makes it easier to export services, change service interfaces, and transplant a system into a new computing environment such as clouds. For example, during the last decade Amazon converted all of the software components in its infrastructure so that they are based on explicit communication interfaces centring on RPC/RMI, HTTP/SOAP and messaging. The framework Amazon has achieved is commercially known as SOA (Service Oriented Architecture).

RPC, RMI and HTTP are all based on the request-reply interaction pattern, which is a distributed version of a function call to data structures and objects. The restriction to the request-reply pattern causes several problems when, like Amazon, one imposes strict communication interfaces among distributed components. These problems include the lack of description or abstraction for overall interaction patterns (e.g. a sequence of interactions constituting a workflow is divided into numerous request-reply interactions) and latency due to the waiting time needed in request-reply. They are closely related to the lack of uniform programming abstractions capturing, among others, the above three forms of communication interfaces and beyond, and the lack of clean programming language support for distributed computing.

### **B.2.2 Production-level programming languages for building distributed services**

In major programming languages usable for production-level application development, such as Java, C#, Python, Javascript and C++, there is no language-level support for distributed communication except as libraries (APIs). An exception is remote method invocation (RMI), which is part of the language specification in Java and C#, but this is one of the lowest-level communication primitives. This lack of language-level support means that communication patterns are not easily visible to the compiler or to program analysis tools. For example, implementing a login service requires a complex sequence of RMIs on constituent services. If this sequence is represented *only* as a series of API calls, with no explicitly-defined high-level structure, then programming errors are easy to make and difficult to detect.

The only relatively widely-used production-level programming language which centres on message passing is Erlang, which is roughly based on Hewitt's actor model. The main programming-level abstraction is the queue, which concerns only messages to a single actor. Another widely-used system, especially in scientific computing, is MPI (Message-Passing Interface). MPI is provided as an API (procedure library) for standard languages such as C and Fortran, and so it suffers from the lack of high-level abstractions mentioned above in relation to RMI. Furthermore, the MPI libraries are tailored for typical computational patterns in scientific number-crunching and are unsuitable for general-purpose distributed programming. Finally, although it is a research programming language, Occam-Pi is highly efficient but is again essentially based on synchronous handshake communication.

Recently Google has been developing two general-purpose programming languages based on its analysis of internal large-scale distributed software development. The Go programming language (<http://golang.org/>) is a system-level, garbage-collected, type-safe programming language which incorporates handshake communication. It is intended as an advance from the C programming language. Dart (<http://www.dartlang.org/>) is a class-based object-oriented language with pluggable static typing and with actor-like queue-based message passing communication. It is intended as an advance from Javascript, and is being developed as part of a large programming environment project. In both languages, communication is provided by language primitives, not by libraries. The underlying communication model used by Dart, based on asynchronous communication, is close to the lower-level communication model adopted in many languages based on behavioural types (see below).

Go and Dart show that in the context of the development of large-scale distributed applications, engineers

have started to realise the need to treat communication seriously. But even in these latest programming languages, there is no language-level abstraction or specification mechanism for communication patterns beyond handshake or request-reply.

### **B.2.3 Origins of behavioural types**

Types articulate computation in a given paradigm. The traditional notion of types offers abstractions for data, objects and operations on them, which are the key building blocks of computation in traditional programming languages. The basic form of behavioural types, specifying interaction patterns in distributed services, articulates the ways in which interactions are programmed and performed. From this basis, it is possible to specify more refined behavioural types and logical assertions. Thus behavioural types are not only a basis for verification, but also a way to organise computation in service-oriented, distributed computing systems. By centring on interaction, behavioural types are often automatically *compositional*. The strategy is to first verify whether or not the interactions of each service satisfy their behavioural types. After that, the behavioural type of a composite service, built from interactions between components, can be verified without re-examining the components.

Type theories for distributed computing systems are a new field largely initiated by the study of behavioural types. They are built on the formal foundations provided by process calculi, such as ACP, CSP, CCS, pi-calculus, and higher-order pi-calculus. Process calculi offer a rigorous mathematical basis upon which communication behaviours of software systems are represented, reasoned about and analysed. Combined with other formalisms such as Petri Nets, they offer a general framework in which to mathematically capture diverse forms of interaction, rigorously and in a distilled form, including synchronous, asynchronous, untimed, timed, and others.

For theories of behavioural types, process calculi with channel passing capability, starting from pi-calculus and higher-order pi-calculus, play a central role. They offer a rich basis for behavioural types because their tiny syntax (a fully expressive asynchronous version of the pi-calculus can be defined by a standard BNF-style grammar in six lines) can nevertheless express fully the whole range of interactional behaviours realisable by programming languages, thus serving as a rich basis for investigation.

However, calculi and programming languages differ in that the latter also need to offer programming abstractions (structuring principles for software) and efficient executions, among other pragmatic concerns. Calculi offer a basis upon which these ideas can be formally expressed and examined, but it is also necessary to build a comprehensive programming framework which offers effective structuring principles and other foundations for building distributed, service-oriented computation.

Research into behavioural types and the associated programming primitives has been addressing these concerns.

### **B.2.4 State-of-the-art in behavioural types**

Representative technical approaches to behavioural types include:

- *Types for sessions* (session types), which are based on the idea that interactions among distributed components and services can naturally be divided into multiple conversations (sessions), each with a specific protocol. The division into sessions enables tractable type-abstraction, making it easy to specify protocols and validate programs.
- *Contracts for interactions*, which give fine-grained specifications stipulating reciprocal responsibilities among interacting parties, often using labelled transition relations. Contracts can be combined using

- various algebraic operators to obtain larger contracts.
- *Behavioural security specifications*, where security properties are specified as behavioural types, often elaborating simpler notions of behavioural types. This enables an explicit specification of information flow as reciprocal responsibilities among communicating parties in a conversation, offering a uniform framework to specify distinct security properties such as secure information flow and access control, centring on types for interaction.

There are diverse sub-approaches in each domain: for example, *conversation types* offer an alternative way to describe sessions, as a set of interactions inside a scope, giving an abstraction of interaction in a common medium. There are also many notions of contract, capturing diverse behavioural properties. Note that contracts differ from simply specifying (say) temporal logical formulae for interactions. They are specified as reciprocal responsibilities among communicating parties, and this global form of contract can be related to individual responsibilities by a process known as *projection*, which is also used in session types for more than two participants.

Thus behavioural types and their associated theories offer a new shape of specification languages centring on interactions and reciprocal responsibilities. But they also enjoy a rich linkage with traditional theories. Contracts are closely related to existing temporal and modal logics, as well as their ramifications. The type structure of a session type is related to an automaton for communicating systems, since it specifies a series of actions among two or more parties. Furthermore, behavioural types give rise to the diversity of type structures cultivated in logics and programming theories, such as polymorphism and dependent types.

All approaches to behavioural types have commonality in that they are closely associated with programming primitives. This is also a key factor for the effective validation of programs against specifications given by behavioural types because the programming primitives offer linkage between programs and type specifications. For example, session types are associated with primitives to initiate a session, communicate through sessions, receive through sessions, and so on. This linkage then leads to a type discipline (a method to validate programs against given types), and its key properties such as principal typing and type/protocol safety (i.e. interactions follow a stipulated protocol), is established. The development of type disciplines closely follows that of traditional typed lambda-calculi and programming languages such as ML and Java. Contracts and logical specifications can then offer richer specifications by elaborating the base type disciplines, which may also make use of spatial localities and their composition for modular specifications.

Many of these studies were initiated in the context of the Sensoria project (Software Engineering for Service-Oriented Overlay Computers), which ran from 2005 to 2010 as part of the EU FP6. The Sensoria Consortium consisted of 14 universities, 1 research institute and 4 companies from 7 countries: Italy, Germany, United Kingdom, Portugal, Denmark, Hungary and Poland. Several of its members are also participants in this Action.

### **B.2.5 Programming languages and tool development based on behavioural types**

These studies give a foundation for the design of new programming languages based on and supporting behavioural types, and for an enhancement of traditional programming languages with behavioural types. From the development viewpoint, the key merits of using behavioural types may be summarised as follows.

- Clean and general structuring principles for implementing, composing and managing communicating distributed systems (distributed services). Given the nature of distributed programming, this aspect covers not only a source program for each endpoint but a configuration of many components, which may be written in different programming languages.



- Suggestions for the design of execution machinery (runtime system) for programming languages for distributed services, through the linkage between programming primitives and behavioural types.
- Simple and intuitive specifications of key properties, which are validated either statically (at compilation time) or dynamically (at runtime), the latter enabling a change of specifications over time. Both static and dynamic validation can make use of efficient algorithms whose correctness is established through the underlying theories.

In recent years, several research projects have started to develop new prototype programming languages based on behavioural types, as well as extensions to existing programming languages. In the latter thread, the languages being extended include Java (the SJ project), ML, BPEL, Scala (the Session Scala project) and C. Microsoft has developed the language Sing#, which uses elements of behavioural types (session types) for systems level programming in the Singularity operating system.

Research prototype languages are often equipped with a full type checker, as well as a lightweight runtime system. Reflecting their experimental nature, runtime and tool support is often elementary. However, through their development, many important results have been obtained regarding, among others: programmability; type checking and type inference algorithms; runtime design; integration of behavioural types with traditional programming language constructs and their validation methods; dynamic validation mechanisms; and design and development experience. Some of these experiments are being further developed to offer more robust production-level support of safety assurance based on behavioural types, often associated with existing frameworks and languages, such as SOA, BPEL, and messaging infrastructure.

### **B.3 Reasons for the Action**

Society is becoming dependent on large-scale distributed software systems, and there is an urgent economic and societal need for improved technologies and methodologies for designing, analysing and building such systems. These technologies and methodologies must be based on a sound theoretical foundation. Many foundational theories and technologies are in place, but delivering them to the software industry and achieving the goal of more reliable large-scale distributed systems will require a co-ordinated effort to ensure a flow of results from foundations to programming languages to applications.

Specific objectives of the Action, at the co-ordination level, include:

- Developing a clear picture of where particular expertise is located, in order to be able to link problems with solutions.
- Supporting short-term scientific missions, especially for PhD students and early-career researchers, in order to facilitate the development of lasting research collaborations which can be supported by national funding.
- Organising scientific meetings and workshops.
- Maximising the value of existing industrial collaborators by putting them into contact with a wider range of participants in the Action.

The networking activities of the Action will result in a strengthened and consolidated European research community with effective working practices to ensure that foundational research informs and is incorporated into effective programming languages and tools, which are in turn delivered to the software industry for use in practical projects.

## B.4 Complementarity with other research programmes

The Action will co-ordinate research that is taking place in around 20 current nationally-funded projects in various European countries.

Three current EU FP7 projects have concerns that are connected with those of this Action to some extent.

- ANIKETOS (Secure Development of Trustworthy Composable Services, 2010 - 2014)

ANIKETOS has a very practical focus on the security of service-oriented systems. Its consortium does not include any of the participants who have been involved in planning this Action. It is likely that the Action will be able to offer appropriate theory to support some of the work of ANIKETOS; conversely, the industrial members of the ANIKETOS consortium should be a good source of practical case studies and potential application areas for the results of this Action. It will therefore be important to develop a relationship between this Action and the ANIKETOS project.

- ASCENS (Autonomic Service-Component Ensembles, 2010 - 2014)

The ASCENS project aims to design methodologies to build “ensembles” of components in which there is a strong emphasis on both service-oriented architectures and a high degree of autonomy, self-awareness and adaptability on the part of components. The project includes the development of theoretical foundations and formal verification techniques, and these goals are very much in line with those of this Action. In particular, there is a clear possibility of applying behavioural type theory to the application domains considered by ASCENS. Some of the participants in this Action are also members of the ASCENS project. It will be useful to build on this connection in order to develop mutually beneficial collaboration.

- HATS (Highly Adaptable and Trustworthy Software using Formal Models, 2009 - 2013)

HATS also aims to improve the reliability of large-scale software systems. Its approach is to develop a rigorous foundation for informal software engineering practices based around the concept of software product families. Behavioural specification formalisms, and tools based on them are important aspects of this project, but it contrasts with the approach of this Action: it does not propose behavioural *types* as a new structuring concept for programs. Some members of the HATS consortium are participants in this Action, and this will be helpful in understanding the results of HATS during the initial stages of this Action.

COST Action IC0901 (Rich Model Toolkit - An Infrastructure for Reliable Computer Systems, 2009 - 2013) aims to improve the reliability of computer systems by developing new frameworks for automated system analysis. In broad terms, the goal of improved reliability is similar to that of this Action. The difference is that this Action has a strong focus on developing new programming languages, whereas COST IC0901 develops external tools to be used with existing programming languages. Nevertheless, it will be useful to establish contact with COST IC0901.

## C. OBJECTIVES AND BENEFITS

### C.1 Aim

The aim of the Action is to use behavioural type theory as a basis for improved programming languages and tools for the implementation of reliable large-scale distributed software systems.

### C.2 Objectives

The main objective of this Action is to develop the domain of certified software for global services, providing languages and tools for automatically checking behavioural properties of concurrent and distributed software systems, specified with simple yet expressive type languages. Successful research in this area will have a deep and broad impact on the practice of software development, and on the scientific theories underlying our understanding of distributed computing systems.

The Action has the following sub-objectives:

O1. Co-ordinate European research activity on the theory and application of behavioural type systems, and the deployment of programming languages and tools based on them, organized into the following themes which will each have a Working Group.

- Theoretical Foundations
- Security
- Programming Languages
- Tools and Applications

O2. Build an effective working community of European researchers in this area, by means of the following activities:

- Regular meetings of the Management Committee and the Working Groups
- Annual workshops (following a successful workshop in April 2011)
- Training schools for PhD students and early-career researchers
- Short-term scientific missions between institutions, especially for PhD students and early-career researchers
- Encourage the use of bi-national joint PhD programmes between institutions in the Action.
- Actively seeking to increase the membership of the Action

O3. Encourage the industrial adoption of advanced programming languages and tools, by working with existing industrial contacts and developing new ones.

### **C.3 How networking within the Action will yield the objectives?**

The scientific objectives will be achieved through close collaboration between the participants in the Action. This will include exchange of PhD students and researchers through Short-Term Scientific Missions, and annual workshops that will report on the latest research results and initiate new collaboration activities. The Action will enable PhD students and early-career researchers to develop and share their expertise, and will contribute to creating new connections between established researchers.

The following specific objectives, with quantitative targets where appropriate, are planned. The targets will be used by the Management Committee to monitor the progress of the Action.

At Action level:

1. The Action will involve at least 36 academic and research institutions in 14 countries (30 institutions in 12 countries have already expressed interest in this Action).
2. The activities of the Action will include participation by at least 9 industrial collaborators (6 companies have already expressed interest) and a large-scale academic development project. Several of the industrial collaborators are actively engaged in the deployment of behavioural types technologies in their development projects.
3. At least 6 collaborative, transnational projects on specific topics relating to the Action will be submitted to national and European programmes.
4. Management Committee meetings will be held every 6 months.
5. A project website listing upcoming and past activities, paper abstracts and presentations will be prepared by the Action participants.
6. There will be an annual workshop which serves as a meeting for all Action participants and as a peer-reviewed symposium.
7. There will be a total of at least 100 publications in international peer reviewed journals and conference proceedings as a result of the scientific projects and annual workshops of the Action.
8. There will be at least 120 Short-Term Scientific Missions within the Action.
9. There will be a final report that describes the outcomes and successes of the Action.

At the level of Working Groups:

1. Working Group meetings will be held annually, in conjunction with the annual workshops.
2. Each Working Group will produce a state-of-the-art report during the first year of the Action, and thereafter an annual progress report.
3. Working groups will organise special sessions, for example tutorials or focussed discussion meetings, at international conferences.

### **C.4 Potential impact of the Action**

This Action will co-ordinate European research in the field of behavioural types for concurrent and distributed software systems. It will facilitate collaboration and knowledge exchange between research centres across Europe, thus maximising the benefits of various national research programmes. This in turn will consolidate the internationally leading position of the European research community in this area.

The software industry (including the open source community), for which programming communication and

concurrency is one of the largest challenges at present and in the near future, will benefit from improved theories, programming languages, methodologies and tools, which will make it easier to develop correct and safety-assured software systems that rely on concurrency, distribution and communication. This benefit will be seen at a range of scales, from single applications running on multi-core chips through to global service-oriented systems.

Teaching programmes at the participating institutions will benefit from the possibility of offering advanced courses based on the latest research in the topics of the Action. This includes the potential for academic staff to deliver courses at other institutions, based on contacts developed through the Action and taking advantage of existing academic exchange schemes such as Erasmus.

Society will benefit from increased reliability of the essential technological infrastructure represented by concurrent and distributed computing systems.

### **C.5 Target groups/end users**

The benefits of the Action, described in the previous section, will be felt by the following groups and end users.

- The scientific community
- The software industry
- Society as a whole

This proposal has been prepared by members of the scientific community, in consultation with representatives of the software industry who are already actively involved in collaborative projects with participants in the Action.

## **D. SCIENTIFIC PROGRAMME**

### **D.1 Scientific focus**

The following specific aims provide a focus for research aimed at addressing the challenges presented by large-scale distributed software systems.

- Develop further theories of behavioural types so that behavioural types which treat different properties (such as interaction patterns, safety, liveness, information flow) can be uniformly specified, integrated and assured.
- Develop meta-theories of behavioural types, enabling description of the features and relationships of a range of theories without duplication of effort.
- Extend behavioural type theory with quantitative features, to allow specification and analysis of non-functional properties such as response time or resource cost, giving a uniform basis for policy specifications and enforcement including service-level agreements.
- Further develop theoretical and practical understanding of sessions, roles and relationships, generalising from an identity-based view to a role-based view of the behaviour of session participants, so that theories can offer solid foundations of a broad range of programming language methodologies.

- Develop theories, tools and methodologies for the specification and implementation of fault tolerance and error recovery, which are essential in service-oriented distributed systems.
- Further develop theories, programming methodologies and tools based on choreography, orchestration and global types, to specify, analyse and assure properties of the collective interaction among software components.
- Develop theories and tools for the analysis of fairness and liveness properties in service-oriented systems.
- Develop theories and tools for the integrated specifications of security properties including secure information flow in a broad sense, and access control based on behavioural types, so that it enables effective runtime verifications.
- Design, develop and examine the dynamic runtime enforcement mechanism for behavioural types, so that we can specify new properties for distributed services at runtime and they can be enforced, with a formal assurance coming from theories of behavioural types.
- Develop mechanisms for negotiation, adaptation, extension and evolution of behavioural contracts and policies within long-lived environments for service-oriented software.
- Develop new production-level programming languages and programming environments based on behavioural types, building on the preceding experiments on prototype programming languages, with formal support for safety assurance and with emphasis on efficiency and scalability.
- Improve support for service-oriented programming in traditional programming paradigms, in particular object-oriented programming languages.

The Action is structured into four themes, each with a Working Group to co-ordinate its activity. These themes are (1) Foundations, (2) Security, (3) Languages, (4) Tools & Applications. Most of the aims listed above need to be addressed in relation to every theme, and the WGs will communicate to ensure that each topic is treated in a consistent way across all themes. The specific tasks of each WG, described in Section D.2, target these aims. Significant scientific innovations are expected through a close collaboration of the best experts in the field, in unification of foundational theories and their transfer into practical technologies and tools for software development.

## **D.2 Scientific work plan - methods and means**

As mentioned above, the scientific work of the Action will be co-ordinated by four Working Groups. Each WG will promote its theme, ensure consistency and avoid duplication of effort, and form links across themes and with industry.

### **Working Group 1: Foundations**

Communication-centred software and service-oriented computing raise new theoretical challenges, involving different research fields and cultural experiences. One of the central issues is the need to specify and validate systems whose components are developed and maintained by third parties who are not necessarily trustworthy, which are dynamically searched and assembled, and which evolve in time. Another central point is the nature of the properties users and administrators wish to ensure for such systems, which are often interactional, reciprocal and global in nature (e.g. developers wish to build a distributed application whose individually developed components interact without deadlock; or users wish to maintain fair and safe usage of a service as a whole).

These problems have been attacked on the basis of several different theories, mainly process algebras, logic

and type theory. The main issue of the WG is that of putting on the same footing apparently heterogeneous aspects of software development and verification, including concurrent and distributed programming disciplines, specification and verification of software components, static and dynamic checking of protocol and policy compliance.

The concepts of session, conversation and choreography have emerged as structuring principles for concurrent and distributed systems. Behavioural types merge ideas from process algebras and type theories, so that they provide a common foundation for structured software development and automatic verification, both statically and dynamically, which are well-suited to concurrent and dynamic systems. The tasks of this WG include:

1. To develop the theory of behavioural types clarifying how abstract process descriptions should be used as computational invariants and module interfaces.
2. To face the problem of expressivity of type languages w.r.t. safety, liveness and fairness of systems, in the context of both theoretical calculi and programming languages, investigated by WG3 of this Action, while keeping the effectiveness and efficiency of the verification procedures (including type systems) through which we can ensure such properties.
3. To define a theory of quantitative behavioural types, viewed as an abstract operational model accounting for non-functional aspects of the system's behaviour.

The present state of the art shows a plethora of formalisms, often sharing the same roots but with significant theoretical differences. The WG will aim at developing a common formal framework, allowing for comparison and study of different approaches currently developed by the participating teams. By focusing on the conceptual and technical problems emerging in the activity of the other WGs, WG1 is aimed at the integration and further development of different domain specific approaches, especially with respect to qualitative and quantitative aspects of system behaviour, and the development of new solutions taking advantage of their strengths in the respective domains of application.

## **Working Group 2: Security**

Enforcement of security properties is an important application domain of behavioural types. In particular, analysis of information flow, which is an essential aspect of both data confidentiality and data integrity and requires a dependency analysis, can greatly benefit from the use of behavioural types, since these also allow for a direct description of the causal dependencies in interprocess communication and the order in which data will flow between processes. From a different viewpoint, instead of specifying safety properties for individual, disparate messages, their specifications become much more amenable in terms of specifiability and effective validations when they are given based on protocols.

Applications of behavioural types to concrete case studies and development of type theories that target specific application domains are among the concrete results and unifying themes of the Action. They are therefore also the focus of this WG. The Action will achieve its objectives through the joint development of theoretical foundations and tools.

The activities include:

1. The development of type theories that make it possible to reason about security properties such as secrecy, authenticity, confidentiality and integrity for communication protocols, and the identification of other forms of security properties that are amenable to type-based analyses using new type theories.
2. The development of algorithmic techniques for type-based analyses of such security properties, for both static and dynamic validations.

The type theories and resulting algorithmic techniques should apply both to mathematical models of computation (such as process calculi) and to existing programming and specification languages.

The working group will build on experience gained from the following past successes of its experts:

- leading major efforts in the development of type theories for security issues;
- developing specialized tools for reasoning using types to reason about security issues.

The Action is therefore in a unique position to develop type theories of this kind and to foster collaboration within the community. It is important that WG2 coordinates its activities with the other WGs of the Action; this will happen through joint meetings and via the involvement of working group members in other working groups.

The infrastructure of WG2 will consist of a repository of type theories for security properties. These type theories will be applied and validated by tool development and case studies within WG4.

### **Working Group 3: Languages**

The study and development of behavioural type theories aims at the integration of behavioural types into mainstream programming languages and at their use as a guideline to design new languages that incorporate them. The working group focuses on research in which we develop theories, principles and key runtime mechanisms for programming with behavioural type systems as intrinsic structures. Actual development of programming environments and tool chains which can be used by developers is covered by WG4 described below. The Languages working group will pursue two distinct lines of research, aimed at the integration of behavioural types in:

1. Mainstream, general-purpose languages
2. Domain-specific languages.

Concerning mainstream languages, it is necessary to consider two orthogonal aspects, namely the programming paradigm and the degree of integration of behavioural types within the languages.

Regarding the former, this WG will target different programming paradigms, including, but not restricted to, object-oriented languages and functional languages. As programs are meant to manipulate data, of particular interest to WG3 are the XML and the JSON data formats.

- Object-oriented languages are relevant for their widespread adoption in the current development of software, for the wealth and popularity of tools that are available, and because objects nicely fit a distribution model to which behavioural types can be applied naturally.
- Functional languages are relevant for their qualities of being easily endowed with high-level type-theoretic and concurrent extensions, for their natural support to parallelism, and since they permit rapid prototyping.
- XML/JSON/YAML data formats are interesting for the role of XML and JSON as the de facto standard exchange paradigms in distributed, service-oriented computing, and need to be integrated as part of description languages for protocols and other global properties.

Within each of these programming paradigms, the WG will investigate the following, increasingly invasive ways to integrate the support for behavioural types:



- By means of explicit code annotations in the form of comments or pragma directives. In this way, (fragments of) behavioural type disciplines can be plugged into existing languages without requiring unreasonable changes to their syntax and semantics. Tools can then be developed for analysing and processing explicitly annotated code, in a similar way to what currently happens with many IDEs for program development. In order to minimise the amount of information provided by the programmer, it may be desirable to confine the use of explicit annotations to APIs dedicated to communication and synchronisation, and to devise suitable type inference techniques for automatically deriving type information in code that uses the APIs.
- By means of syntax extensions to be handled by pre-processing. In this way, high-level programming constructs (for communication, synchronisation, coordination) are more easily associated with high-level type information than the corresponding implementations in terms of low-level communication primitives.
- By means of conservative extensions of existing languages. Extensions of core research-oriented languages such as, for object-oriented programming, Featherweight Java, will be considered first. Then, results will be validated in fully-featured languages.

Regarding the first point, note that the applicability of behavioural types to traditional languages includes both statically and dynamically typed languages; in the latter case, for example Python, through runtime verification.

Concerning domain-specific languages (DSLs), the WG will follow two distinct research tracks:

- The adoption of behavioural types within existing DSLs designed for application areas and use-cases to which behavioural types naturally apply, such as service-oriented programming, distributed security protocols, distributed queries for semi-structured data.
- The development of new domain specific languages explicitly targeted at exploiting the characteristics of behavioural types and aimed at the coordination and/or orchestration of wide-area distributed services.

#### **Working Group 4: Tools & Applications**

This WG will co-ordinate work in two directions.

1. The development of software tools including programming language environments, as a linkage between behavioural type theories and development practice for large-scale distributed systems.
2. The organisation of case studies to apply theories, languages, methodologies and tools developed throughout the Action, to the practice of large-scale software development, centring on the above mentioned software tools.

The aim is to develop tools with which developers can build, analyse and maintain systems incorporating behavioural types. They include analysis tools that are applicable to already running applications built with standard languages and development methods e.g by automatically deriving and annotating code with type information in communication APIs as described in WG3 above. They also include tools for extensions of existing languages and new domain specific languages developed in WG3. Concrete examples of such software tools include:

- A concrete embodiment of developed programming principles as concrete programming environments. This includes, if it is a new programming language, production-level language processors (parsers) and runtime; or, if it is for existing programming languages, as APIs and runtime libraries. These

environments include bindings to major communication transports such as TCP, UDP, HTTP, AMQP and OMQ, as well as messaging formats such as XML, JSON and Google protocol buffer.

- A tool chain in which a developer can specify contracts for inter-component communications, project such a contract onto each endpoint through projection algorithms, and, through multiple language bindings, use these endpoint constraints to validate source code written in a concrete programming language, to ensure that, when these components are executed and interact with each other, they conform to the original contract.
- A complementary tool chain for regulating behaviours of distributed applications, which enables specifications of behavioural constraints at various levels, from bare protocols to more complex policies. This will generate data for runtime monitoring of the interaction behaviour of applications, in order to check compliance with the specified contracts.

The central point of having these software tools is to make it possible for developers to effectively benefit from the safety and liveness guarantees for the constructed software systems provided by the theories of behavioral types without having to deal with the complexity of understanding the theories themselves.

Regarding point (2) above, we will first identify the main strengths and application areas of the tools developed in research projects lead by members of the Action. This effort will produce a suite of case-studies emphasising the benefits of the various tools, at the same time contrasting and comparing them. Afterwards, we will organise use-case explorations, development/industrial experiments and feedback from industry, through the contacts that have been cultivated by participants in the Action during their own nationally-funded research projects during the last several years. Involving these development partners in the Action, via the activities of this WG, will maximise the benefit for them and for the Action, by putting them into contact with a broader group of researchers that will provide valuable feedback on each existing application and/or tool. The programming languages developed or adapted by WG3, and the tools and software development environments developed by this WG, will be channelled towards the industrial partners for trial in suitable case studies and large-scale explorations. Conversely, the requirements of application domains such as healthcare, financial services, e-science and development of large-scale distributed software infrastructure, will be fed into the activities of all WGs in order to maintain a focus on the needs of end-users.

## **E. ORGANISATION**

### **E.1 Coordination and organisation**

The Action will be co-ordinated by the Management Committee (MC), which as usual will have a Chair, a Vice-Chair and representatives from the participating countries. The scientific work of the Action will be overseen by four Working Groups (WGs), as described in Sections D.1 and D.2. Each WG will have a co-ordinator who will participate in the MC meetings in order to report on progress. There will also be a website manager and a co-ordinator of Short-Term Scientific Missions (STSMs).

#### **Management Committee**

The main responsibilities of the MC are:

- election of the Chair and Vice-Chair, and appointment of the WG co-ordinators, the website manager

- and the STSM co-ordinator
- ensuring that the planned activities of the Action (MC meetings, WG meetings, STSMs, workshops, training schools) take place in such a way as to meet the objectives of the Action
  - producing the annual reports and final report of the Action
  - ensuring that the activities of the WGs are co-ordinated
  - ensuring that the scientific results of the Action are disseminated appropriately, including via the website

The MC will have an annual meeting, which will be part of the annual workshop of the Action (see below). The website of the Action will have a discussion forum for the MC, which will be used for continuous discussion of the progress, direction and activities of the Action. At half-year points, between the annual meetings, the MC will have a tele-conference for more formal discussion of any necessary business.

### **Working Group Meetings**

Each WG will use a continuous process of discussion, led by its co-ordinator, in order to achieve its objectives. This discussion will take place in a forum on the Action's website, by email, and by Skype audio/video calls. Also, at the annual workshop of the Action, there will be discussion sessions related to the theme of each WG, and a more formal meeting of the members of each WG.

### **Workshops**

There will be an annual workshop of the Action, at which all participants will meet to exchange ideas, report on progress and plan further collaboration. The workshops will also be open to participation by any interested people.

### **Training Schools**

The Action will organise two training schools, in Years 2 and 4. These will be aimed mainly at PhD students and junior researchers. Courses will be contributed largely, but not exclusively, by participants in the Action.

### **Short-Term Scientific Missions**

STSMs are an important part of the activity of the Action. They will consist of short visits, typically of one or two weeks but exceptionally of longer duration, between participating sites. STSMs are primarily intended for PhD students and early-career researchers to visit other research groups in order to acquire new expertise or to contribute their expertise to particular projects. They will facilitate the exchange of ideas between participating institutions and will help PhD students and early-career researchers to develop contacts and collaborators which will be of benefit to their later careers.

### **Website**

The website of the Action will be set up at an early stage, and will be used extensively for discussion in a collection of forum areas for the MC and each WG. The website will also be essential for dissemination of the results of the Action, as described in Section H.

### **Milestones**

The reports produced by the WGs and the MC, and the published material from the annual workshops and the training schools, constitute a series of milestones for the Action. Specifically:

- M1 (end of Year 1): State-of-the-art reports by each WG. First Annual Report by each WG. First Annual Report of the Action. Published proceedings of the first workshop.
- M2 (end of Year 2): Second Annual Report by each WG. Second Annual Report of the Action. Published proceedings of the second workshop. Published tutorial material from the first Training School.
- M3 (end of Year 3): Third Annual Report by each WG. Third Annual Report of the Action. Published proceedings of the third workshop.
- M4 (end of Year 4): Final report by each WG. Final report of the Action. Published proceedings of the fourth workshop. Published tutorial material from the second Training School.

These milestones will enable the MC to monitor the progress of the Action towards the achievement of its objectives, and to plan changes in the work programme if necessary.

## **E.2 Working Groups**

The Action will have four Working Groups, which are described in Sections D.1 and D.2. A co-ordinator for each WG will be appointed by the MC. Membership of the WGs will be determined by the interests of the participants in the Action. It is expected that many participants will belong to more than one WG, and this will help to ensure compatibility between the directions being taken by the WGs.

## **E.3 Liaison and interaction with other research programmes**

The Action will liaise with the EU FP7 projects mentioned in Section B.4. In the case of ASCENS and HATS, which have participants in common with the Action, it will be straightforward to establish contact and propose exploratory joint meetings. In the case of ANIKETOS, an initial approach will be made by the Chair of the Management Committee and the co-ordinators of the Working Groups. Representatives of those projects, as well as others that may be identified during the course of the Action, will be invited to participate in relevant meetings and workshops of the Action. It is expected that the greatest benefit of interaction with the projects mentioned in Section B.4 will come from access to a wider range of practical case studies and scenarios.

The Action will also develop links with COST Action IC0901, also mentioned in Section B.4. The open and extensible nature of COST Actions means that mutual participation in workshops and working groups should be straightforward.

## **E.4 Gender balance and involvement of early-stage researchers**

This COST Action will respect an appropriate gender balance in all its activities and the Management Committee will place this as a standard item on all its MC agendas. The Action will also be committed to considerably involve early-stage researchers. This item will also be placed as a standard item on all MC agendas.

This COST Action will respect an appropriate gender balance in all its activities and the Management Committee will place this as a standard item on all its MC agendas. The Action will also be committed to considerably involve early-stage researchers. This item will also be placed as a standard item on all MC agendas.

It is well known that computer science does not have a good gender balance. This applies to every level from undergraduate courses through PhD students and junior researchers to university faculty members and senior researchers. However, the scientific community proposing this Action has a better gender balance than computer science as a whole. The list of experts contains 12 female members from a total of 64, i.e. just under 20%, which compares favourably with a typical ratio of 15% or less among the faculty members of many university departments. These female members include early-stage researchers, mid-career scientists and senior researchers/professors. There are also a significant number of female PhD students among the research groups participating in the Action. Building on this strong initial position, the Action will ensure that female participants are actively involved in all areas, including membership of the MC and the WGs, participation in Short-Term Scientific Missions, presentation of results at the annual workshop, and participation in the Training Schools as both lecturers and participants.

The list of experts also contains many early-stage researchers, and all of the research groups involved in the Action have PhD students who are likely to become post-doctoral researchers during the Action. This Action will ensure that early-career researchers have the opportunity to participate fully in scientific meetings associated with the Action. In particular, short-term scientific missions will benefit early-career researchers. Training schools organised by the Action will benefit PhD students and early-career researchers as participants. Also, when appropriate, early-career researchers will be invited to present courses at training schools, which will benefit their career development. Similarly, active involvement in the WGs and in preparation of their annual reports will be valuable experience for early-career researchers.

## F. TIMETABLE

The duration of the Action will be 4 years.

The Action will begin with a kick-off meeting at which the MC will be elected. The MC will appoint co-ordinators of the WGs, and the website manager and STSM manager.

There will be an annual workshop for all participants in the Action: years 1, 2, 3, 4.

There will be two training schools, combined with the workshops in years 2 and 4.

There will be an annual meeting of the MC, and annual meetings of the WGs, combined with the annual workshops. The format of a WG/MC meeting will be half a day for the WGs to meet for discussion and for a business meeting, and half a day for the MC meeting (including the co-ordinators of the WGs).

	Time (Years)			
	1	2	3	4
<b>Kick-off Meeting</b>	*			
<b>MC &amp; WG Meeting</b>	*	*	*	*
<b>Workshop</b>		*	*	*
<b>Training School</b>			*	*

## G. ECONOMIC DIMENSION

The following 12 COST countries have actively participated in the preparation of the Action or otherwise indicated their interest: Denmark, France, Germany, Ireland, Italy, Netherlands, Poland, Portugal, Romania, Serbia, Sweden, United Kingdom. On the basis of national estimates, the economic dimension of the activities to be carried out under the Action has been estimated at 48 Million € for the total duration of the Action. This estimate is valid under the assumption that all the countries mentioned above but no other countries will participate in the Action. Any departure from this will change the total cost accordingly.

## H. DISSEMINATION PLAN

### H.1 Who?

The results of this COST Action will be disseminated to the following target audiences:

- Participants in the COST Action.
- Researchers in related fields of computer science.
- Software developers within the area of large-scale distributed systems.
- Computing researchers and practitioners more generally.
- The general public with an interest in technology.

### H.2 What?

*For participants in the Action*, the following dissemination methods will be used.

- An internal, protected website with description of completed, current and future activities. This includes working documents, publications, reports from workshops and information about grant applications involving members of the Action
- An internet-based communication network for interaction between participants, with special focus on the needs of PhD students and early-career researchers.
- Short-Term Scientific Missions by PhD students and early-career researchers within the Action.
- Two training schools for PhD students and early-career researchers.
- An annual workshop which serves as a meeting for all Action participants and as a peer-reviewed symposium.
- Publications in international peer reviewed journals and conference proceedings as a result of the scientific projects and annual workshops of the Action.
- Annual reports by the Working Groups and by the Management Committee.
- A final report that describes the outcomes and successes of the Action.

*For researchers in related fields of computer science*, the following dissemination methods will be used.

- Publication of state-of-the-art reports, annual reports, case study reports, workshop proceedings, software documentation, and the final report.
- A public website with general information about the Action as well as a repository of publications by the Action.
- The annual workshops, and the two training schools, organised by the Action.
- Contributions to other national and international conferences and symposia in the form of tutorials and satellite events.
- Articles in peer-reviewed scientific and technical journals and conferences.

*For software developers within the area of large-scale distributed systems*, the following dissemination methods will be used.

- Invitation to participate in trials of novel programming languages, tools and methodologies developed by the Action.
- Release of open-source software tools for system analysis, based on the research carried out by the Action.
- Encouraging PhD students within the Action to apply for internships with companies who may benefit from software produced by the Action.
- A public website with general information about the Action as well as a repository of publications by the Action.

*For computing researchers and practitioners in general*, the following dissemination methods will be used.

- Posting of general information on the public website.
- Less technical articles in general computer science publications such as Communications of the ACM, IEEE Spectrum, and magazines of national computing societies.
- Articles in publications by national funding agencies.

*For the general public*, the following dissemination methods will be used.

- Articles in general science and technology publications.
- Participation in events aimed at public understanding of science.
- Posting of information and non-technical articles on the public website.

### **H.3 How?**

The dissemination methods used within the Action will be used to present the research insights obtained. Research publications are a traditional form of disseminating knowledge both within a specific research community and to a wider research setting. Apart from this form of dissemination,

- The *website* will have a public part that will be used to present published papers and to support stable citation and the correct attribution of work done. This public part will disseminate knowledge to the wider research community. A password-protected part of the website will allow for communication of ongoing research within the Action in the form of online discussion fora as well as draft manuscripts.

- The *annual workshop* will act as a forum in which members of the Action as well as other researchers can present their work.
- The *exchange visits* by PhD students will, when seen as a means of disseminating knowledge, serve to increase knowledge of work in progress.
- The *final report* will disseminate the results of the Action to the community as a whole.

DRAFT MOU



## Part II - Additional Information (This part will not be element of the MoU)

### Part II-A . LIST OF EXPERTS

Total number of participants 64

Gender balance: female 12 of 64 (18.75%)

#### COST Participants

##### DE - Germany

Prof. Jakob REHOF  
Fraunhofer Institute for Software and Systems  
Engineering  
Director  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Prof. Peter THIEMANN  
Universität Freiburg  
Institut für Informatik  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

##### DK - Denmark

Mr Morten MARQUARD  
Exformatics A/S  
Chief Technology Officer  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Marco CARBONE  
IT University of Copenhagen  
Programming, Logic and Semantics  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Dr Thomas HILDEBRANDT  
IT University of Copenhagen  
Programming, Logic and Semantics  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Dr Hans HUTTEL  
University of Aalborg  
Department of Computer Science  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

##### FR - France

Dr Giuseppe CASTAGNA  
CNRS & Université Paris 7  
PPS Laboratory  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT


Dr Marco GIUNTI  
École Polytechnique  
Laboratoire d'informatique (LIX)  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Daniel HIRSCHKOFF  
ENS de Lyon  
LIP Laboratory  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Nils GESBERT  
Grenoble INP – Ensimag  
laboratoire d'informatique de Grenoble  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Damien POUS  
INRIA Rhône-Alpes

<p>SARDES [Proposal Participant] [WG Member] Expertise: ICT</p>	<p>Dr James LEIFER INRIA Rocquencourt INRIA-MSR [Proposal Participant] [WG Member] Expertise: ICT</p>
<p>Dr Ilaria CASTELLANI INRIA Sophia Antipolis Méditerranée INDES [Proposal Participant] [Potential MC Member] [WG Member] Expertise: ICT</p>	<p>Dr Pierre-Loïc GAROCHE Onera, Centre Midi-Pyrénées TIS/DTIM/ISC [Proposal Participant] [WG Member] Expertise: ICT</p>
<p>Dr Samir CHOUALI Université de Franche-Comté LIFC [Proposal Participant] [WG Member] Expertise: ICT</p>	<p>Dr Sebti MOUELHI Université de Franche-Comté LIFC [Proposal Participant] [WG Member] Expertise: ICT</p>

 IE - Ireland

Prof. Matthew HENNESSY  
Trinity College Dublin  
School of Computer Science and Statistics  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

 IT - Italy

Mr Cesare ACCOMAZZO  
Etica s.r.l.  
President  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Maria Grazia BUSCEMI  
IMT Lucca Institute for Advanced Studies  
Computer Science and Applications  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Gian Luca CATTANI  
Maps S.p.A.  
Software Engineer  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Mario BRAVETTI  
Università di Bologna  
Dipartimento di Scienze dell'Informazione  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Dr Elena GIACHINO  
Università di Bologna  
Dipartimento di Scienze dell'Informazione  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Ivan LANESE  
Università di Bologna  
Dipartimento di Scienze dell'Informazione  
[Proposal Participant] [WG Member]  
Expertise: ICT

Prof. Cosimo LANEVE  
Università di Bologna  
Dipartimento di Scienze dell'Informazione  
[Proposal Participant] [WG Member]  
Expertise: ICT

Prof. Davide SANGIORGI  
Università di Bologna  
Dipartimento di Scienze dell'Informazione  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT


<p>Dr Gianluigi ZAVATTARO  Università di Bologna  Dipartimento di Scienze dell'Informazione  [Proposal Participant] [WG Member]  Expertise: ICT</p>	<p>Dr Franco BARBANERA  Università di Catania  Dipartimento di Matematica e Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>
<p>Dr Lucia ACCIAI  Università di Firenze  Dipartimento di Sistemi e Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>	<p>Dr Michele BOREALE  Università di Firenze  Dipartimento di Sistemi e Informatica  [Proposal Participant] [Potential MC Member]  [WG Member]  Expertise: ICT</p>
<p>Dr Roberto BRUNI  Università di Pisa  Dipartimento di Informatica  [Proposal Participant] [Potential MC Member]  [WG Member]  Expertise: ICT</p>	<p>Dr Marco ALDINUCCI  Università di Torino  Dipartimento di Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>
<p>Dr Viviana BONO  Università di Torino  Dipartimento di Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>	<p>Dr Sara CAPECCHI  Università di Torino  Dipartimento di Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>
<p>Prof. Mario COPPO  Università di Torino  Dipartimento di Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>	<p>Dr Ugo DE'LIGUORO  Università di Torino  Dipartimento di Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>
<p>Prof. Mariangiola DEZANI-CIANCAGLINI  Università di Torino  Dipartimento di Informatica  [Proposal Participant] [Potential MC Member]  [WG Member]  Expertise: ICT</p>	<p>Dr Luca PADOVANI  Università di Torino  Dipartimento di Informatica  [Proposal Participant] [WG Member]  Expertise: ICT</p>

 NL - Netherlands

Dr Peter WONG  
Fredhopper  
Research Software Engineer  
[Proposal Participant] [WG Member]  
Expertise: ICT

 PL - Poland

Dr Pawel T WOJCIECHOWSKI  
Poznan University of Technology  
Institute of Computing Science  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

 PT - Portugal

<p>Dr Luís CAIRES Universidade Nova de Lisboa Departamento de Informática [Proposal Participant] [WG Member] Expertise: ICT</p>	<p>Dr Carla FERREIRA Universidade Nova de Lisboa Departamento de Informática [Proposal Participant] [WG Member] Expertise: ICT</p>
<p>Dr Jorge PEREZ Universidade Nova de Lisboa Departamento de Informática [Proposal Participant] [WG Member] Expertise: ICT</p>	<p>Dr António RAVARA Universidade Nova de Lisboa Departamento de Informática [Proposal Participant] [Potential MC Member] [WG Member] Expertise: ICT</p>
<p>Dr Hugo VIEIRA Universidade Nova de Lisboa Departamento de Informática [Proposal Participant] [WG Member] Expertise: ICT</p>	<p>Dr Dimitris MOSTROUS University of Lisbon Departamento de Informática [Proposal Participant] [WG Member] Expertise: ICT</p>
<p>Prof. Vasco VASCONCELOS University of Lisbon Departamento de Informática [Proposal Participant] [Potential MC Member] [WG Member] Expertise: ICT</p>	

 RO - Romania

<p>Prof. Gabriel CIOBANU “Alexandru Ioan Cuza” University of Iasi &amp; Institute of Computer Science Faculty of Computer Science [Proposal Participant] [Potential MC Member] [WG Member] Expertise: ICT</p>	
---	--

 RS - Serbia


<p>Prof. Silvia GHILEZAN University of Novi Sad Centre for Mathematics and Statistics [Proposal Participant] [Potential MC Member] [WG Member] Expertise: ICT</p>	<p>Prof. Jovanka PANTOVIC University of Novi Sad Centre for Mathematics and Statistics [Proposal Participant] [WG Member] Expertise: ICT</p>
---	--

 SE - Sweden

<p>Dr Johannes BORGSTROM University of Uppsala Department of Information Technology [Proposal Participant] [WG Member] Expertise: ICT</p>	<p>Prof. Joachim PARROW University of Uppsala Department of Information Technology [Proposal Participant] [Potential MC Member] [WG Member]</p>
---	---

Expertise: ICT

Dr Björn VICTOR  
University of Uppsala  
Department of Information Technology  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

 UK - United Kingdom

Prof. Steve ROSS-TALBOT  
Cognizant Technology Solutions Ltd.  
European Technology Officer  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Pierre-Malo DENIELOU  
Imperial College London  
Department of Computing  
[Proposal Participant] [WG Member]  
Expertise: ICT

Prof. Sophia DROSSOPOULOU  
Imperial College London  
Department of Computing  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Dr Raymond HU  
Imperial College London  
Department of Computing  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Nobuko YOSHIDA  
Imperial College London  
Department of Computing  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Dr Romain DEMANGEON  
Queen Mary University of London  
School of Electronic Engineering and Computer  
Science  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Luca FOSSATI  
Queen Mary University of London  
School of Electronic Engineering and Computer  
Science  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Kohei HONDA  
Queen Mary University of London  
School of Electronic Engineering and Computer  
Science  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Dr Aybek MUKHAMEDOV  
Queen Mary University of London  
School of Electronic Engineering and Computer  
Science  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Gary BROWN  
Red Hat Inc.  
Senior Software Engineer  
[Proposal Participant] [WG Member]  
Expertise: ICT

Dr Simon GAY  
University of Glasgow  
School of Computing Science  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

Dr Emilio TUOSTO  
University of Leicester  
Department of Computer Science  
[Proposal Participant] [Potential MC Member]  
[WG Member]  
Expertise: ICT

## Non-COST Participants

None

## European Commission Participants

None

## European Bodies Participants

None

### **Part II-B. HISTORY OF THE PROPOSAL**

For several years, there had been increasing activity on the topic of behavioural types in distributed and service-oriented systems. There were several research groups as well as a number of more isolated researchers. Some links had developed between particular groups, but there was not a strong sense of an overall community. A workshop was held in April 2011, which was very effective in bringing researchers together and in identifying new researchers in the field as invitations spread via individual contacts. At the end of the workshop there was a general desire to consolidate the community and develop a more coordinated approach towards further research on this topic. This led to two initiatives: (1) a proposal for a special issue of a journal, to focus on behavioural type systems; (2) the decision to apply for a COST Action.

The special journal issue is going ahead and at the time of writing, submissions are being reviewed. The call for papers was distributed widely and resulted in submissions from authors who had not previously been in contact with any participants of the workshop. Several of these authors are now potential participants in the proposed Action. The preliminary proposal was developed during the summer, by a collaborative process in which a small writing group produced drafts which were commented on by many members of the community. The development of the full proposal has followed a similar process but with more people involved in writing; again, drafts were reviewed by other potential participants in the Action.

The objectives of the Action were decided by a collaborative discussion, from which emerged a shared understanding of the most important challenges facing the field of distributed service-oriented systems and the most effective ways in which behavioural types can be applied.

### **Part II-C. PRELIMINARY WORK PROGRAMME**

#### **Relationships between the Working Groups**

The sequence of Working Groups: WG1 (Foundations), WG3 (Languages) and WG4 (Tools & Applications) forms a chain from theory to practice. As such, there will be significant interaction between WG1 and WG3, and between WG3 and WG4.

WG2 (Security) addresses a cross-cutting concern, and as a result it will have significant interaction with the other three Working Groups.

What this amounts to is strong interaction between every pair of Working Groups, with the exception that WG1 (Foundations) will have less direct interaction with WG4 (Tools & Applications). Most of the interaction between WG1 and WG4 will be mediated through WG2 (Security) and WG3 (Languages).

# Work Programme and Schedule

## WG1: Foundations

### Year 1

During the first year of the Action, WG1 will prepare a state-of-the-art report on behavioural type theory and its existing use as a foundation for programming languages. The report will provide a comprehensive picture of the research community and its achievements to date, and identify areas in which different approaches to behavioural type theory are conflicting, overlapping or deficient. This knowledge will be used both to plan further foundational research and as a basis for responding to the needs of WG2 (Security) and WG3 (Languages), during subsequent years.

### Year 2

The second year will be devoted to expressivity issues, by promoting a workshop on the themes of safety, liveness and fairness properties of concurrent and mobile systems, as well as information flow analysis, and by discussing ideas about how such properties can be expressed by means of behavioural types. This involves the question of how experience and knowledge gathered by means of different models than process algebra and type theory, like true concurrency models and model checking, can be imported in terms of behavioural types, and to what extent known problems from the world of distributed computing can be faced in the setting of type systems.

At the end of the second year of the Action, the language designs developed in WG3, possibly applied to important scenarios/case studies arising from WG4 (Tools), will be used as test beds to check the effectiveness and suitability of the formal techniques developed in WG1. This will generate topics for WG1's activity during the second half of the Action, e.g. modifications/extensions to the models considered or to the verification techniques.

### Year 3

The third year will be primarily dedicated to the problems of probabilistic and quantitative properties of concurrent and distributed systems, and of their formalizations within the theory of behavioural types. The work will start with the preparation of a report on the most promising attempts at embodying such aspects into process algebras and behavioural type systems, and will be concluded by a meeting on the same themes. Quantitative and probabilistic analysis are relevant both to WG2 (Security) and to WG4 (Tools), giving the opportunity of testing the results from WG1 against case studies and concrete applications.

### Year 4

In the final year, WG1 will work towards an overall assessment of the results obtained both theoretically and practically. This will lead to a final report surveying the achievements as well as collecting the most interesting open questions that the scientific community working in the area should address in the future.

## WG2: Security

## **Year 1**

During the first year of the Action, WG2 will prepare a state-of-the-art report on security mechanisms for large-scale distributed and service-oriented systems. The report will consider security from the perspective of each of the other three Working Groups, in consultation with them: (1) what are the practical security issues in large service-oriented systems, and how well are they handled at present? (2) how is security supported by the programming languages presently used for large-scale distributed software development, and is this support considered to be adequate? (3) to what extent does the theoretical basis for security analysis in large-scale distributed and service-oriented systems allow security to be reliably certified? This report will form the basis for detailed planning of the activities of WG2 and its interaction with the other Working Groups during subsequent years.

## **Year 2**

In its study of new type theories for security properties, WG2 will build on the foundational work in WG1 and keep interacting with WG1 throughout the project. Security properties such as authenticity and secrecy can be understood as secure information-flow properties based on the notion of non-interference. There is already work on information flow in the related area of control flow analysis and also type-based analyses that use correspondence assertions. Also, there is a considerable body of work on the relationship between secure information flow and trust management. In the second year of the Action and onwards, WG2 will therefore examine how behavioural types can be used to give more precise analyses of the secure information flow properties that have been identified in Year 1. WG2 will also address security issues that are not easily understood as information flow properties, and investigate how these they can be captured by type theories. Finally, WG2 will relate its work on type theories to that pursued in WG1 and to existing work on security analyses that use a type-like approach, including control flow analyses.

The practical usefulness of behavioural types for security will depend on their algorithmic properties. In Year 2 and onwards, WG2 will therefore interact with WG4 to determine how security properties can be checked algorithmically. In particular, type checking must be decidable within reasonable complexity bounds and type reconstruction should be possible.

## **Year 3**

Security properties, and in particular those based on non-interference, are important in actual program development. In Year 3, WG2 will therefore examine how behavioural types can be used to guarantee security properties in real programming languages, both when designing new languages, which should be security-minded from their very inception, and when extending existing ones. This work will benefit and interact with the work on type theories for session calculi based on the pi-calculus or the psi-calculus. More specifically, WG2 will examine how behavioural types can be used to guarantee security properties in real programming languages, both when designing new languages, which should be security-minded from their very inception, and when extending existing ones. As such, the work of WG2 will be strongly connected with that on programming language design developed by WG3, which takes correctness concerns into consideration.

Starting in Year 3, WG2 will also work in close connection with WG4 to discover the security requirements of specific application domains, and match these requirements with suitable formalisms for security. An important activity will be to investigate concrete cases (including ones studied by WG4 in Year 2) and determine how behavioural types can be used to reason about security properties for them.

## **Year 4**



The final year of the Action will be devoted to developing software tools that make use of the algorithmic properties of behavioural types. A litmus test for the usefulness of behavioural types for security will be the possibility of applying them to the domains and case-studies studied in Year 3.

## **WG3: Languages**

### **Year 1**

During the first year of the Action, WG3 will prepare a state-of-the-art report on behavioural types in programming languages. This will be approached from two directions: (1) a survey of current experiments and prototypes in adding behavioural types to existing languages or designing new languages that incorporate behavioural types; (2) an assessment of current practice in software development for distributed and service-oriented systems, in the absence of support for behavioural types, and identification of the areas in which behavioural types are most likely to prove beneficial. The report will provide a starting point for detailed planning of a strategy to exploit behavioural types in practical programming languages during the remainder of the Action.

### **Year 2**

In year 2, WG3 will explore non-invasive ways of enriching existing programming and domain-specific languages with behavioural type disciplines allowing for a greater number of correctness properties to be automatically verified. *Non-invasive* means that behavioural types will not, at this stage, be deeply integrated into programming languages; rather, they will be handled at a pre-processing or external tool level. This work will include the study of program annotations (in the form of comments or pragma directives), the definition of program macros to be expanded during pre-processing, as well as the definition of APIs specialized to communication, synchronization, and co-ordination making use of the augmented expressiveness given by behavioural type disciplines.

### **Year 3**

During the third year, WG3 will elaborate on the experience of the previous year and on the feedback received from WG4 regarding the application of behavioural type disciplines to the case studies. In doing so, it will identify the expressiveness gap between the constructs available in current programming languages and the needs that have emerged in the scenarios taken under consideration, and will devise high-level language constructs allowing, on the one hand, for a higher level of abstraction in designing and developing distributed and/or cooperative applications and, on the other hand, for more precise type checking.

### **Year 4**

The final year of the Action will be devoted to studying conservative extensions of both general-purpose and domain-specific languages, with a particular focus on the implementation of (possibly pluggable) components for compilers and run-time environments. This work will also require the study of adapters and interfaces for inter-working between legacy code and programs using these new constructs. The increased expressiveness and precision provided by behavioural types will be exploited for the automatic or assisted generation of adapters and interfaces. The final report of WG3 will summarise and reflect on progress, and indicate directions for future work by continuing projects in this area.

## **WG4: Tools & Applications**

## **Year 1**

During the first year of the Action, WG4 will prepare a report on potential case studies that can be undertaken in collaboration with the industrial participants. This will begin with an assessment of real software systems and the problems that companies face in developing them, and then identify a range of systems or sub-systems that exemplify particular issues. If possible this will include case studies related to systems with a high social impact, such as integrated European health services or voter management systems for electronic elections. The report will be used to plan application-focused work during the remainder of the Action.

## **Year 2**

During the second year, WG4 will select a few case studies based on systems identified during the first year, and try to use the available languages (identified in the state-of-the-art report by WG3) and tools to re-implement or re-structure them, or to implement exemplar systems that incorporate the key features, including security aspects, of the case studies. The second year report will use this experience as the basis for a critique of the available programming languages and tools for behavioural types, and will compile a wish-list for WG2 (Security) and WG3 (Languages) to work towards.

## **Year 3**

By the beginning of year 3, WG3 will have produced designs for programming language extensions based on the idea of using annotations or pragma directives to express behavioural typing properties. WG4 will spend the third year developing compiler extensions and analysis tools based on these extensions. Towards the end of the year it will be possible to test these compilers and tools on the same case studies used in the previous year's work, and assess their effectiveness. The report at the end of this year will reach conclusions about the benefits and limitations of this non-invasive approach.

## **Year 4**

The final year will be devoted to the implementation of programming languages in which behavioural type systems are an intrinsic aspect of the design. This stage will be closely connected to the results from WG3 during the third year, and also to the ongoing work of WG3 during the final year. As far as possible, prototype programming languages will be tested on the same case studies as before, and there will be considerable feedback from the experience of the case studies into the design and implementation of libraries and APIs. The final report, when combined with that of WG3, will present a comprehensive view of the experience of the design, implementation and test cycle for novel programming languages incorporating behavioural type systems.

In this last year, WG4 will organise a workshop to present to industrial partners the solutions to the case studies developed previously. This workshop will show how the tools and the language-based approaches selected in the other WGs can address the real problems and software development challenges the industry faces. Furthermore, WG4 will prepare joint pilot projects with the industrial partners to be developed in their companies and project applications to funding agencies to pursue the work of the Action.

## Part II-D. RECENT PUBLICATIONS

- [AB09] L. Acciai and M. Boreale: Deciding safety properties in infinite-state pi-calculus via behavioural types. Proc. ICALP'09. LNCS 5556:31-42, 2009.
- [ABZ10] L. Acciai, M. Boreale and G. Zavattaro: Behavioural contracts with request-response operations. Proc. Coordination'10. LNCS 6116:16-30, 2010.
- [BBDL11] M. Boreale, R. Bruni, R. De Nicola and M. Loreti: CaSPiS: A Calculus of Sessions, Pipelines and Services. To appear in MSCS, 2012.
- [BCCDC11] V. Bono, S. Capecchi, I. Castellani, and M. Dezani-Ciancaglini: A Reputation System for Multirole Sessions. Proc. TGC'11, LNCS, 2012. To appear.
- [BCDCM11] M. G. Buscemi, M. Coppo, M. Dezani-Ciancaglini, and U. Montanari: Constraints for Service Contracts. Proc. TGC'11, LNCS, 2012. To appear.
- [BCDDDY08] L. Bettini, M. Coppo, L. D'Antoni, M. De Luca, M. Dezani-Ciancaglini, and N. Yoshida: Global Progress in Dynamically Interleaved Multiparty Sessions. Proc. CONCUR'08, LNCS 5201:418-433, 2008.
- [BCDGV08] L. Bettini, S. Capecchi, M. Dezani-Ciancaglini, E. Giachino, and B. Venneri: Session and Union Types for Object Oriented Programming. In Concurrency, Graphs, and Models. LNCS 5065:659-680, 2008.
- [BCdL10] F. Barbanera, S. Capecchi, and U. de' Liguoro: Typing Asymmetric Client-Server Interaction. Proc. FSEN'09, LNCS 5961:97-112, 2010.
- [BdL10] F. Barbanera and U. de' Liguoro: Two Notions of Sub-behaviour for Session-based Client/Server Systems. Proc. PPDP'10, ACM SIGPLAN, 155-164, 2010.
- [BH12] G. Bernardi and M. Hennessy: Modelling Session Types using Contracts. Proc. 27th Symposium on Applied Computing, 2012. To appear.
- [BHJRVÅP11] J. Borgström, S. Huang, M. Johansson, P. Raabjerg, B. Victor, J. Åman Pohjola, J. Parrow: Broadcast Psi-calculi with an Application to Wireless Protocols. In G. Barthe, A. Pardo & G. Schneider (ed.) Software Engineering and Formal Methods, pp74-89, Springer 2011.
- [BHTY10] L. Bocchi, K. Honda, E. Tuosto and N. Yoshida: A Theory of Design-by-Contract for Distributed Multiparty Interactions, Proc. CONCUR'10, LNCS 6269:162-176, 2010.
- [BJPV09] J. Bengtson, M. Johansson, J. Parrow and B. Victor. Psi-calculi: Mobile processes, nominal data, and logic. Proc. LICS, 2009.
- [BLZ08] M. Bravetti, I. Lanese, G. Zavattaro: Contract-Driven Implementation of Choreographies. Proc. of TGC 2008: 1-18, 2008.
- [BM09] M. G. Buscemi and H. Melgratti: Abstract Processes in Orchestration Languages. Proc. ESOP 2009. LNCS 5502:301-315, 2009.
- [BMP11] V. Bono, C. Messa, and L. Padovani: Typing Copyless Message Passing. Proc. of ESOP'11, LNCS 6602: 57-76, 2011.
- [BZ08] M. Bravetti, G. Zavattaro: A Foundational Theory of Contracts for Multi-party Service Composition. Fundamenta Informaticae 89(4):451-478, 2008.
- [CBDHY11] T.-C. Chen, L. Bocchi, P.-M. Denielou, K. Honda and N. Yoshida: Asynchronous Distributed Monitoring for Multiparty Session Enforcement. Proc. TGC'11, LNCS, 2012. To appear.
- [CCDCR10] S. Capecchi, I. Castellani, M. Dezani-Ciancaglini, and T. Rezk: Session Types for Access and Information Flow Control. Proc. CONCUR, LNCS 6269:237-252, 2010.
- [CCDDG09] S. Capecchi, M. Coppo, M. Dezani-Ciancaglini, S. Drossopoulou, and E. Giachino: Amalgamating Sessions and Methods in Object Oriented Languages with Generics. Theoret. Comp. Sci. 410:142-167, 2009.
- [CDC09] M. Coppo and M. Dezani-Ciancaglini: Structured Communications with Concurrent Constraints. Proc. TGC'08, LNCS 5474:104-125, 2009.
- [CDCP11] G. Castagna, M. Dezani-Ciancaglini, and L. Padovani: On Global Types and Multi-party Sessions. Proc. FMOODS/FORTE, LNCS 6722:1-28, 2011.
- [CDG08] M. Coppo, M. Dezani-Ciancaglini, and E. Giovannetti. Types for Ambient and Process Mobility. Math. Struct. Comp. Sci. 18:221-290, 2008.

- [CDGP09] G. Castagna, M. Dezani-Ciancaglini, E. Giachino, and L. Padovani: Foundations of Session Types. Proc. PPDP, 2009.
- [CGY10] S. Capecchi, E. Giachino, and N. Yoshida: Global Escape in Multiparty Sessions. Proc. FSTTCS 2010. LIPIcs 8:338-351, 2010.
- [CK11] G. Ciobanu and M. Koutny: Timed Migration and Interaction with Access Permissions. Proc. FM, LNCS 6664:293-307, 2011.
- [DCdL10] M. Dezani-Ciancaglini and U. de' Liguoro: Sessions and Session Types: an Overview. Proc. WS-FM'09, LNCS 6194:1-28, 2010.
- [DCDMY09] M. Dezani-Ciancaglini, S. Drossopoulou, D. Mostrous, and N. Yoshida: Session Types for Object-Oriented Languages. Inf. and Comp. 207(5):595-641, 2009.
- [DCGJP11] M. Dezani-Ciancaglini, S. Ghilezan, S. Jaksic, and J. Pantovic: Types for Role-Based Access Control of Dynamic Web Data. Proc. WFLP'10, LNCS 6559:1-29, 2011.
- [DCGPV08] M. Dezani-Ciancaglini, S. Ghilezan, J. Pantovic, and D. Varacca: Security Types for Dynamic Web Data. Theoret. Comp. Sci. 402:156-171, 2008.
- [DGDY07] M. Dezani-Ciancaglini, E. Giachino, S. Drossopoulou, and N. Yoshida: Bounded session types for object-oriented languages. Proc. FMCO. LNCS 4709:207-245, 2007.
- [DHS10] R. Demangeon, D. Hirschhoff, D. Sangiorgi: Termination in Impure Concurrent Languages. Proc. CONCUR 2010, LNCS 6269:328-342, 2010
- [DY10] P.-M. Denielou and N. Yoshida: Buffered Communication Analysis in Distributed Multiparty Sessions. Proc CONCUR'10, LNCS 6269: 343-357, 2010.
- [DY11] P.-M. Denielou and N. Yoshida: Dynamic multirole session types. Proc. POPL 2011, 435–446, 2011.
- [DY12] P.-M. Denielou and N. Yoshida: Multiparty Session Types Meet Communicating Automata. Proc. ESOP 2012, LNCS, 2012. To appear.
- [GL11] E. Giachino and C. Laneve: Analysis of deadlocks in object groups. Proc. FMOODS/FORTE, LNCS 6722:168-182, 2011.
- [GLMZ09] C. Guidi, I. Lanese, F. Montesi, G. Zavattaro: Dynamic Error Handling in Service Oriented Applications. Fundam. Inform. 95(1):73-102, 2009.
- [GV10] S. J. Gay and V. T. Vasconcelos: Linear Type Theory for Asynchronous Session Types. J. Func. Prog. 20(1):19-50, 2010
- [GV10a] M. Giunti and V.T. Vasconcelos. A linear account of session types in the pi calculus. Proc. CONCUR, LNCS 6269:237-252, 2010.
- [GVRGC10] S. J. Gay, V. T. Vasconcelos, A. Ravara, N. Gesbert and A. Z. Caldeira. Modular Session Types for Distributed Object-Oriented Programming. Proc. POPL, 2010
- [H11] H. Hüttel: Typed Psi-Calculi. Proc. CONCUR, LNCS 6901, 2011.
- [HDKS11] H. Hüttel, M. Dahl, N. Kobayashi and Y. Sun: Type-Based Automated Verification of Authenticity in Asymmetric Cryptographic Protocols. Proc. 9th International Symposium on Automated Technology for Verification and Analysis. LNCS, 2011.
- [HKPYH10] R. Hu, D. Kouzapas, O. Pernet, N. Yoshida and K. Honda: Type-Safe Eventful Sessions in Java, Proc. ECOOP'10: LNCS 6183: 329-353, 2010.
- [HMBCY11] K. Honda, A. Mukhamedov, G. Brown, T.-C. Chen and N. Yoshida: Scribbling Interactions with a Formal Foundation, Proc. ICDCIT'11: LNCS: 55-75, 2011.
- [HYC08] K. Honda, N. Yoshida and M. Carbone, Multiparty Asynchronous Session Types. Proc POPL'08, 273–284, 2008.
- [HYH08] R. Hu, N. Yoshida and K. Honda. Session-Based Distributed Programming in Java. Proc. ECOOP 2008. LNCS 5142: 516–541, 2008.
- [KS10] N. Kobayashi, D. Sangiorgi: A hybrid type system for lock-freedom of mobile processes. ACM Trans. Prog. Lang. Syst. 32(5):1-49, 2010.
- [KYH11] D. Kouzapas, N. Yoshida and K. Honda: On Asynchronous Session Semantics. Proc. FORTE/FMOODS'2011. LNCS 6722, 2011.
- [LGMZ08] I. Lanese, C. Guidi, F. Montesi, G. Zavattaro: Bridging the Gap between Interaction- and Process-Oriented Choreographies. Proc. SEFM 2008:323-332, 2008.

- [LP10] H. A. Lopez and J. A. Perez: Time and Exceptional Behavior in Multiparty Structured Interactions. Proc. WS-FM'11, LNCS, 2012.
- [MV11] D. Mostrous and V. T. Vasconcelos: Session typing for a featherweight Erlang. Proc. Coordination, LNCS 6721:95-109, 2011.
- [MY09] D. Mostrous and N. Yoshida, Session-Based Communication Optimisation for Higher-Order Mobile Processes, Proc. TLCA'09. LNCS 5608:203–218, 2009.
- [NYPHK11] N. Ng, N. Yoshida, O. Pernet, R. Hu and Y. Kryftis: Safe Parallel Programming with Session Java. Proc. COORDINATION, LNCS 6721:110-126, 2011.
- [P10] L. Padovani: Contract-Based Discovery of Web Services Modulo Simple Orchestrators. Theoret. Comp. Sci. 411:3328-3347, 2010.
- [P11] L. Padovani: Fair Subtyping for Multi-Party Session Types. Proc. Coordination, LNCS 6721:127-141, 2011.
- [PCPT12] J. A. Perez, L. Caires, F. Pfenning and B. Toninho: Linear Logical Relations for Session-Based Concurrency. Proc. ESOP, LNCS, 2012. To appear.
- [V09] Vasco T. Vasconcelos: Session Types for Linear Multithreaded Functional Programming. Proc. PPDP, 2009.
- [YDBH10] N. Yoshida, P.-M. Denielou, A. Bejleri and R. Hu, Parameterised Multiparty Session Types. Proc. FOSSACS 2010. LNCS 6014: 128-145, 2010.

## **Part II-E. FURTHER REMARKS**

### **Early-Stage Researchers**

As explained in Section I.E, the list of participants of this Action includes many early-stage researchers, and they will have the opportunity to engage fully with all aspects of the Action. This includes: membership of Working Groups; participation in Short-Term Scientific Missions; participation in Training Schools as attendees and, where appropriate, lecturers; presentation of research results at the annual workshops; involvement in preparing the annual reports of the Working Groups.

### **Gender Balance**

As explained in Section I.E, the gender balance among the participants in this Action is better than that typically found in a university computer science department. The Management Committee and Working Groups will constantly strive to ensure a good gender balance in all activities of the Action.

### **Range of Countries**

The list of participants includes experts from 12 countries, including several of the newer member states of the EU. This means that already a significant proportion of the COST countries are included in the Action. The Management Committee and the Working Group co-ordinators will actively seek to identify interested researchers in other COST (and, if appropriate, non-COST) countries and encourage them to join the Action.

### **Monitoring and Evaluation of Objectives**

Section I.C includes a detailed list of objectives, including, where appropriate, quantitative targets for the level of activity. Progress towards the objectives, and measurement of outputs and activities against targets, will be monitored continuously by the Working Group co-ordinators and formally checked by the Management Committee at its annual meetings.

Section I.E includes a series of milestones, explicitly stating what the Management Committee expects to receive each year from the Working Groups and as a result of the workshops and Training Schools.

## **Application and Exploitation of Results**

The interaction between the Working Groups, set out in Section II.C, is designed to ensure a flow of results towards their areas of application. For example, theoretical research carried out within WG1 (Foundations) will be planned to address points identified by WG2 (Security) and WG3 (Languages). Similarly, programming language design work carried out within WG3 (Languages) will be aimed at providing desirable language features and constructs identified by WG2 (Security) and WG4 (Tools & Applications).

The ultimate application of the results of this Action will be to industrial software development methods and processes. This application will be ensured by the involvement in the Action of several representatives of software companies (see below). During the Action, the Management Committee and the Working Group co-ordinators will actively look for additional software companies who might be interested in participating.

## **Industrial Involvement**

The following companies have agreed to participate in the Action. Their representatives are included in the list of experts.

Cognizant Technology Solutions Ltd. (UK)

Exformatics A/S (Denmark)

Etica s.r.l. (Italy)

Fredhopper (Netherlands)

Maps S.p.A. (Italy)

Red Hat, Inc. (UK)