European Cooperation
in Science and Technology
- COST -

Secretariat

Brussels, 4 July 2012

COST 4154/12

MEMORANDUM OF UNDERSTANDING

Subject : Memorandum of Understanding for the implementation of a European Concerted Research Action designated as COST Action IC1205: Computational Social Choice

Delegations will find attached the Memorandum of Understanding for COST Action as approved by the COST Committee of Senior Officials (CSO) at its 185th meeting on 6 June 2012.

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

COST Action IC1205
COMPUTATIONAL SOCIAL CHOICE

The Parties 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 4154/11 "Rules and Procedures for Implementing COST Actions", or in any new document amending or replacing it, the contents of which the Parties are fully aware of.

2. The main objective of the Action is to coordinate European research in Computational Social Choice and to address the fundamental challenges in the design of methods for collective decision making raised by recent advances in information technology.

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 EUR 76 Million in 2012 prices.

4. The Memorandum of Understanding will take effect on being accepted by at least five Parties.

5. The Memorandum of Understanding will remain in force for a period of 4 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.

A. ABSTRACT AND KEYWORDS

The COST Action on Computational Social Choice will address the fundamental new challenges in the design and analysis of methods for collective decision making raised by recent technological advances in areas such as social networks, electronic commerce, webpage ranking, and e-governance. Computational Social Choice is a novel and innovative research trend that is now gathering substantial momentum, especially in the European research arena. It combines methods from Computer Science with insights from Economic Theory. The Action will provide a much needed framework for coordinating research efforts in this important new field at the European level. The scientific programme will focus on four thematic areas: (1) voting and elections, (2) fair division, (3) information merging, and (4) matching mechanisms. In organisational terms, the Action will particularly emphasise the close involvement of Early-Stage Researchers.
Keywords: Computational Social Choice; Economics and Computer Science; Artificial Intelligence; Multiagent Systems; Collective Decision Making.

B. BACKGROUND

B.1 General background

The COST Action on Computational Social Choice (COMSOC) will address, at the European level, the fundamental new challenges in the design and analysis of methods for collective decision making raised by recent technological advances in areas such as social networks, electronic commerce, webpage ranking, and e-governance. Computational Social Choice is a novel area of research that recently formed at the interface of, on the one hand, a research tradition in Economic Theory known as Social Choice Theory and, on the other, a broad range of concepts and techniques originating in Computer Science.

Social Choice Theory addresses a fundamental question: How should we aggregate the preferences of a group of individuals so as to arrive at an adequate collective preference that can serve as the basis for making acceptable group decisions? Not only is this question relevant to the quality of everyday interactions between people, but it strikes at the very heart of our understanding of democracy. Indeed, both the importance of the question and the significance of the research tradition that has sought to address it are widely recognised, as witnessed, for instance, by the Nobel prizes awarded to social choice theorists such as Kenneth Arrow and Amartya Sen.

Arrow's doctoral thesis, published in 1951 under the title "Social Choice and Individual Values", is often regarded as the starting point of modern Social Choice Theory. By providing a rigorous mathematical formulation of the problem of fairly aggregating the preferences of a group of individuals into a collective preference order, Arrow made this problem amenable to systematic scientific analysis. Sixty years on, in the light of extraordinary technological advances, we now have to recast the fundamental question of social choice. It is not anymore just a matter for groups of people who have to choose between a small number of alternatives. Instead, a wide range of modern applications, many of them in the domain of Information and Communication Technologies (ICT), require agents, both human and artificial, to make collective decisions, often choosing from a huge number of alternatives and under conditions where some of the classical assumptions made in Social Choice Theory. For instance, both the nature of preferences and the concept of fairness may have to be reconsideered in this novel context. Examples for such applications include the integration of search results produced by multiple Internet search engines, the assignment of junior doctors to hospitals in view of both their preferences and their qualifications, the amalgamation of user preferences to issue recommendations for other users in online recommender systems, the support of group decision making technologies for social networks, and a multitude of problems in e-commerce and e-governance. Computational Social Choice provides the scientific foundations for tackling these challenges.

The development of Computational Social Choice as a research area is, to a large extent, a European endeavour: while there are a small number of strong groups overseas, most research takes place in Europe. The first three editions of the International Workshop on Computational Social Choice (http://www.illc.uva.nl/COMSOC/), which has attracted steadily increasing numbers of participants, also took place in Europe (Amsterdam 2006, Liverpool 2008, Duesseldorf 2010) and the fourth will be held in Krakow in September 2012. There now are sufficient critical mass and momentum to move from the bottom-up community building activities of recent years to a coordinated effort at a European scale. A COST Action provides the perfect framework to achieve this transition.

The interdisciplinary nature of the research programme required to address the challenges raised by this new perspective on collective decision making necessitates the close cooperation of experts from different areas. This naturally fits the COST approach: while research on the various subdomains will continue to be funded locally, the COST Action on Computational Social Choice will enable coordination at the European level.

B.2 Current state of knowledge

Computational Social Choice is an interdisciplinary research area at the interface of Social Choice Theory and Computer Science, promoting an exchange of ideas in both directions. On the one hand, it is concerned with the application of techniques developed in Computer Science, such as complexity analysis or algorithm design, to the study of social choice mechanisms, such as voting procedures or fair division algorithms. On the other hand, Computational Social Choice is concerned with importing concepts from Social Choice Theory into computing, for instance to appropriately model the problem of ranking webpages or to support applications in e-governance.

An important line of work in Computational Social Choice applies the tools of computational complexity to social choice settings, particularly to the analysis of voting rules. For instance, while it is computationally easy to compute the winner for most of the voting rules in regular use, there are also rules for which this problem turns out to be computationally intractable (i.e., NP-hard). Computational complexity has also been suggested as a barrier against manipulation in elections.

Classical results from Social Choice Theory show that for any voting rule to choose between three or more candidates that is not dictatorial (meaning that one voter always determines the outcome), there will be situations in which some voters may have an incentive to manipulate the election (in the sense of voting untruthfully). To address this problem, it has been suggested that it may be possible to identify voting rules that make it computationally hard to manipulate. While there are several positive results establishing NP-hardness for certain rules, the most recent discussion has concentrated on the question to what extent such worst-case complexity results offer sufficient protection in practice, and whether or not meaningful average-case complexity results might be achievable as well. Besides manipulation in the aforementioned sense, the complexity of other forms of election control, e.g., by strategically entering additional candidates into the race, has also been studied.

Known methods for collective decision making and classical results from Social Choice Theory may not always be applicable when the number of alternatives from which to choose is very large. This may, for instance, be the case when these alternatives have a combinatorial structure, as for the fair allocation of indivisible goods (where the number of bundles an agent may obtain is exponential in the number of goods) or committee elections (where the number of possible committees is exponential in the number of seats to be filled). For such combinatorial problems, the mere representation of the preferences of individuals over different alternatives becomes a non-trivial problem. Another important direction for research in Computational Social Choice is therefore the application of techniques developed in Artificial Intelligence, particular the subfield known as Knowledge Representation and Reasoning, for the compact representation of preferences to this kind of problem.

Another example for the application of tools typically used in Computer Science to problems stemming from Economics and Social Choice Theory is the use of logic for the formal specification and verification, or more generally analysis, of social procedures. In the same way as computer scientists have long been using logic to formally specify the behaviour of computer systems, so as to allow for the automatic verification of certain desirable properties of such systems, suitable logics may be used to specify social procedures such as voting protocols or fair division algorithms.

While classical Social Choice Theory largely concentrates on the aggregation of individual preferences, a similar methodology can also be applied to the aggregation of other types of information coming from individual agents. In the Artificial Intelligence literature, for instance, the problem of belief merging (building on important earlier work in logical belief revision) has received much attention. In Economic Theory, the problem of judgment aggregation has become an important topic in recent years. In both frameworks, information is modelled as a set of formulas in propositional logic.

Matching problems are a prominent class of social choice problems in which a set of agents has to be matched to another set of agents or objects. Important examples include centralised entry level job markets (e.g., medical students being assigned to hospitals), the assignment of children to public schools, and transplant organ allocation systems. The computational social choice aspect of Matching Theory manifests itself in the (re)design of real-world matching algorithms in order to accommodate efficiency, fairness, and nonmanipulability objectives. Most European countries employ matching algorithms for assignments as mentioned above, although to date cooperation on matching market design at the European level is still limited, and thus opportunities for enhancing our understanding of the spectrum of matching algorithms that can and should be used to accommodate social and economic objectives have not yet been explored in depth.

In recent years there has been a growing interest in applying the methodology of Computational Social Choice to a broad range of new domains. For instance, researchers in agent-based computing have begun to explore voting procedures as a means of achieving coordination between autonomous agents. In the field of recommender systems, researchers working on collaborative filtering have used the axiomatic method of Social Choice Theory to give a formal characterisation of some of their methods, similar to the axiomatic characterisation of preference aggregation mechanisms, as a means of choosing the best method for a concrete problem at hand. In the context of Internet search engines, the same axiomatic method has been used to provide formal foundations for the notion of significance of a webpage, which is exploited by state-of-the art search algorithms, and in the field of bibliometrics it has been used to gain deeper insights into some of the methods used to analyse the productivity and reputation of scientists and research institutes. More generally, this approach can be used for the analysis of online reputation systems.

The state of the art in Computational Social Choice, as outlined here, provides a good platform from which to address the new challenges mentioned above. European researchers have contributed significantly to the development of this state of the art, even if to date there have been no coordinated research programmes at the European level specifically addressing the challenges in this area. The COST Action on Computational Social Choice will build on this knowledge and combine the expertise coming from different corners of the European research community.

B.3 Reasons for the Action

At the scientific level, the Action will provide a platform for consolidating results obtained in nationally funded research activities, concerning important foundational questions, such as the complexity of elections, the design of interaction protocols for fair division, the refinement of matching algorithms, the role of incomplete knowledge and communication requirements in voting, and the extension of core social choice techniques to broader domains, such as aggregating highly structured information from different sources.

At the technological level, the Action has substantial potential in benefiting a range of high-profile application domains, including electronic commerce systems enabling meaningful interaction between human users and automated selling and buying agents, the transition of social networking platforms from communication tools to powerful group decision support frameworks, fair work organisation (timetabling, team formation), and the development of e-governance systems with increased expressivity and flexibility. By their very nature, several of these technological benefits have the potential to also translate into long-term societal benefits.

The Action will also contribute to interdisciplinary research in Europe. It will facilitate synergies by bringing computer scientists into close and regular contact with economists and political scientists.

Finally, there are important benefits for education and training. Currently, only a handful of European universities are in a position to offer full courses on Computational Social Choice. The Action will provide a platform for the exchange and further development of the scarce educational resources currently available.

B.4 Complementarity with other research programmes

While a number of individual research groups have been successful in attracting funding for their own projects at the national level, there is currently no coordinated funding programme for Computational Social Choice (or closely related fields) anywhere in the world, nor has there been any such programme in the past.

Small national projects awarded to individual groups in Europe include grants funded by, amongst others, DFG, ANR, EPSRC, NWO, ISF, TUBITAK.

A small number of European groups have previously (2008-2011) cooperated in a collaborative research project coordinated by the European Science Foundation (funded at the national level) as part of the EUROCORES programme "LogICCC: Modelling Intelligent Interaction". This project made important contributions to advance the state of the art in Computational Social Choice, but the restriction to a small number of participants did not allow for coordination at the European level.

Like this Action, COST Action IC0602 on Algorithmic Decision Theory (2006-2011) brought together expertise from different scientific fields, associated broadly with Computer Science, on the one hand, and the Social and Economic Sciences, on the other. The crucial difference with respect to this Action is that Action IC0602 focussed on decision-theoretic problems faced by a single agent, while social choice concerns phenomena related to collective decision making. Nevertheless, some of the insights gained in Action IC0602 will be relevant for this new Action.

The FET Proactive Initiative on "Fundamentals of Collective Adaptive Systems" (launched in 2012) is complementary to the concerns of the COST Action on Computational Social Choice. It addresses phenomena that may arise in "systems that are constructed as a collective of heterogeneous components and that are tightly entangled with humans and social structures". The Call for Proposals for this new programme does not address the contribution that insights from Economics, and social choice in particular, might be able to make and focusses instead on issues such as evolution and adaptation.

C. OBJECTIVES AND BENEFITS

C.1 Aim

The main objective of the Action is to coordinate European research in Computational Social Choice and to develop the European capacity to address the fundamental challenges in the design and analysis of methods for collective decision making raised by recent advances in information and communication technologies allowing users to interact in complex new ways.

C.2 Objectives

The scientific objectives of the Action fall under two broad headings:

1. To apply the methods of Computer Science to the study of social choice so as to arrive at a clear understanding of the types of collective decision making procedures required to serve new technological challenges, in which computational concerns take centre stage.

2. To use insights from Social Choice Theory to derive better models of large distributed ICT systems in which autonomous entities, driven by different incentives and endowed with varying capabilities, can interact, in both cooperative and competitive environments.

At the organisational level, the main objectives of the Action are also twofold:

1. To create a strong network of researchers in Europe that can tackle the challenge of developing and analysing methods for collective decision making and that extends beyond the classical disciplinary boundaries of Computer Science and Economics.

2. To establish an inclusive research community for Computational Social Choice in Europe that creates career opportunities for Early-Stage Researchers, that includes scientists from a wide range of European countries, and that actively promotes gender balance.

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

The objectives stated above will be achieved through the following networking activities to be organised and coordinated by the Action:

- **Regular Meetings of the Working Groups:** The core of this COST Action are its four Working Groups. These groups will have the opportunity to organise regular meetings during the management meetings of the Action to be held twice a year. This will ensure that research activities taking place at the different sites can be coordinated appropriately and that information on the latest advances can be exchanged in a timely manner.
- **Short-Term Scientific Missions:** The meetings of the Working Groups will be complemented by a programme of Short-Term Scientific Missions (STSMs). These are research visits by individuals allowing for an in-depth exchange of ideas and collaboration. Both senior researchers and Early-Stage Researchers will be able to benefit from this programme, but the latter group will be particular encouraged and supported in this respect.
- **Training Events for Early-Stage Researchers:** Once a year the Action will organise a training event, such as a summer school, at which Early-Stage Researchers, particularly doctoral students, have the opportunity to receive intensive training from international experts. The theme of the training event will change each year and highlight a particular "hot topic". Besides their educational value, these training events will also provide an important means for networking amongst Early-Stage Researchers themselves.
- **Action Website:** The website maintained by the Action will provide an important source of information for the members of the Action, the broader research community, and the public at large. Making all relevant information available in one place will directly contribute to shaping the European research community in Computational Social Choice.

C.4 Potential impact of the Action

The most immediate impact will be on the research community itself. The Action can be expected to significantly influence the research landscape at the interface of Computer Science and Economics in Europe. While there is currently a lot of research activity in the field of Computational Social Choice, much of it very successful and highly visible, this research trend currently lacks coordination, which the Action will be able to provide.

In the long term, the work of the Action can be expected to have an important impact on a range of high-profile application areas. These include social networks, Internet search engines, e-governance, electronic commerce, and transplant organ allocation systems. The close involvement of several participants from the industrial sector will ensure an appropriate orientation towards such application areas from the beginning.

The technological impact of the work of the Action will have a ripple-on effect and influence policy makers (e.g., for transplant organ allocation) and society at large (e.g., via new developments in e-governance).

C.5 Target groups/end users

Due to the fundamental nature of the problems addressed by the Action, an important target group is the European research community. The Action will serve this community by providing a common framework for integrating advances made in nationally funded research projects.

Beyond the scientific community, also developers (and eventually, users) of applications that rely on collective decision making technology will be interested in and able to benefit from the activities of the Action. Examples include developers of electronic commerce systems or policy makers deciding on matching procedures to assign, for instance, interns to hospitals.

Stake-holders from industry will be involved in the Action from the beginning. This includes, in particular, representatives of the research departments of some of the world's best known high-tech companies, all of which have an immediate interest in the type of technology developments advanced by the fundamental work of this Action.

D. SCIENTIFIC PROGRAMME

D.1 Scientific focus

The scientific objectives stated in Section C.2 will be served by a scientific programme organised along three dimensions: (a) the nature of the social choice problem addressed, (b) the scientific methods employed, and (c) the long-term applications the work contributes to.

In terms of the first dimension (type of problem), the work coordinated by the Action falls into four areas:

1. voting and elections;
2. fair division, of both divisible and indivisible goods;
3. information merging, including judgment aggregation, belief revision and merging, ontology merging, and peer prediction methods; and
4. matching mechanisms.
5.

In terms of the second dimension (methods to be employed), besides the classical axiomatic method developed in Social Choice Theory itself, the Action will promote the innovative use of the tools of modern Computer Science, in particular: algorithm design and engineering to develop efficient social choice procedures, including approximation algorithms; complexity theory, including parameterised complexity theory, to understand fundamental limitations of algorithmic solutions; mathematical logic to provide fully formalised accounts of social choice problems, thereby making them amenable to automatic verification; and knowledge representation techniques developed in artificial intelligence to compactly model social choice problems with very large numbers of highly structured alternatives to choose from.

Regarding the third dimension (applications), the Action will address the challenges created by the fact that individuals may interact through the medium of a social network, the fact that both human users and automated software agents may be involved in joint decisions, and the fact that many applications give rise to domains of alternatives with a combinatorial structure. The Action will also coordinate the generation, collection, and distribution of libraries of benchmark problems derived from these applications.

This way of structuring the research programme both allows for efficient interaction between participants and is sufficiently open and flexible to enable scientists from countries that may not be represented in the Action from the start to join at a later stage and make a contribution.

D.2 Scientific work plan methods and means

The scientific work programme of the Action will be structured in terms of four Working Groups, according to the thematic threads outlined above:

- **WG1:** Voting and Elections
- **WG2:** Fair Division
- **WG3:** Information Merging
- **WG4:** Matching Mechanisms

The first Working Group, WG1, will concentrate on questions arising in the context of voting and elections. An election is the archetypal mechanism for making a social choice: to choose between several possible alternatives, each member of society is asked to provide a preference ordering over these alternatives, and a voting rule is used to compute a "best" alternative given this preference information. A wide range of voting procedures have been proposed in the literature, and many of them have been used in practice, both for large-scale political elections and for a multitude of other tasks where several agents have to choose from a set of alternatives. Voting technology, if suitably adapted, is relevant to a broad range of problems in ICT. For instance, in the context of Internet search engines, we may interpret a link from one website to another as a "vote"; this analogy shows that principles of social choice can play a role in the selection of high-quality web-content for display to the user. A second example are social networks, where a user may wish to call an election on a given question amongst her peers; while similar to a classical election scenario, the inherent structure of the social network raises new questions and allows for the formulation of new properties that we might want to see satisfied for the election mechanism employed. These new types of applications thus demand a fresh view at voting and elections that requires the expertise of political scientists, social choice theorists, and computer scientists.

The second Working Group, WG2, will focus on fair division. Fair division is the problem of dividing one or several goods amongst the members of a group in a manner that satisfies a suitable fairness criterion. Such fairness criteria may, for instance, refer to the well-being of the agent that is worst off, or it may promote solutions in which no agent envies any of the others. Fair division is a problem of social choice in the sense that every possible division of the goods constitutes a social alternative that the agents involved will have preferences over. Unlike standard problems in voting, however, fair division problems come with a high degree of internal structure, which is both a challenge and an opportunity. The structure stems from the fact that it is often reasonable to assume that agents would rather want more than fewer goods, and that an agent is likely to be indifferent between two alternative divisions that allocate her the same bundle. Mechanisms for fair division need to exploit this structure. Fair division problems with indivisible goods are examples for social choice problems with a combinatorial structure: even a relatively small number of goods will give rise to a very large number of potential divisions (e.g., there are over one million ways of distributing ten items amongst four agents, even if we do not allow for the possibility of dividing or sharing individual items). WG2 will address the algorithmic challenges raised by this combinatorial explosion in fair division problems.

The third Working Group, WG3, will cover a broad area that may best be described as information merging. Preferences are not the only type of information that we may wish to aggregate. Another important example is judgment aggregation. Here, each individual is asked to make a judgment about which of a list of propositions is true and which is false, and we are faced with the task of aggregating this information to come to a collective judgment. While originating in legal theory, the problem is clearly relevant to a range of technological challenges. A closely related approach is belief merging, where the goal is to compute a global representation of the beliefs of a group of agents. Methods of this kind have applications in crowd-sourcing tasks on the Internet. Yet another example is ontology merging, the task of integrating ontologies for the same domain but developed independently from each other in a manner that yields a consistent overall ontology but that also respects the input of the various contributors as much as is feasible. This question has important applications in the development of the Semantic Web. Ontology merging can naturally be cast as a social-choice theoretic question, a connection that has yet to be exploited. Finally, WG3 will also cover reputation systems as well as peer prediction methods, including prediction markets, where experts are asked to make an estimate regarding a question at hand and this information needs to be aggregated in a manner than encourages sincere estimates from those experts.

The fourth Working Group, WG4, will address the design and analysis of matching mechanisms. Examples for matching problems include internship allocation, public school choice, higher education admissions, or kidney exchanges (between patients, each associated with an incompatible donor that might be compatible with some other patient). Matching market design is an interdisciplinary research area of Economics, Game Theory, and Computer Science. Based on the particular characteristics of the matching markets at hand, a first task is to determine the desirable properties that a satisfactory algorithm should satisfy. These considerations may include criteria of efficiency, fairness, and non-manipulability and the importance and formulation of each criterion might differ depending on the socio-economic and cultural context. A second task is to establish the existence of a (class of) algorithm(s) that satisfy the requirements (i.e., prove the compatibility of the requirements), and determine their computational complexity. Additional techniques to design a matching algorithm for specific markets are experiments and simulations. WG4 will address the challenge of designing good matching mechanisms at all of these levels.

E. ORGANISATION

E.1 Coordination and organisation

The COST Action for Computational Social Choice will follow the standard rules and procedures for implementing COST Actions.

On some occasions the Action will seek to co-locate meetings with other scientific events, such as major international conferences. Such co-location will allow for an exchange of ideas with other research communities and provide additional opportunities for dissemination.

Once a year the Action will organise a training event for Early-Stage Researchers, particularly doctoral students. Teachers will be drawn both from the Action itself and from outside.

A further important instrument to be used by the Action are Short-Term Scientific Missions (STSMs) between different countries. An appropriate proportion of the budget will be reserved for STSMs by Early-Stage Researchers. The application procedure implemented to award STSM grants will conform to COST regulations and be kept as simple as possible. Decisions regarding STSM awards will be made by an STSM Committee appointed by the MC. There will be three deadlines each year at which the STSM Committee solicits and collects applications.

The Action will maintain a public website to inform about its activities and to make deliverables, including benchmark problems and educational resources, available to the research community and the public at large. The MC will nominate one of its members to ensure the Action website is kept up to date. The maintenance of the website will be made a regular agenda item for MC meetings.

Progress of the Action will be evaluated against the following milestones:

- **M0** (after 1 month): Election of members of the Steering Committee and nomination of MC member responsible for the Action website.
- **M1** (after at most 6 months): Ensuring broad awareness of the existence of the Action in the relevant scientific communities, through public launch of the Action website and wide dissemination of the announcement of the first plenary meeting of the Action.
- **M2** (after 12 months): First annual progress report, including summary of activities of each Working Group and the Action as a whole.
- **M3** (after 24 months): Second annual progress report (*idem*).
- **M4** (after 36 months): Third annual progress report (*idem*).
- **M5** (after 48 months): Final report, summarising activities over the full period.
- **M6** (after 48 months): Publication of an edited book with contributions from across the participants, highlighting key achievements of national and trans-national research programmes coordinated by the Action.

In addition, each working group will define a number of scientific milestones on an annual basis, in close interaction with the national and trans-national research programmes coordinated by the Action.

E.2 Working Groups

The activities of each of the four Working Groups will be coordinated by a Working Group Chair. Working Groups will meet twice a year during the general meetings of the Action. During these meetings, a small number of members of each group will present recent work, both within the group and to other interested members of the Action. The main body of the work will be carried out throughout the duration of the Action in nationally funded projects, enhanced by individual research visits of scientists supported by the Action in the form of Short-Term Scientific Missions.

E.3 Liaison and interaction with other research programmes

The COST Action on Computational Social Choice will liaise with the COST Action on Agreement Technologies (IC0801) on topics such a negotiation between autonomous agents. To quote from the Memorandum of Understanding of COST Action IC0801, "Agreement Technologies refer to computer systems in which autonomous software agents negotiate with one another, typically on behalf of humans, in order to come to mutually acceptable agreements". While the COST Action on Computational Social Choice focuses on high-level issues, such as the appropriate modelling of the notion of fairness in the context of making a multi-party agreement, the COST Action on Agreement Technologies bundles expertise on how to physically realise such agreements on a network of geographically distributed computers and end-users and how to appropriately model important concepts such as obligation and trust arising in this context. Taking into account the fact that COST Action IC0801 is now in its final year of operation, the liaison between the two COST Actions will focus on an exchange of information.

COST Action IC0901 ("Rich-Model Toolkit: An Infrastructure for Reliable Computer Systems") bundles European expertise in automated reasoning and aims to make this important technology applicable to a wider range of application domains. Computational Social Choice gives rise to new challenging problems for this technology: an important problem, which to date has only received limited attention and made relatively little progress, is the automated verification of properties of social choice mechanisms. The COST Action on Computational Social Choice will therefore liaise with COST Action IC0901 to initiate joint research in appropriate national programmes on this topic.

This liaison will be organised in terms of an exchange of information and meetings at the level of Working Groups. Specifically, the COST Action on Computational Social Choice will be in a position to provide new types of benchmark problems for COST Action IC0901.

The COST Action on Computational Social Choice will also seek contact with other COST Actions, including those that may start at a later time.

A possible point of contact with the FP7 programme of the European Commission is the recently launched FET Proactive Initiative on Fundamentals of Collective Adaptive Systems. This initiative addresses the operating and design principles underlying "systems that are constructed as a collective of heterogeneous components" in which "interaction between the [components] may lead to the emergence of unexpected phenomena". Some of these phenomena share similarities with the kind of phenomena studied in social choice, where the interaction of individuals may result in sometimes undesirable decisions taken at the level of the collective. The Action will offer its expertise to the research projects funded under this initiative.

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.

Concretely, care will be taken to appropriately represent female researchers in both the Management Committee itself and in the Working Groups. For instance, several of they key positions in the Action, such as chairs of Working Groups, will be held by female researchers.

The Action will organise a training event for doctoral students and other Early-Stage Researchers on an annual basis. Also amongst the participants in these events an appropriate gender balance will be respected. The same is true for the Short-Term Scientific Missions (STSMs) to be facilitated by the Action, which are specifically aimed at Early-Stage Researchers.

F. TIMETABLE

The estimated duration of the COST Action on Computational Social Choice is four years. It will start its operations with a kick-off meeting of the MC. During this kick-off meeting the MC will elect the members of the Steering Committee, who will ensure the proper implementation of the work programme and the timetable from then on.

The timetable for each year will be similar in structure: two Plenary Action meetings per year, each including both an MC meeting and a scientific meeting of the Working Groups, and one training event for Early-Stage Researchers. The details are summarised in the following table:

Year 1:
Month 1: Kick-off MC meeting. Forming of Working Groups. Election of Steering Committee.
Month 3: 1st Plenary Action Meeting (MC and Working Groups). Public launch of website.
Month 9: 2nd Plenary Action Meeting (MC and Working Groups). 1st Training School.
Year 2:
Month 3: 3rd Plenary Action Meeting (MC and Working Groups).
Month 9: 4th Plenary Action Meeting (MC and Working Groups). 2nd Training School.
Year 3:
Month 3: 5th Plenary Action Meeting (MC and Working Groups).
Month 9: 6th Plenary Action Meeting (MC and Working Groups). 3rd Training School.
Year 4:
Month 3: 7th Plenary Action Meeting (MC and Working Groups).
Month 9: 4th Training School.
Month 12: 8th Plenary Action Meeting (MC and Working Groups). Launch of edited volume withresearch highlights.

Working Groups are free to organise smaller meetings above and beyond the scheduled Action plenary meetings. STSMs will take place throughout the duration of the Action. The MC will aim at distributing these evenly across the four years.

Progress reports will be compiled in line with the plan set out in Section E.1 above. Deliverables required by the COST Office, such as financial statements and annual budget plans, will be produced according the the common timetable.

G. ECONOMIC DIMENSION

The following COST countries have actively participated in the preparation of the Action or otherwise indicated their interest: AT, BE, CH, CZ, DE, EL, ES, FI, FR, HU, IE, IL, IT, NL, PL, PT, SK, TR, UK. On the basis of national estimates, the economic dimension of the activities to be carried out under the Action has been estimated at EUR 76 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 target audience for the dissemination of the results of the Action consists of three groups:

1. The first group is the scientific community. This includes both researchers working in areas close to those addressed by the Action, such as Computer Science, Artificial Intelligence, Economics, and Formal Political Theory, and researchers working in other fields, who may not themselves conduct research on collective decision making but who may be able to make use of this kind of technology for their own research. An important goal for this group is community building and dissemination of knowledge regarding the state of the art in Computational Social Choice. As members of the scientific community are often also (university) teachers, these activities include the dissemination of educational materials. Special emphasis will be given to ensuring that the information provided by the Action will reach young researchers in the field.
2. The second group are companies in the industrial sector that are active in the area of Information and Communication Technology. Developers of social networking applications or electronic commerce systems, which increasingly involve forms of collective decision making, will be interested in the results of this Action.
3. The third group are policy makers who have to decide on collective decision making methods to be used for a number of different tasks in the public interest, such as rules for implementing organ transplant chains or rules for matching students to schools.

H.2 What?

The public website maintained by the Action will take a central role in all dissemination activities. All events organised by the Action will be announced here; teaching materials developed will be made available for download; and research reports will be linked from this website. This dissemination channel will be used to communicate with all three target groups identified above.

Dissemination activities for the first target group (the scientific community) also include traditional scientific activities such as publishing the results in peer-reviewed journals and proceedings of major international conferences, as well as organising workshops and other scientific events. Finally, events organised by the Action, such as workshops or summer schools, will also provide important opportunities for dissemination.

H.3 How?

The website will be regularly updated by a designated MC member and its content will be reviewed by the full MC.

All scientists belonging to the Action will regularly publish their findings in appropriate peer-reviewed journals and present them at workshops and conferences.

All four Working Groups will regularly be involved in the organisation of scientific events related to the thematic focus they are covering.

The training events organised by the Action can be expected to attract new young researchers to get involved with the topics covered by the Action.

All dissemination activities will be regularly monitored by the MC and the dissemination plan as outlined here will be updated whenever needed to ensure maximum visibility of the Action.