160th CSO Meeting 1st-2nd December 2004
Proposal for a new COST Action

COST 295

'DYNAMO'

Dynamic Communication Networks: Foundations and Algorithms

Contact Person: Pierre Fraigniaud
CNRS
Laboratoire de Recherche en Informatique (LRI)
Université Paris Sud
91405 Orsay cedex, France
Tel: +33 169 156 906
Fax: +33 169 156 586
Mail: pierre@lri.fr
http://www.lri.fr/~pierre

COST National Coordinator: Michel Gorlicki
Ministère de la Recherche et des Nouvelles Technologies
Bureau de la Coordination Européenne et des Relations Internationales
1, rue Descartes
75231 Paris Cedex 05
FRANCE
Tel:+33 (0) 1 5555 9972
Fax:+33 (0) 1 5555 9628
michel.gorlicki@technologie.gouv.fr

Rapporteur: Anders Hedin
VINNOVA
Liljeholmsvägen 32
10158 Stockholm
Sweden
Tel: +46 8 473 3000
Fax: +46 8 473 3005
Anders.Hedin@vinnova.se

DRAFT

Memorandum of Understanding

For the implementation of a European Concerted Research Action

designated as

COST 295

'DYNAMIC COMMUNICATION NETWORKS: FOUNDATIONS AND ALGORITHMS'

The Signatories of 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 the document COST/400/01 "Rules and Procedures for Implementing COST Actions", the contents of which the Signatories are fully aware of.
2. The main objective of this Action is to provide foundations, models, algorithms, and general tools for dynamic communication networks.
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 M€ 38 in 2004 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 4 years, unless the duration of the Action is modified according to the provisions of Chapter 6 of the document referred to in Point 1 above.

Technical Annex

COST 295

‘DYNAMO’
Dynamic Communication Networks: Foundations and Algorithms

A. Background

The usual parameters considered to measure the characteristics and the performances of networks (such as inter-node mean distance, cut-width, connectivity, diameter vulnerability, throughput, etc.) are no longer valid in the context of highly-dynamic environments, like peer-to-peer systems and the world-wide web at the application level, or ad hoc and sensor systems at the network level. Classical network analysis (cf. the seminal book by Ford and Fulkerson in 1962, or the book by Frank and Frisch in 1971), as well as stochastic models are by essence limited, and cannot capture the wealth of new and unexpected phenomena experienced in strategic real-world dynamic networks. This lack of foundations in the area of networks in general, and in dynamic networks in particular, is so problematic that a recent workshop of the USA National Science Foundation on Fundamental Research in Networking recommended “the development of *reference models* or benchmarks to ease the burden of reproducible experiments”. The conclusion of this workshop also stated that “emerging applications and directions require *fundamental* understanding on how to design and control networks on an even grander scale than before”.

This COST Action, called DYNAMO, is thus motivated by the need to supply a convincing theoretical framework for the analysis and control of all modern large networks induced by the interactions between decentralised and evolving computing entities, characterised by their inherently dynamic nature.

In Europe, research in the area of dynamic networks has been done in a scattered manner, both geographically and scientifically, with only a few exceptions like: (1) the 5th FP project COSINE, which gathered physicists focusing on the use of statistical mechanics and computer simulations, to study “complex networks”, and (2) the European Commission’s Future and Emerging Technologies initiative on “complex systems”. Moreover, the EU Industry is doing little fundamental research in this area, although a better knowledge of these networks would boost the creation of new tools and would stimulate the use of the new technologies (e.g., search engines, detection of communities, data mining, hybrid networks, etc.).

Moreover, research on fundamental and algorithmic aspects of dynamic networks is not only scattered around Europe, but is also fragmented by fields of expertise. Mathematicians, computer scientists, and engineers often approach the same problems with different techniques and tools from their own field of expertise. Results of these studies are published in various journals that are specialised in subjects as far apart as mathematics (e.g., Elsevier's Discrete Applied Mathematics), operations research (e.g., Wiley's Networks or INFORMS Journal on Computing), computer science (e.g., Academic Press' Journal of Parallel and Distributed Computing), information theory (e.g., World Publishing's International Journal of Foundations of Computer Science) and engineering (e.g., IEEE Transactions on Communications). The same fragmentation holds for conferences and workshops in these areas.

As a result, both the visibility of these studies and the integration of techniques from different fields are limited. The DYNAMO Action will coordinate the efforts toward the construction of the foundations of dynamic networks, efforts that are currently only carried out at a national level. Looking at the current structure of the research community in Europe, we identify a number of shortcomings that affect the strength and competitiveness of European industry, and that could be addressed by DYNAMO.

- Research on the fundamental and algorithmic aspects of dynamic communication networks is mainly carried out in small groups at university level. In Europe, hardly any industrial research lab exists where advanced network design, analysis, and provisioning is studied. In contrast, in the USA, large research groups funded by companies and universities are working in this area. Such is the case of, e.g., AT&T and Lucent's Bell Labs, or of the IRIS project at MIT.

- Within the scientific research community, research efforts in the foundations of dynamic networks are not coordinated at a European level. Potential gains are lost by the unawareness of recent results and intentions. The lack of collaboration and coordination often duplicates the effort, leading to a waste of resources.

- Within the scientific research community, cross-fertilisation among the different fields of expertise is limited. Although the fields of Algorithms, Discrete Mathematics, Mathematical Optimisation, and Distributed Computing are closely related and could benefit of each other's progress, exchange of theoretical and practical results among these areas is rare in general.

- For historical reasons, collaborations between academia and industry have been restricted (if any) predominantly at a national level. Companies all over Europe, however, are facing the same challenges in implementing new concepts and technologies but cannot benefit from progress relevant to them. Integrating the efforts at the European level is necessary to compete on a par with the USA and Japan.

- Specialist knowledge and a deep understanding of dynamic communication networks are necessary to develop world-wide scalable implementations of truly distributed services that are beginning to emerge (like peer-to-peer and ad-hoc networks, grid computing, or web services). Their commercial impact is exemplified by companies like Akamai (which, incidentally, was founded by a group of renowned US algorithmic scientists). Right now, this knowledge and understanding is scattered over European research groups, and practically lacking within Europe's industry. This implies that the next wave in Internet-based application development will be almost exclusively driven (and capitalised upon) by US companies.

- Last but not least, the strategic factor should not be underestimated. Gaining strategic knowledge about dynamic networks, like the web-graph structure (for search-engines) and the Internet topology (for security), is of paramount importance. The current status shows that US companies, government, and researchers, have a large advantage over their European counterparts.

In addition to the above, we observed that research on the design and analysis of networks has been traditionally based on methodologies drawn from statistical science (e.g., queuing theory) and on ad-hoc engineering approaches. It is presently widely acknowledged that algorithmic approaches to the core problems in networks are quite relevant and that we are actually witnessing a paradigm shift in the approach toward networks in general. Therefore, COST offers an excellent opportunity

for starting a much needed European co-operation at the frontier of Computer Science and Networking, focusing on Dynamic Networks, because the strong algorithmic European research community will only be able to play a major role in the field if it is integrated at a European level.

Related European initiatives

DYNAMO shares some of its scientific background with COST 293: Graphs and algorithms in communication networks. The main objective of COST 293 is to elaborate advances in the design of communication networks through close cooperation between mathematicians and researchers from the networking community. In contrast, DYNAMO focuses on dynamic networks, and considers all aspects of these networks (not only design). Several other COST Actions are focused on various aspects of networks, such as 272, 279, and 290 on networking technologies, and 273, 285 and 289 on broadband and services. Again, these Actions focus on network design and/or provisioning from the networking point of view. DYNAMO aims at federating research on fundamental aspects of dynamic networks, from the Internet to the Web, from sensor to peer-to-peer, in order to bring up algorithmic technologies that can be applied to a wide range of problems in the dynamic network framework.

Concerning ongoing European projects, there are three IST Networks of Excellence (NoE) which are related to this Action: EuroNGI (http://www.eurongi.org/), E-Next (http://www.ist-e-next.net/), and NEWCOM (ftp://ftp.cordis.lu/pub/ist/docs/directorate_d/cnt/newcom_en.pdf). The former is on the design and engineering of the next generation Internet, the second is aiming at allowing the integration of computers and networks into the everyday environment, and the last one is on the deployment of wireless networks. Although these NoEs do some research in the area of dynamic networks, their approach is networking and engineering oriented, whereas DYNAMO is oriented towards foundations and algorithms. In some sense, DYNAMO can be viewed as a complement of these NoEs.

With a more multidisciplinary approach, one finds the IST-FET NoE DELIS (http://delis.upb.de), whose main goal is to develop methods, techniques and tools to cope with challenges imposed by large scale systems. DELIS is taking an interdisciplinary approach to the issue by combining the fields of Computer Science, Physics and Economy. Finally, the FP5 Project COSIN (http://www.cosin.org) is also related to the DYNAMO Action. This project aims at developing statistical models based on agent interactions, and inspired by the theory of self-organisation and fractal growth, to describe network evolutions, with applications to economic networks.

DYNAMO will actively look for synergy with these related European scientific networks. In particular, it is planed that leaders of these Actions and projects will be invited to participate at meeting(s) of DYNAMO, so that our Action can:

- benefit from the results of these networks, as far as the foundational and algorithmic aspects of dynamic networks are concerned, and
- avoid performing investigations that may have been already carried out for other contexts, using a different vocabulary, or for the purpose of different objectives.

B. Objectives and Benefits

The main objective of this Action is to provide foundations, models, algorithms, and general tools for dynamic communication networks. These new decision-support tools will favour the study and the efficient design of applications for networks of decentralised interacting and evolving entities, experiencing possibly brutal modifications of their environments.

DYNAMO will integrate the state-of-the-art scientific knowledge from fields relevant to dynamic communication networks, in order to develop decision-support network technologies under the guidance of its local industrial partners. These sophisticated decision-support tools will help overcome the increased complexity of questions induced by the dynamics of advanced decentralised systems. Among the main problems to be tackled we find: placement of equipment, spectrum assignment, bandwidth allocation, infrastructure design and distributed systems design.

Through the constitution of a COST Action in this field, DYNAMO will provide to the European ICT Industry the scientific foundations needed on dynamic communication networks design and provisioning. This will help towards optimising the deployment and usage of high-end technology communication networks, like mobile networks, radio networks, cellular networks, satellite networks, ad hoc networks and sensors networks, as well as of societal-driven emerging information technologies, like the World Wide Web and peer-to-peer networks.

The advent of a global ICT economy and the staff reduction in private R&D centres triggered a brain-drain toward the USA, where salaries are higher, professional conditions are better, and scientific research (and ensuing patents) is highly praised by companies like AT&T, Lucent Technologies, or yet Telcordia Technologies. As a result, the influence of Europe on technological development may decrease, whereas the dependence on non-European solutions may increase. Moreover, financial investments involved in emerging communication technologies are nowadays so high, that failures can have serious consequences for the employment in these areas in Europe, as witnessed by the recent race toward UMTS licenses. Other specific domains where European ICT companies could benefit from an integrated European research within DYNAMO include, but are not limited to: security, fair pricing, (dynamic) infrastructure design, dynamic resource allocation, traffic and mobility management, flow and congestion control, routing, energy consumption, reliability, etc. For many of these topics, it is conceived that not only European businesses may benefit from advanced research on dynamic networks, but also the society as a whole. In this context it is also important to point out that many applications in the area of dynamic networks have a strong social impact with particular emphasis on: people security, anti-terrorism defence structure, emergency operations, data acquisition in disaster areas in connection with earthquakes or volcanic eruptions, etc.

The DYNAMO Action will be able to build on the previous and existing European-funded, small-scale cross-cooperation in the area, in order to span more than 50 European organisations. With our strategy based on scientific excellence and industrial relevance, we will be able to concentrate the European know-how in our domain, in a way not far from the USA's scientific societies (ACM, IEEE, SIAM, INFORMS, etc.), but more tightly connected and, above all, steered to respond to European needs.

C. Scientific Programme

The scientific objectives of DYNAMO are focused around the following main aspects:

1. Characterisation of common properties shared by large-scale dynamic complex networks (and their communication aspects) arising from the interaction of basic entities, including Internet, sensor, cellular and ad-hoc networks, and social networks (e.g., the WWW, peer-to-peer, etc.).

2. Design of new dynamic network models (and general communication schemes), optimised with respect to the new measures resulting from the above, while improving other system characteristics (reliability, communication delays, etc.).

3. Construction of distributed protocols (in the standard setting, or combined with the use of mobile agents) capable of extracting and processing local and global network information.

4. Design of basic practical algorithms (for routing, control, searching, etc.) to be used in the development of network-based applications.

5. Validation of the results using data extracted from real self-organising dynamic networks, by the use of reproducible simulations and benchmarks.

From the point of view of the methods considered to achieve the expected results for the Action, we will use classical research techniques from applied mathematics, networking, and computer science, such as modelling, analysis of algorithms, simulations, etc. We will make extensive use of our knowledge on graph theory, analysis of algorithms, statistical tools, communication networks (topology issues, embeddings, routing, communication schemes) and network technologies.

DYNAMO will be organised into four major Working Groups (WG), and three transversal Co-ordination Activities (CA).

C.1. Working Groups

The proposed four Working Groups of DYNAMO will focus on:

WG1: Small Worlds

This WG focuses on mapping and modelling large complex and continuously evolving networks (e.g., Internet or the Web). The models should capture the small-world and scale-free nature of these networks. The mapping (i.e., crawling) uses specific agents for performing measures and extracting information from large dynamical networks. In addition, this WG is in charge of developing communication primitives and basic algorithms for ranking and searching large dynamic networks. A key objective of this WG is to produce data sets from both real and simulation dynamic networks, to allow validation of the small world models and algorithms. Hence, WG1 aims at providing practical measures, models, methods and algorithms to (1) deal with large, complex and continuously evolving networks, and (2) help in the design of new content network applications.

WG2: Wireless Networks: Sensor, dust, and ad-hoc networks
Sensor and ad-hoc networks represent emerging technologies in the field of wireless networks. Sensor networks connect a large number of distributed autonomous electronic sensor devices gathering data for environmental screening, security surveillance, and many other applications. While sensor networks connect automatic devices, ad-hoc networks address also personal devices like laptops, mobile phones, personal digital assistance (PDA). The key difference between ad-hoc networks and cellular networks is the fact that ad-hoc networks work without any infrastructure, like base stations. To allow significant progress in this area it is necessary to close the gap between the nearly unlimited physical and hardware possibilities and the little knowledge about the design and distributed construction methodologies of such networks. This WG concentrates its effort on the design of low-energy algorithmic solutions for: self-organisation, routing, gathering, etc. This includes not only the design of new protocols, but may require original solutions based on new addressing and labelling techniques. Although very different in nature, LEO satellite networks also fit in the WG for they are constantly moving, which implies the need of sophisticated routing and control protocols.

WG3: Peer-to-Peer Networks
Peer-to-peer (P2P) applications will be the mechanism allowing people, on the long term, to harness the enormous potential of decentralised systems, yielding vastly increased computational capabilities, robustness and high scalability. The multitude of peers represent a resource of both CPU and storage capacity with scaling and resilience properties unequalled by any client-server system offered today. Research in WG3 concentrates on developing uniform architectures and platforms for powerful, scalable, fault resilient, and secure P2P systems. More broadly, this WG deals with the development, adaptation and testing of algorithms for P2P distributed systems (searching, publishing, etc.). The strengths and limits of solutions based on attractive technologies such as Distributed Hash Tables (DHT) are specifically addressed. Additional objectives of this WG are to transfer knowledge to Industry, develop better understanding, more uniform treatment and more standardised deployment of P2P systems, leading to improved, convenient and reliable utilisation of global resources. (With respect to the dynamic properties of P2P networks, this WG has close relationships with some topics considered in WG2 on wireless networks, e.g., with ad-hoc networks.)

WG4: Emerging Algorithmic Technologies
In traditional theoretical computer science, the algorithm designer usually makes the assumption that the participating computational entities (agents) will follow the prescribed algorithm or will play against each other. In game theory, instead, the participating strategic agents are neither obedient nor adversarial and act in a selfish way so that, although one cannot assume that the agents will follow the prescribed algorithm, it is reasonable to assume that they will respond to incentives. Game theory is a powerful tool that is already being used to analyse network protocols and distributed algorithms: within this theory, mechanism design is a somehow inverse task, whose aim is to design algorithms (games) so that agents' selfish behaviour results in the desired system-wide goals. This WG studies the practical impact on the interacting entities involved in dynamic networks of new emerging algorithmic technologies that address incentive and computational complexity simultaneously. These technologies provide solutions that contain both an algorithmic ingredient and a “reward” ingredient that motivates the entities involved in the algorithms (e.g., ad hoc nodes). The goals of this WG include (in collaboration with the other related WGs) developing mechanisms for routing, power assignment in radio networks, web caching, peer-to-peer file sharing, overlay networks, and distributed task allocation.

C.2. Co-ordination Activities

The four aforementioned WGs cover a large scientific domain. Hence, the Management Committee will co-ordinate and integrate the activities of the WGs (cf. Figure 1) via three Co-ordination Activities, which are transversal to all WGs. Specifically, DYNAMO will include the Co-ordination Activities described below. Their operational details will be decided by the Management Committee, depending also on the expertise of its nominated members.

CA1: Cross-fertilisation

This CA will help researchers from different WGs to share their common expertise, including algorithmic tools and mathematical techniques. New advances in any of the WGs will be communicated to the other WGs in order to avoid duplication of research and efforts. Exchanges between the WGs will be promoted, typically by financially supporting inter-WG visits, and main results of every WG will be presented at the plenary meetings of DYNAMO.

CA2: Large Networks for Benchmarks/Test-Beds

In each of the WGs, and even within a same WG, partners of the DYNAMO Action use their own communication networks examples (real-world or artificially created) when they reach the stage of confronting to the real-world the algorithms they design. Hence, DYNAMO will gather all these networks files, collect additional files from some of DYNAMO's industrial partners and from existing sources on the Internet, in order to design a coherent, organised, and standardised repository of large communication networks data files that can be used as a source for comparative studies. This repository of real-world graphs and networks will be maintained and made accessible on the Internet to all DYNAMO partners and to the research community. It will be used as a source of benchmarks for performance evaluation of algorithms for communication networks.

CA3: Convergence with other Aspects of Dynamic Networks

Most advances in the study of dynamic networks demonstrate that networks arising from different fields (communication networks, sociological relationships, virus propagation, economical exchanges, etc.) share common properties, often summarised as characterising “complex networks”. Although DYNAMO mostly aims at investigating dynamic *communication* networks, we also foresee cooperation with scientists from other domains such as sociology, medicine, physics, and economy. Therefore, we will promote the convergence of the DYNAMO community with other communities dealing with “complex networks”.

D. Organisation

DYNAMO integrates researchers with diverse scientific backgrounds together with industrial partners, thus creating a research body that can address fundamental problems in dynamic communication networks of great practical importance, devise accurate, broadly agreed upon models and produce comprehensive and efficient algorithmic solutions. Such a strong coordinated effort will create a focal point that will attract additional researchers in algorithms to work in this direction, increasing the visibility of the achievements to industry, and by that, strengthening the position of Europe in this field.

The integration and coordination is achieved by means of a number of instruments, as follows.

1. **Meetings, workshops and schools.** Management Committee meetings will be mostly devoted to orientation and management issues, while workshops and schools aim towards disseminating knowledge and exploiting results. Workshops will focus on the identification of those trends within the area whose developments are particularly interesting to industry, and that deserve further study within DYNAMO. Particular topics will be studied with explanatory talks, brainstorm sessions, and assessments to steer further research. At the same time, research results achieved in the previous period are presented and their applicability is discussed. Finally, schools will guarantee a regular training of the involved researchers, in particular young researchers and engineers, on recent developments, not only in their own field of expertise but also in related fields to foster cross fertilisation. The proposed schedule of all these events is shown in Table 1.

2. **Short Term Scientific Missions.** These reciprocal visit exchanges among the partners are intended to accelerate research on specific topics in dynamic networks theory, as well as to disseminate deep insight knowledge within the Action.

3. **A web-site.** Through an internal web-site, DYNAMO members can track research trends, learn about the latest results, identify promising directions for further research, discuss open problems, etc. Through an external web-site, research results are disseminated to potential users in academia, industry, and government. This web-site will have the potential to serve as an entry-point for industry to obtain information about state-of-the-art mathematical research results and acts as an intermediate between industry and academia establishing contacts with the best-suited research group(s) to answer the needs of industry.

The implementation of each of these instruments will have a considerable impact on the way research is done and organised in Europe. The consulting expertise of the DYNAMO network members may help establish a reliable and efficient pan-European bridge between industrial companies and academic institutions. Moreover, a substantial progress in theoretical and applied research is expected from the direct collaboration between most of the partners with industry. All of the DYNAMO activities above will be under the supervision of the DYNAMO Management Committee.

Figure 1 represents the role of the CAs for the coordination of the WGs. As mentioned above, the MC shall decide on the best way to implement such activities.

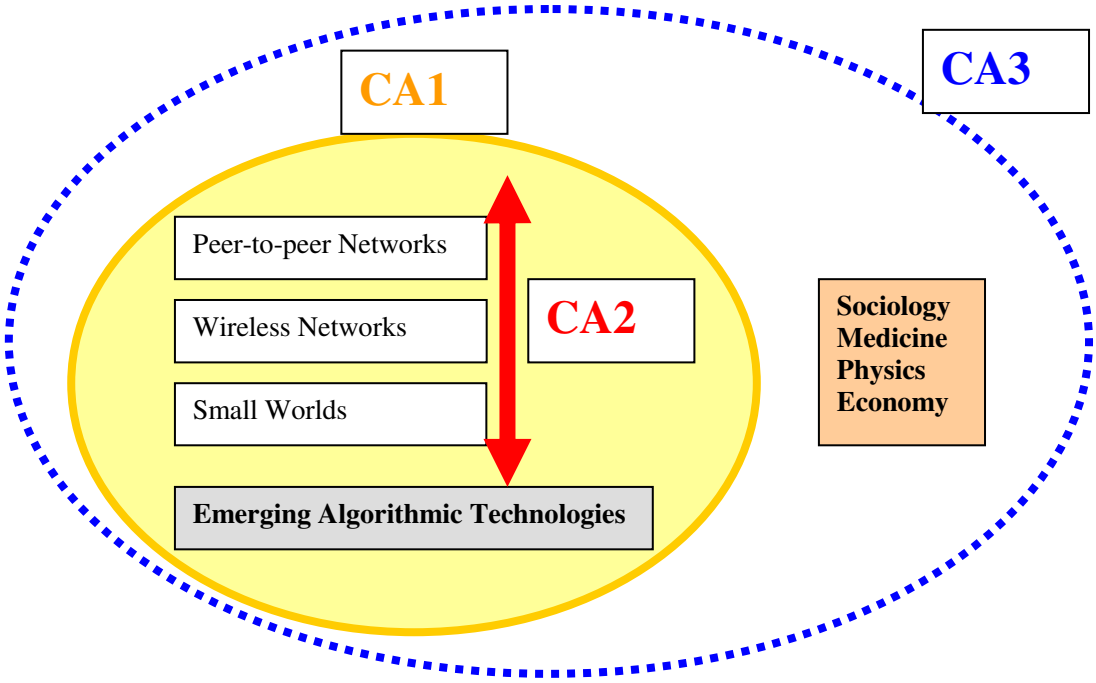

Figure 1. DYNAMO WGs and their co-ordination.

E. Timetable

Figure 2 describes the global scheduling of the WGs and CAs.

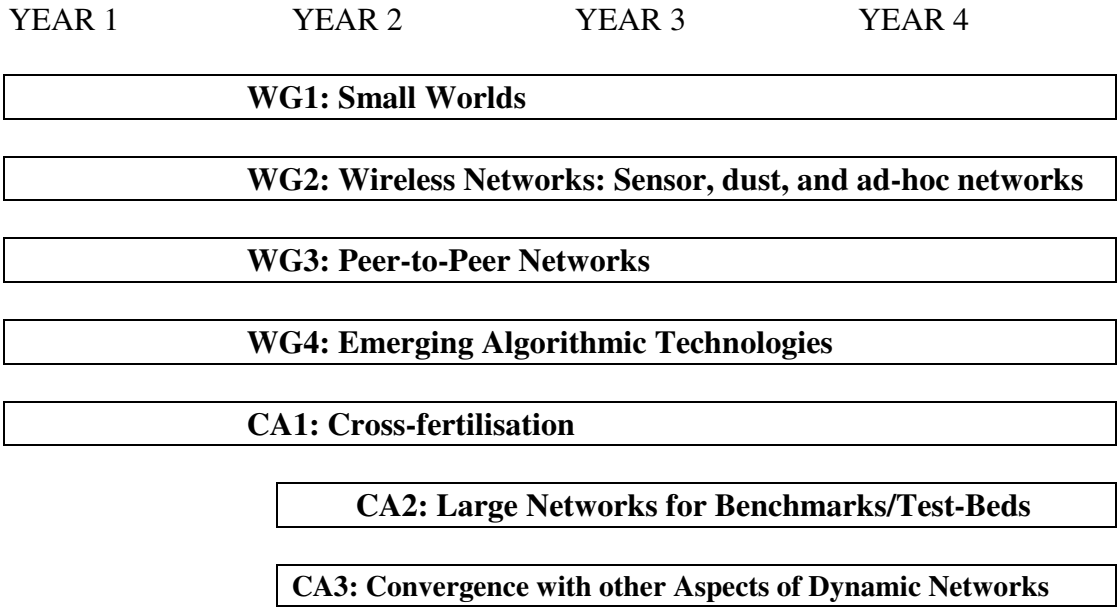

Figure 2. Global scheduling of Dynamo WGs and CAs.

The Management Committee will meet twice a year. In order to be economically efficient, each of these meetings will be accompanied by WG meetings, which are organised in parallel. Immediately after the parallel WG meetings, a plenary meeting shall be scheduled.

The Management Committee meetings aim at controlling and co-ordinating the activities of the Working Groups, deciding the scientific content of the schools and open workshops, managing the budget, allocating support for the STSMs (DYNAMO plans to support around 8 short visits/year), etc. One main objective of the plenary meeting is to ensure cross-fertilisation between the several WGs, for sharing expertise and knowledge. In particular, the plenary meetings aim towards identifying common problems from different WGs for which the same algorithmic techniques can be used.

The WG meetings gather researchers involved in every corresponding Working Group. These meetings, organised in parallel, precede each plenary meeting of the Action, during which the main recent results obtained by the Action, in each Working Group, are presented to all participants, increasing awareness and co-ordination. The Management Committee will ensure that the integration of the research carried out by the different WGs is progressing as expected (expertise sharing, common benchmarks and test-beds, etc.).

DYNAMO will pay attention to organise its meeting at the same place as major scientific events (such as annual international conferences, e.g., ICALP, STACS, ESA, etc.), immediately before or after these events. As a consequence, (1) the travel expenses of the DYNAMO meeting participants will be limited, and (2) the public awareness of the Action will be promoted.

DYNAMO will organise four major yearly events, alternately: two training schools and two workshops. These events will be open to all participants and widely advertised. A final workshop, within three months of the end of the Action will also be organised. The global timetable for the four years of DYNAMO can be found in the Table 1, below (every year will follow roughly the same planning).

Table 1. Timetable for the four years

Year 1	Kick-off meeting
Year 1	1st DYNAMO meeting (2.5 days)
Year 1	2nd DYNAMO meeting (2.5 days)
Year 1	1st DYNAMO training school (1 week)
Year 2	3rd DYNAMO meeting (2.5 days)
Year 2	4th DYNAMO meeting (2.5 days)
Year 2	1st DYNAMO Workshop (1 week)
Year 3	5th DYNAMO meeting (2.5 days)
Year 3	6th DYNAMO meeting (2.5 days)
Year 3	2nd DYNAMO Workshop (1 week)
Year 4	7th DYNAMO meeting (2.5 days)
Year 4	8th DYNAMO meeting (2.5 days)
Year 4	2nd DYNAMO school (1 week)

F. Economic Dimension

The following COST countries have actively participated in the preparation of the Action or otherwise indicated their interest: Belgium, Cyprus, Czech Republic, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland, and UK. On the basis of national estimates provided by the representatives of these 19 countries, the economic dimension of the activities to be carried out under the Action has been estimated, in 2004 prices, at roughly **38 Million Euros**.

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.

G. Dissemination Plan

The dissemination activities planned within DYNAMO are as follows.

1. Two training schools on specific fields of investigation will be organised in order both to disseminate DYNAMO's expertise to researchers coming from outside the Action, and to train DYNAMO's participants. These schools will be widely advertised, and participants outside DYNAMO are particularly welcome.
2. DYNAMO is aiming at promoting the public awareness of the Action by means of traditional and modern communication media (Web, peer-reviewed journal papers, presentations at international conferences, etc.). The research community in the domain of foundational and algorithmic aspects of dynamic networks, and the European companies in the communication sector may eventually consider the DYNAMO network as the hub of information in this domain.

In the longer term, DYNAMO might even set up a quality label for European activities related to foundational and algorithmic aspects of dynamic networks. In this way, workshops, conferences, colloquia and schools which are already regularly, but scattered, organised, will have a common and improved visibility.

Close relation with Industry. The multi-disciplinarity of the DYNAMO Action is only one example of its richness. The close relations that most members nourish with industry through formal cooperation are another guarantee of the success of our approach. DYNAMO members have contacts with the following IT companies (list non-exhaustive): Aerocomputer, Alcatel France, Alcatel Italia, Alenia, Bouygues Telecom, British Telecom, Deutsh Telekom, Ericson, Etnoteam, Exensus AB, France Telecom, IBM, Infineon, Kowalk, KPN, Microsoft, Neta, Nexse, Nokia, Siemens, Spazio, Swisscom, Telenor R&D, Teleste, Telia, TIM, UK Radiocommunications Agency, Vimercate, and WIND. Several of these companies have already accepted to be external partners of the DYNAMO Action, and most have expressed a deep interest in our Action. DYNAMO may eventually become the one-stop counter, where Industry can find the right academic partner(s) to resolve their problems.

COST 295
‘DYNAMO’

Dynamic Communication Networks: Foundations and Algorithms

ADDITIONAL INFORMATION
NOT PART OF THE MOU

H. Appendix - Additional Information

H.1. More about the economic dimension of DYNAMO

As mentioned in Section F, the economic dimension of the activities to be carried out under the Action has been estimated, in 2004 prices, at roughly **38 Million Euros**. This accounts for 2 M Euro per Signatory, for the four-year period of the Action, i.e., 500 K Euro/year per Signatory. These costs represent around 10 persons/year, some part-time, plus travel expenses for each Signatory of the Action.

H.2. History of the Action

DYNAMO was motivated by the lack of coordination between the different scientific actors involved in fundamental research on dynamic networks. As already mentioned, these actors are currently spread out geographically (i.e., small groups all over Europe), scientifically (i.e., from different scientific fields), and economically (i.e., little cooperation between industry and academia). An important step towards the integration of this community was done in France, where the researchers have organised themselves around a CNRS Action named TAROT[1], and some projects with the Ministry of Research. The DYNAMO Action follows up several discussions among scientific European leaders in the field of dynamic communication networks. During a meeting, organised in France in February 2003, it was decided to set up a pan-European initiative able to manage research on dynamic networks, including the promotion of the most important fundamental results, and the coordination of the European scientific efforts in this field.

The participants of the February 2003 workshop faced, however, the lack of European support towards fundamental research in networking, as witnessed by the topics of the FP6 call. This was very surprising, almost shocking even, since the complexity of the current communication networks (from Internet to sensors, from the Web to peer-to-peer systems) requires new tools for their analysis, as well as new ideas for their management. These ideas and tools cannot result from a single scientific domain, but will result from the coordinated efforts of different communities. This COST Action aims at being the coordinator of such efforts.

H.4. List of Experts

We list hereafter most of the industrial and academic potential partners contacted during the writing of this Action. The Action will be open to new participants, according to COST rules and procedures.

List of Industrial Experts

Company	Country	Chief Scientist(s)	Domain of Application
Alcatel Italia	Italy	C. Spinelli	Broadband network design
Bouygues Telecom	France	F. Chauvet	Wireless network design
Ericsson	Sweden	U. Bengtner, S. Korsback	Peer-to-Peer Systems

[1] TAROT stands for *Techniques Algorithmiques et Réseaux pour l'Optimisation des Télécommunications.*

MostlyTek Ltd.	Israel	M. Ran	Broadband network provisioning
NETA	Italy	A. Alleva, M. Secone	World Wide Web
T-Mobile	Germany	R. Hager	Wireless network provisioning
TNC-Telecom	Germany	M. Hellebrandt	Wireless network provisioning

List of Potential Academic Partners

Country, City (Organization)	Chief Scientist	E-address
Belgium, Louvain (U Catholique)	V. Blondel	<blondel@inma.ucl.ac.be>
Canada (*), Gatineau (UQO)	A. Pelc	<pelc@uqo.ca>
Canada (*), Ottawa (U, Carleton)	P. Flocchini	<flocchin@site.uottawa.ca>
Cyprus, Cyprus (U)	M. Mavronicolas	<mavronic@ucy.ac.cy>
Czech Rep., Prague (TU)	P. Tvrdik	<tvrdik@sci.felk.cvut.cz>
Finland, Kuopio (U)	M. Penttonen	<penttonen@cs.joensuu.fi>
France, Lyon (INRIA, INSA)	E. Fleury	<Eric.Fleury@inria.fr>
France, Paris (CNRS, U. Orsay)	P. Fraigniaud	<pierre@lri.fr>
Germany, Paderborn (U, Nixdorf)	C. Schindelhauer	<schindel@upb.de>
Germany, Freiburg (U)	S. Albers	<albers@informatik.uni-freiburg.de>
Greece, Patras (CTI)	C. Kaklamanis	<kakl@ceid.upatras.gr>
Greece, Ioannina (U)	P. Fatourou	<faturu@cs.uoi.gr>
Hungary, Budapest (U of Technology)	R. Vida	<Rolland.Vida@tmit.bme.hu>
Israel, Haifa (Technion)	S. Kutten	<kutten@ie.technion.ac.il>
Israel, Rehovot (Weizmann)	M. Naor	<moni.naor@weizmann.ac.il>
Italy, Florence (U)	P. Crescenzi	<piluc@dsi.unifi.it>
Italy, Roma (La Sapienza)	A. Marchetti-Spaccamela	<alberto@dis.uniroma1.it>
Netherlands, Twente (U)	J.-H. Hoepman	<jhh@cs.kun.nl>
Netherlands, Utrecht (U)	J. van Leeuwen	<jan@cs.uu.nl>
Norway, Bergen (U)	F. Fomin	<Fedor.Fomin@ii.uib.no>
Poland, Warsaw (U)	K. Diks	<diks@mimuw.edu.pl>
Portugal, Lisbon (U)	L. Rodrigues	<ler@di.fc.ul.pt>
Slovakia, Bratislava (Acc.of Sc.)	I. Vrto	<vrto@savba.sk>
Spain, Barcelona (UPC-CS)	M. Serna	<mjserna@lsi.upc.es>
Sweden, Gothenburg (Chalmers)	M. Papatriantafilou	<ptrianta@cs.chalmers.se>
Sweden, Lund (U)	A. Lingas	<Andrzej.Lingas@cs.lth.se>
Switzerland, Geneva (U)	J. Rolim	<Jose.Rolim@cui.unige.ch>
Switzerland, Zurich (ETH)	R. Wattenhoffer	<wattenhofer@inf.ethz.ch>
UK, Liverpool (U)	L. Gasieniec	<leszek@csc.liv.ac.uk>
UK, Warwick (U)	M.Paterson	<msp@dcs.warwick.ac.uk>

Note: Participants from Canada, marked with a * symbol, are expected to participate using their own funds.