European Cooperation
in Science and Technology
- COST -

Secretariat

Brussels, 2 July 2010

COST 4144/10

MEMORANDUM OF UNDERSTANDING

Subject :	Memorandum of Understanding for the implementation of a European Concerted Research Action designated as COST Action IC1001: Transactional Memories: Foundations, Algorithms, Tools, and Applications (Euro-TM)

Delegations will find attached the Memorandum of Understanding for COST Action IC1001 as approved by the COST Committee of Senior Officials (CSO) at its 178th meeting on 25 May 2010

MEMORANDUM OF UNDERSTANDING

For the implementation of a European Concerted Research Action designated as

COST Action IC1001

TRANSACTIONAL MEMORIES: FOUNDATIONS, ALGORITHMS, TOOLS, AND APPLICATIONS (EURO-TM)

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 4159/10 "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 aim of the Action is to enhance the quality, visibility and impact of European research in the area of TMs by establishing a network of ICT researchers and developers to address the TM theoretical underpinnings and to facilitate the development and showcase of TM platforms and applications.

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 48 million in 2010 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

Parallel programming (PP) used to be an area once confined to a few niches, such as scientific and high-performance computing applications. However, with the proliferation of multicore processors, and the emergence of new, inherently parallel and distributed deployment platforms, such as those provided by cloud computing, parallel programming has definitely become a mainstream concern. Transactional Memories (TMs) answer the need to find a better programming model for PP, capable of boosting developers' productivity and allowing ordinary programmers to unleash the power of parallel and distributed architectures avoiding the pitfalls of manual, lock based synchronization. It is therefore no surprise that TM has been subject to intense research in the last years.
This Action aims at consolidating European research on this important field, by coordinating the European research groups working on the development of complementary, interdisciplinary aspects of Transactional Memories, including theoretical foundations, algorithms, hardware and operating system support, language integration and development tools, and applications.
Keywords: Transactional Memories Concurrent, Parallel and Distributed programming Multi-threaded applications Multi-core Processors Cloud Computing

B. BACKGROUND

B.1 General background

Over the past 30 years, the advent of a new generation of CPUs (Central Processing Units) has regularly corresponded to a free performance gain for most software applications. Since the beginning of the last decade, however, the architectural trend toward ever-increasing CPU clock speeds has taken a sharp turn. As the performance gains from increasing the processors' operating frequency have steadily decreased till hitting the point of diminishing returns, the major CPU manufacturers had eventually to resort to the adoption of multi-core architectures in order to keep on chasing the Moore's Law.

This architectural trend will allow the forecoming generations of processors to continue on achieving regular gains in terms of maximum achievable throughput, and some experts believe that, in contrast to commercially available four and eight core machines in 2009, by 2017 embedded processors could sport 4,096 cores, server CPUs might have 512 cores and desktop chips could use 128 cores. Differently from the past, however, the increase in the speed of execution of straight-line instruction flows is going to be only marginal.

The advent of multi-core architectures has therefore decreed the end of the free performance gains' era for single-threaded applications, which happen to represent the vast majority of today's software. Hereafter, unleashing the full potential of multi-core processors demands a radical shift in the way software is developed, moving parallel programming from the niche of scientific and high performance computing to mainstream application domains.

Unfortunately, writing scalable parallel programs using traditional lock-based synchronization primitives is well known to be a hard, time consuming, and error-prone task, mastered by only a minority of programmers. Thus, to bring parallel programming into the mainstream of software development, we are in urgent need of better programming models.

Building on the abstraction of atomic transactions, and freeing programmer from the complexity of conventional synchronization schemes, Transactional Memories (TMs) promise to respond exactly to this need, simplifying the development and verification of concurrent programs, enhancing code reliability and boosting productivity.

Over the last five years, research on TMs has received a huge attention in computer science and engineering, with hundreds of papers published in prestigious international conferences and journals addressing a wide range of complementary aspects including hardware and operating systems (OSs) support, language integration, as well as algorithms and theoretical foundations. Moreover, whereas TMs were first introduced in the context of shared-memory multiprocessors, they have been recently extended to distributed environments with the same goal of simplifying the programming model of such systems.

The growing interest towards TMs has not only been confined to academic environments. On the industrial side, some of the major players of the software and hardware markets have been up-front in the research and development of prototypal products providing support for TMs. Further, TMs are starting to be used with success in real-world systems, and the preliminary feedbacks gathered from the field are confirming the achievement of the promised reductions of the time to market's and bug density.

Overall, the research on TMs appears to have an enormous potential in affecting the way today's and tomorrow's software is going to be developed, but there is still a long road ahead before TMs can turn into a widespread and really effective tool. On the performance front, TM systems still lack the right hardware and software support to provide scalability not only in terms of number of cores, but also in terms of code size and complexity. On the theory side, it is only very recently that the scientific community has started to address the problem of formalizing the precise guarantees that TM implementations should provide. At the language level, the lack of standard programming APIs (and semantics) is severely hampering the portability of TM-based applications and discouraging the penetration of this technology in the software industry.

Unfortunately, up to date, research in the area of TMs has been done in a scattered manner, both geographically and in terms of fields of expertise. Research on TMs spans a number of different areas such as hardware design, operating systems, compilers, theory of concurrent and distributed computing, each of which having its own specialized publication venues (journals, conferences, and workshops). The result of this fragmentation is that both the visibility of these studies and the integration of techniques from different areas are limited. Also, the lack of an effective of coordination of research is leading to overlap or even redundancy, and to progress in small increments.

Today, research on TMs is also hampered by a dangerous vicious circle. Being a young research field, driven mainly by academic researchers, only a very limited number of real-world, complex applications have been up to date developed for TMs. Unfortunately, this forces researchers to design their solutions based on synthetic (micro-)benchmarks, whose representativeness of real workloads is at least arguable. This raises sensible concerns on the robustness of existing TM solutions once employed in production systems, refraining the penetration of this technology in the software industry and, ultimately, the development of realistic benchmarking applications that would provide valuable feedbacks and insights to researchers. This Action aims at tackling the above issues by establishing a pan-European network bringing together almost all leading academic and industrial researchers in TMs with software industry practitioners active in the development of applications for multi-core and cloud-computing environments. This Action will discuss and establish broad theoretical underpinnings, and develop and showcase algorithms, tools and applications for Transactional Memories, defining a commonly agreed research agenda and paving the way for related research projects at national and European levels.

B.2 Current state of knowledge

Since the seminal paper on Transactional Memory from Herlihy and Moss, and the later proposal of a software realization of the same idea by Shavit and Touitou, the research on Software Transactional Memory remained mostly dormant until 2003, when the advent of the first generation of multi-core processors spurred again the interest in the area.

The numerous alternative TM solutions since then proposed can be coarsely classified depending on whether they are implemented fully in software (STM), in hardware (HTM), or rely on hybrid schemes that synergically leverage on both hardware and software mechanisms (HySTM).
Having grown out of research into non blocking concurrent data structures, the first proposed STMs provided strong progress guarantees such as lock-freedom. In a lock-free STM, each transaction is required to complete after a finite number of steps system-wide have been executed. This guarantee of system-wide progress is usually satisfied by requiring a process that experiences contention to help the conflicting transaction to complete before continuing its work. This is a very different approach from that taken by lock-based approaches, where transactions are forced to spin or block until the contending transactions have completed. Lock-free solutions are typically quite complex and impose considerable performance penalties. A second generation of STMs ensured the weaker progress guarantee of obstruction-freedom, guaranteeing progress only in case of a thread that is running in isolation for long enough and relying on an out of band contention manager to probabilistically avoid the occurrence of livelocks in the presence of contention. As a key benefit, obstruction-free STMs are easier to implement efficiently. A third generation of STMs adopts a lock-based scheme that, while lacking theoretical progress guarantees, was shown to achieve better performance in a wide range of realistic workloads, avoiding extra indirections for data accesses and allowing for improved cache locality and faster algorithms.

On the performance side, even though over time STMs have become increasingly competitive, they still incur in two main sources of overhead when compared with classic lock-based synchronization schemes: conflict detection, which requires keeping track of transactions' read-sets and write-sets, and version management, which requires to store at least the new and old values of the memory items updated by a transaction.

By implementing these mechanisms directly at the hardware level, HTMs permit to avoid their costs, achieving significant performance boosts with respect to STMs. The first TM proposal, presented by Herlihy and Moss in 1993 was, in fact, a purely hardware implemented solution. This approach, however, imposed quite stringent programming constraints, supporting only "bounded" transactions that do not access more than a fixed, architecture-specific number of cache lines. Since that date, a number of "unbounded" HTMs have been designed, seeking to minimize both their overhead and their intrusiveness with respect to the pre-existing processors' architecture. The main issue with TMs fully implemented in hardware is that they impose complex, risky, alterations of the processors' organization. This is likely to discourage their inclusion in commercial processors in the near future, at least before the TM-based programming model builds a critical mass of users in the software industry.

An emerging, alternative solution consists of hybrid approaches that rely on best-effort bounded HTM supports to accelerate transactions' execution whether possible, otherwise falling back to software implemented mechanisms. This added flexibility considerably simplifies the task of integrating HTM into a commercial processor, because the processor can respond to difficult conditions by simply aborting the transaction. Further, it predisposes applications to take advantage of future improvements of the HTM supports.

A considerable body of research has also been performed on the theoretical aspects of TMs, including the formalization of correctness criteria capturing the unique features of the TM programming model, the establishment of optimality and complexity bounds results, and the formal verification via mechanical tools of TM algorithms and of their implementations.

Finally, a recent evolution of the research area on TMs is represented by Distributed STMs (DSTMs): DSTMs enrich the traditional TM model to breach the boundaries of a single machine and transparently leverage the resources of commodity, shared-nothing clusters and achieve higher scalability and dependability levels. This is a very recent research topic, which represents, in some sense, the confluence of the research areas on TM, distributed shared memory (DSM) and database replication. Interestingly, the few currently available DSTMs have shown promising, preliminary results, highlighting how the reliance on the atomic transaction abstraction allows avoiding the well-known performance limitations of classical DSM systems, while providing strong consistency guarantees and scalability up to hundreds of nodes. These features, combined with the simple and familiar interface of DSM systems, make DSTMs an attractive candidate to become the reference programming paradigm for large scale cloud computing platforms, whose popularity has been growing at an incredibly rapid pace in recent years.

B.3 Reasons for the Action

The past five years have seen an explosion of activities in TM research over a range of cross-disciplinary areas including processors' hardware design, operating systems, compilers and debuggers, theory of concurrency and of distributed systems. However, much of today's research seems hectic and uncoordinated, leading to overlap or even redundancy, and to progress in small increments. In order to push the research front faster, it is essential to pursue a better coordination of research. On the other hand, to ensure that the research efforts are actually aimed to meet the requirements of ICT practitioners it is first of all crucial to foster the adoption of TMs in the software industry via effective showcase activities. This represents the first step necessary to create a virtuous circle that guarantees not only the timely dissemination of the research results, but also the prompt gathering of precious feedbacks from real-world TM based applications.

The COST Action is the most appropriate means to implement this strategy. It ensures the right level of flexibility to its participants, thus preserving the necessary level of freedom in individual research activities, while offering sufficient means for the constitution of a coordinated "community" out of the set of research groups involved in TM. The Action allows reaching the critical mass that is necessary to be recognized as the primary reference point for European research in the field of TM (note that as many as 12 countries have expressed interest in this Action, with 37 institutions involved).

The Action does not aim exclusively at pursuing scientific/technological advance in the TM's research field. It is also aimed at enhancing competitiveness of European ICT software industry, which is largely composed of SMEs. In fact, the knowledge transfer enacted by this Action will facilitate the shift of European software industry towards the TM-based programming model, which has been shown capable of sensibly boosting the productivity of the development process of concurrent applications. From a strategic perspective, gaining technical knowledge on the development and usage of TMs is of paramount importance for the European ICT society to respond to the challenges (and take advantage of the opportunities) created by the adoption of today's multicore architectures and of tomorrow's massively-parallel systems.

B.4 Complementarity with other research programmes

In Europe, research in the area of TMs has been done in a scattered manner, both geographically and scientifically. The only exceptions are the 7th FP IP projects VELOX and TERAFLUX, and the Marie Curie ITN TransForm, which represent, to the best of our knowledge, the only EU funded projects coping (at different extents) with TMs. Up to date, however, there has not been any initiative aimed at building a pan-European network for coordinating the research activities in such a relevant research area.

Other related European research programmes are:

FP7 Apple-CORE: "Apple-CORE will develop compilers, operating systems and execution platforms to support and evaluate a novel architecture paradigm that can exploit many-core chips to the end of silicon. It adopts a systematic model of concurrency implemented as instructions in the processors' ISA (developed in the EU FP6 AETHER project). This has enormous potential but is disruptive, as this paradigm shift requires a new infrastructure of tools."
Coordination activity: Explore shared interests in the area of multi-core programming and concurrency theory.

FP7 Jeopard: "The strategic objective of the (Jeopard) project is to provide the tools for platform-independent development of predictable systems that make use of SMP multicore platforms. These tools will enhance software productivity and reusability by extending processor technology already established on desktop systems for the specific needs of multicore embedded systems"
Coordination activity: Explore the issues associated with the integration of TM technology in real-time systems (e.g. embedded systems).

COST Action IC0805: "The main objective of the Action is to develop an integrated approach for tackling the challenges associated with heterogeneous and hierarchical systems for High Performance Computing (HPC)."
Coordination activity: Jointly explore the possibility of, and the issues associated with, integrating the TM abstraction in HPC oriented data parallel and message passing programming paradigm such as OpenMP and MPI.

FP7 HIPEAC Network of Excellence: "HiPEAC's mission is to steer and increase the European research in the area of high-performance and embedded computing systems."
Coordination Activity: As for the FP7 Jeopard and the COST Action IC0805 programmes, the Euro-TM Action will coordinate with HiPEAC to address the issues related to the integration of TM tools and libraries within programming frameworks for the high-performance and embedded computing domains.

C. OBJECTIVES AND BENEFITS

C.1 Main/primary objectives

The main objective of this Action is to enhance the quality, visibility and impact of European research in the area of TMs by establishing a network of ICT researchers and developers to address the TM theoretical underpinnings and to facilitate the development and showcase of TM platforms and applications. This will allow coordinating and shaping the research efforts in this new and rapidly evolving research field.

C.2 Secondary objectives

The networking activities enacted by the Action will also:

Facilitate and initiate the development of novel TM algorithms, platforms, and verification tools by:

- continuously tracking the activities in the field, and identifying the trends and needs
- fostering standardization procedures in the field
- defining standard benchmarks and evaluation methodologies for the assessment of TM systems

Incentivise knowledge transfer to the scientific community and practitioners by:

- promoting dissemination of results through joint, multidisciplinary, and multinational high-quality scientific publications
- developing showcases that demonstrate the value of research results in the TM area to ICT industry
- increasing the prominence of the research area for both practitioners and the general public by targeting publications in the mainstream media

Create an interdisciplinary European research networking on TM that achieves sustainability beyond the Action by:

- enlarging the network of involved experts beyond the individuals and countries originally involved in the preparation of this Action
- establishing sustainable secondary networks of experts and organizations that can persist beyond the lifetime of the Action
- educating early-stage researchers through Training Schools and Young Researchers Forums
- facilitating visits of young researchers to foreign laboratories

C.3 How will the objectives be achieved?

The following, concrete instruments will be employed to pursue the Action's objectives:

- Setup a portal and associated electronic collaborations tools (thematic discussion forums and mailing lists) targeting also external audience;
- Organize regular coordination meetings (at least two per year) giving particular emphasis to technical discussions and presentations; representatives from EU projects and externals experts will be regularly invited to participate to the technical discussion;
- Organize a series of yearly workshops, inviting key speakers and lecturers from both academy and industry;
- Release a series of Annual Reports (total of four) extensively surveying the state-of-art on TM and reporting case studies on the adoption of TM technology in real-world, complex applications;
- Organize summer schools (total of three) and seminars for young researchers.
- Organize short visits for senior participants (<1 month), and long-term visits for young researchers (1-3 months) to be allocated on a competitive basis.

C.4 Benefits of the Action

The main foreseeable benefits of the Action can be summarized as follows:

- Reduced fragmentation of research through the definition of a common research agenda on TM;
- Increased quality and level of inter-disciplinarity of TM research;
- Improved communication and know-how sharing between researchers and practitioners by stimulating innovation transfer and real-world application of the research results, and facilitating the gathering of feedbacks from field deployment toward research community;

- Enhanced competitivity of European software industry by fostering the adoption of the TM-based programming model, capable of sensibly boosting the productivity of the development process of complex, concurrent applications;
- Incentivising and supporting trans-national collaborations and exchanges, especially involving young researchers;
- Achieving a common framework of tools and benchmarks, and foster standardization efforts for what concerns integration of TM APIs across the most popular programming languages and development frameworks (Java, C++, C#, OpenMP etc).

C.5 Target groups/end users

The target users of the scientific results developed within this Action belong to two main categories. A first category comprises ICT researchers, who will benefit both from the interdisciplinary collaboration with other researchers working on complementary aspects of TM systems, and from the exposure to the real-world data, benchmarking applications, and requirements provided by the industrial partners participating in the Action.

The second category of target users will consist, as already discussed, by ICT industry practitioners (ranging from software houses to hardware manufacturers) who will benefit from having prompt access to novel algorithms, methodologies, and tools developed by the researchers involved in the Action.

D. SCIENTIFIC PROGRAMME

D.1 Scientific focus

The experts who will participate in this Action represent a mix of academic and industrial researchers, as well as of ICT professionals, who will implement a research network bringing together research competences as well as practical knowledge in the development of concurrent applications in parallel and distributed systems.

The Action will focus on four thematic areas:

1. Theoretical foundations and algorithms for TM systems.
2. Hardware level supports for TMs.
3. Integration of the TM model within existing, or emerging, programming languages.
4. Definition of best practices, methodologies, and tools for the development, verification, and performance evaluation of TM-based applications.

Tasks Addressed:

1. Establish the state of the art: As previously highlighted, TM has become a hot research topic only recently, with the advent of multi-core processors, and even more recent is the investigation of the TM model in the context of distributed systems. On the other hand, some of its theoretical roots date back several decades, e.g. to the work carried out in the database area on serialization theory and concurrency control or in the distributed systems community on replication and distributed shared memories. In the current explosion of research activities, work is taking place in an isolated fashion in many groups worldwide. Hence, it will be an important task of the Action to establish the state of the art in the four thematic areas addressed, trying to organise and structure the available body of knowledge, critically review existing approaches, programming practices, and benchmarking methodologies.
2. Define application requirements and the research agenda: the network of this Action unites academic researchers with ICT practitioners, enabling the definition of application requirements and formulation of a research agenda with a solid rooting in application needs.
3. Provide research tools for the scientific community: this Action will provide research tools that are useful to the scientific community, such as an online bibliography, dedicated mailing lists and on-line forums serving as reference points for focused discussions and exchange of ideas, benchmarks and evaluation methodologies.

4. Develop showcases for researchers and European ICT companies: the showcases will serve not only to foster cross-fertilisation among different scientific communities and to enable better coordination of the European research activities on TMs. They will also permit to demonstrate the usefulness of the developed scientific results to the ICT practitioners, and to allow them to directly influence the direction of future research efforts by providing researchers with feedback on the adoption of research results in realistic settings.

Human resources will be required to undertake duties in the following areas:

- scientific (network management, project management);
- administrative (scientific secretary, Webmaster, organization of meetings and exchanges);
- editorial (analysis of results, drafting and distribution of recommendations);

Funding will also be needed to cover:

- scientific meetings and events;
- training schools and young researchers forums;
- assistance with publications (proofreading, editing, printing);
- invitations to international experts;
- large-scale dissemination of information (workshops, showcases, academic and non-academic publications).

D.2 Scientific work plan methods and means

The Action will be organized according to the following Working Groups:

- WG1: Cross-WG Activities, Showcases
- WG2: Theoretical foundations & Algorithms
- WG3: Hardware's & Operating System's Support
- WG4: Language Integration & Tools
- WG5: Applications & Performance Evaluation

WG1: Cross-WG Activities, Showcases

WG1 plays a special role within this Action by bundling cross-WG activities in order to ensure their cohesion and boost interdisciplinary collaborations. Specifically, this WG will be in charge of the integration of the remaining four WGs so that researchers in these different fields can share their common expertise, favouring cross-fertilization and minimizing fragmentation of efforts. New advances in any of the WGs will be communicated to the others WGs in order to avoid duplication of research and efforts; moreover, exchanges between the WGs will be promoted.

Furthermore, this WG will oversee tasks related to dissemination towards end users, including the design and development of the showcases of the Action.

WG2: Theoretical foundations & Algorithms

Unlike database transactions, in a TM transactions do not execute in a sand-boxed environment. This exposes TM transactions to hazardous execution scenarios that would not be filtered out by the safety and liveness properties normally ensured in a DBMS (Database Management System) environment.

On the safety side, for instance, if, in a TM, a transaction observes an inconsistent system it may generate unexpected memory access violations or arithmetic exceptions, causing the failure of the whole application, or get trapped in infinite loops, never reaching the commit phase. This kind of phenomena is absent in a database system where transactions observing stale state can be eventually aborted without harmful side effects (except of course from a performance perspective). Further, in a TM, application's state can be concurrently accessed by both transactional and non-transactional code, a scenario clearly not modeled by classical database serialization theory. On the progress side, lock-based concurrency control schemes are extremely popular in DBMSs, where deadlocks can be lazily detected incurring simply in a performance penalty; in a TM system, conversely, lock-based solutions are often deemed unacceptable as the consequences of deadlocks (e.g. involving interrupt handlers) may lead to a halt of the whole system.

This WG will carry out foundational research aimed, firstly, at precisely formalizing the correctness criteria that TM implementations should provide. Secondly, based on such formalization, the WG will work on the design, verification and performance analysis of TM algorithms, considering both the scenarios of deployment on a single cache coherent shared memory system and over a set of distributed, shared-nothing nodes.

The performance analysis carried out within this WG will be of a twofold nature. From a theoretical perspective, they will aim at establishing fundamental complexity and optimality results, shedding light on the inherent costs of the various trade-offs in the TM's design space. These theoretical results will be complemented by performance evaluation studies of software based implementations using, whenever possible, the benchmarking applications that will be produced by WG5.

WG3: Hardware's & Operating System's Supports

Albeit TMs can be purely implemented in software, there is a growing consensus that some form of hardware support is desirable to improve performance. By providing hardware level support for conflict detection and version management, in fact, the bookkeeping overhead incurred by STMs can be drastically reduced, making TMs' performance superior even to that of hand-crafted, fine grain, locking. On the other hand, the implementation of these mechanisms in hardware is way more problematic than in software, as this can entail invasive, risky modifications of crucial components of existing processors such as cache, TLB (Translation Look-aside Buffer) and bus protocol. An alternative solution, that appears more viable in the near future, is to rely on hybrid approaches that require simpler (and therefore less disruptive) best-effort hardware mechanisms, capable of boosting transactions in common execution scenarios, and to resort to software implementations only whether needed.

The researchers participating in this WG will explore the design space of hardware supports (hybrid and not) for TMs, and investigate the relations between hardware, operating system and user level applications. This includes: understanding if, or at what extent, modern OSs can take advantage of various TM hardware supports (e.g., for synchronizations within the kernel) to enhance their performance and/or reliability; evaluating pros and cons of the integration of TM APIs at the OS level; devising efficient OS level mechanisms allowing TMs' transactions to encompass interactions with hardware devices and/or system call invocations, one of the major limitations of current TM systems.

WG4: Language Integration & Tools

There has been an increasing number of papers illustrating the potential of transactions in parallel programs, introducing TM implementation techniques, and proposing hardware mechanisms to accelerate common case behavior. Nevertheless, TMs are still not widely used by application developers as the TM technology has not yet been integrated with any of the widely used software development frameworks. It is unlikely that programmers will abandon the many advantages of mature frameworks in order to try this new technology. As a result, there are very few large applications that use transactions and there is no clear assessment on whether TM actually improves programmability. In addition, there is little convergence or sharing of technology between TM researchers. Each group uses a different programming model, code generation, and TM system approach. This leads to incompatible and immature modules that are unlikely to attract significant attention.

The activities carried out within this WG will be aimed at removing this major technology roadblock by focusing on the issues related to the integration of TM supports within mainstream languages and development frameworks. This encompasses the development of a set of extensions to compilers (e.g. the GNU gcc compiler for C/C++ languages), virtual machines for languages relying on managed runtime systems (e.g. the Jikes RVM for the Java language), as well as companion tools for assisting the development process such as debuggers and profilers.
This WG will also endeavor to boost the international standardization process of a basic set of APIs to be provided at the language level to support execution of TM-based applications.

WG5: Applications & Performance Evaluation

TMs appear to have a huge potential in simplifying the development of parallel applications. On the other hand, being the research field on TMs still in its infancy, the number of complex benchmarks for TMs currently available is still very limited, and the real-world applications pioneering the adoption of TMs are probably even less. This represents a major impairment not only for realistically evaluating the performance of the various TM solutions proposed in literature, but also for precisely assessing the usability of TMs in complex, large scale applications and across the many different application domains that could potentially benefit from their adoption (including web-based applications, video-games, CAD systems, stream processors, financial, healthcare, video-editing, navigation systems, simulators, graph analysis toolkits, just to mention a few).

Leveraging on the research results developed by WG2, WG3 and WG4, this WG will unite researchers and ICT practitioners working on the development of TMs and/or TM-based applications. This will permit, on one hand, to ensure timely interactions among developers of TM-based applications with the purpose of establishing and disseminate best-practices; on the other hand, it will create bi-directional communication channels between developers of TM-based applications and developers of TMs, to assist the former ones in the development and tuning of applications and provide the latter ones with valuable feedback from real-world settings. Overall, this will bring the twofold benefit of fostering the adoption of TMs in the software industry, and of driving future research efforts to actually meet the requirements of real-world applications.

An important research goal of this WG is also to establish automatized methods (based on analytical models/simulations or both) assisting the application's developers in the choice of the TM implementation maximizing the performance of their own application, forecasting the scale-up achievable by acquiring additional or more powerful hardware (e.g. in a cloud-computing settings), and identifying the performance bottlenecks both at the application level and TM level.

E. ORGANISATION

E.1 Coordination and organisation

This Action will focus on unifying research efforts in the area of Transactional Memories, in order to drive them in a common direction and overcome fragmentation. Acting as a coordination effort across the research activities of the existing groups, the Action will incentivise synergic international cooperation with the ultimate purpose of accelerating the development of standard TM platforms capable of effectively meeting the actual requirements of real-world systems.
Following the COST Guidelines, the Management Committee (MC) will be in charge of the overall supervision of the Action, ensuring that the scientific objectives of this proposal are achieved. The MC will also endorse the editorial role of organising the publication strategy and reviewing the content of the publications of the Action. To support the MC in the review of publications, the MC may appoint an Editorial Board. Each year, the MC will plan activities for the coming year (workshops, schools...), appointing, where needed, Coordinators in charge of promoting and supervising specific activities. The nominating COST countries are encouraged to nominate young researchers as MC members.

The MC will nominate a special task force to setup and maintain the Euro-TM portal and the other electronic collaboration tools. In this meeting, the MC will also nominate the Coordinator of each Working Group (WG).

Each Working Group will nominate a WG Steering Board (WGSB) which will coordinate the WG specific activities, under the supervision of the WG Coordinator appointed by the MC. The WGSB will organize the technical meetings that will normally take place after each MC meeting.

Half of the MC meetings will contain a Workshop, possibly connected to a suitable conference or held jointly with another event. All Workshops will be open to the general scientific/engineering public.

A Training School will be organised every year starting from Year 2, conditioned to budget availability for such initiative. Target audience are PhD students and young researchers willing to explore issues related to the usage and development of TM platforms and applications during their research activities.

The Euro-TM Website will be set up and continuously maintained throughout the lifetime of the Action in order to publish results, innovations and applications generated by members of the Action's network. It is intended to enhance the visibility of the Action, optimising and accelerating dissemination of results to researchers, practitioners and public funding agencies. While most of the information will be made available publicly, a password-protected section will allow COST Action network members to share confidential (e.g. non-final) and internal information.

The Action's Website will be launched within 6 months since the start of the Action and will be regularly updated. In order to increase traffic flows from other relevant websites, the EuroTM Website will be registered with appropriate portals. Also, statistics will be maintained to monitor the actual usage of the Website.

The purpose of the unrestricted access section will be to give a maximum visibility to the project and reach as wide an audience as possible. It will host regularly updated information, documents and deliverables in a variety of media, plus links to other relevant networks, sites and programs and a list of participants.

In the password-protected section, there will be a collaborative work platform (e.g. a wiki), giving WG members an interactive space where they can maintain databases with address details of relevant researchers; communicate within and between WGs; keep track of the progress of the Action; communicate and review internal reports, meeting minutes, and drafts of publications; use tools for WG coordination (e.g. meeting scheduling and calendars).

E.2 Working Groups

As detailed in Section D.2, the Action will be structured into the following WGs:

- WG1: Cross-WG Activities, Showcases
- WG2: Theoretical foundations & Algorithms
- WG3: Hardware's & Operating System's Supports
- WG4: Language Integration & Tools
- WG5: Applications & Performance Evaluation

Each WG will be managed by a WG Steering Board under the supervision of a WG Coordinator, to be appointed by the MC. The WGs will coordinate the research work of their members that is funded and carried out at the national level. They will also prepare the WG meetings, devoted to presenting and reviewing research findings among their members, as well as planning future directions.

The first meeting will be devoted to the presentation, by each member, of the research interests that fit in this Action. Therefore, after the first meeting, synergies between each participant will be defined. The goal of further meetings will be to provide a forum for exchange and develop new ideas on the subject, serving as cornerstones for implementing national and international research coordination.

Thanks to these Working Group meetings, bilateral common research topics will be identified. Collaborations between research groups will be supported and incentivised through Short-Term Scientific Missions. Lasting normally several weeks, Short-Term Missions represent indeed an ideal framework to carry out joint research efforts on specific topics.

As a general guideline for the Action, the structure of the WGs will be as flexible as possible, in order to enable other countries and partners to join the Action and therefore extend its visibility and impact.

E.3 Liaison and interaction with other research programmes

Liaison is primarily directed towards other COST Actions and FP7 programmes, as listed in Section B.4. Concerning FP7 programmes, the MC carries a special responsibility to monitor and report relevant activities, using for instance the restricted-access part of the Website. Regarding the complementary COST Actions (see Section B.4) invitations will be forwarded to the chairs of the relevant Actions in order to make possible their representation at the kick-off meeting, thus ensuring coordination at an early stage. During the life of the Action the MC will ensure that the relevant WG is in contact with, and invites experts from, complementary Actions and FP7 projects. It is noteworthy to highlight that the List of Experts participating to the Action already includes representatives and institutions involved in most of the referenced initiatives. This would facilitate the communication and coordination. At the same time, it represents an indication that the Action does have the critical mass to play a central role in the research community.

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 foster involvement of early-stage researchers. This item will also be placed as a standard item on all MC agendas.

Early-stage researchers will take an active role in the management and organisation of the Action itself, e.g. leading task-forces, taking on coordination responsibility. Young scientists will be solicited to act as MC members or otherwise participate to the MC meetings as auditors. The Action will favour budget allocation for STSMs of young researchers (namely PhD students and postdoctoral scientists and, where possible, undergraduates). In all cases, encouragement of applications from all sectors of the community (including minorities) will be clearly stated. There will be no discrimination based on sex, race, religion or citizenship. Conversely, gender balance will be a criterion in acceptance of STSM applications.

F. TIMETABLE

The Action will last four years. A minimum of two MC Meetings, two WG Meetings, and one Workshop will take place every year. To minimise travel and maximize the chances of cross-fertilization and synergies, the meetings and the Workshops will normally be synchronised and held jointly with relevant conferences or other events. The planning of Short-term Scientific Missions (STSMs) will start after the first Working Group meetings, once participants will have identified common research interest.

One Training School will be organized every year starting from Year 2. Over the last two years, two showcases will be organized aimed at reporting on the experiences gained from the adoption of the research results achieved throughout the Action, placing particular emphasis on presenting case studies based on complex, realistic applications.

The Final Conference will be organised over the course of Year 4 and held at the very end of the Action. The final report (or edited book) will be generated during the final year.

The planned schedule of activities, and corresponding milestones, of the Action is reported in Table 1.

Month	What	Notes and Milestone
0.	Kick-off Meeting	Elect WG Coordinators, decide on the upcoming activities, nominate coordinators for these activities.
6	MC&WG Meeting	Presentation of Euro-TM Portal, Opening of the Short-term Scientific Missions call
12	MC&WG Meeting	Delivery of 1st Annual Report
18	MC&WG Meeting	
from 18 to 24	Training School	
24	MC&WG Meeting + Workshop	Delivery of 2nd Annual Report
30	MC&WG Meeting	
from 30 to 36	Training School + Showcase	
36	MC&WG Meeting	Delivery of 3rd Annual Report
42	MC&WG Meeting	
from 42 to 48	Training School + Showcase	
48	MC&WG Meeting + Final Conference	Delivery of Final Report

Table 1. Schedule of Activities of the Action.

G ECONOMIC DIMENSION

The following COST countries have actively participated in the preparation of the Action or otherwise indicated their interest: Denmark, France, Germany, Greece, Ireland, Israel, Italy, Portugal, Spain, Sweden, Switzerland, 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 target audiences for the dissemination of the Action results are:

- The international research community involved in the TM field;
- Industry players involved in the development of concurrent and parallel applications for multi-cores and/or distributed platforms, e.g. video-games, CAD, media editing, financial, health care, cloud-computing applications;
- EU projects and initiatives in the area of parallel and distributed programming;
- Standardization bodies, particularly those active in the area of programming languages' specification (e.g. ANSI/ISO for C++, or the Java Community Process for Java);
- Ultimately, everyday users of today's desktop computers, who will benefit from the availability of robust applications capable of fully exploiting the potentialities of modern multi-core processors.

H.2 What?

The dissemination of the results achieved by this Action will be based on the following activities:

- Euro-TM portal, thematic mailing list and discussion forums;
- Research reports and peer-reviewed articles authored by members of the Action;
- Participation to forums and booths at major events (conferences, symposia, etc.);
- Series of Workshops (one per year), with published proceedings;
- On-line bibliography;
- Working Group reports;
- Contribution to standardization bodies;
- Press statements.

H.3 How?

WebSite: The Euro-TM Portal will be a key dissemination tool for the Action. Besides providing a centralized collaboration platform for the Action's participants (e.g. for document sharing), the portal aims at becoming the natural online reference for experts in the field, from both academy and industry. The portal will host all the documents output by the Action, from scientific articles to technical reports, including presentations and meetings' minutes. It will also serve as an interactive Web 2.0 platform hosting discussion forums, comments on published articles, thematic mailing lists and a wiki.

A dedicated task force will be appointed at the kick-off MC meeting to set up and maintain the TM-portal, assuring it is regularly updated. The Action will monitor regularly the popularity and usage level of the various sections of the portal through quantitative indicators, which will be revised at each MC meeting.

The Action will foster the active participation to the portal resources by its members and promote the role of the Euro-TM portal as a central hub for other European initiatives in this research area (e.g. FP7 and national projects).

Media Promotion: The Action, and in particular WG1, will produce press releases detailing newsworthy findings for distribution to the media.

Showcases: Showcases highlighting the benefits of TM technology in tackling practical problems will be organized, so to demonstrate the results achieved by the Action to industrial practitioners and decision makers. WG1 will be responsible for to coordinate the organization of these showcases. Showcases will be disseminated through every possible channels, including the Website, Workshops, as well as publications presenting case studies on the use of TM technology in realistic applications.

Workshops: The Action will contribute to the organization of workshops jointly targeting researchers and practitioners pioneering the adoption of TM technology. These workshops will be open also to scientists who are not members of the Action and will invite external experts, particularly scientists from other disciplines and representatives from software industry, thus serving as a dissemination channel towards them. Workshops will generate reports that will be published on the Euro-TM portal. Several experts involved in this Action are involved in the editorial boards of well-known journals in the field as well as organizers of scientific events. This represents a guarantee of the dissemination of the Action activities on a large scale.

Publications (Flyers): Flyers summarizing the main facts about the Action will be made available at conferences and other relevant public events. This will advertise the Action's activities to a wide audience also through non-electronic channels.

Publications (Non-academic): To enhance visibility of research results among non-academic researchers and practitioners, whenever possible, articles will be published in non-academic outlets such as newspapers or trade magazines (possibly also in the electronic media).

Publications (Academic): Research results will be published in international scientific journals and high-quality peer-reviewed conferences listed in relevant literature databases (ISI, DBLP etc.), in order to maximise visibility within the scientific community.

Online Bibliography: An annotated bibliography collecting the relevant literature generated worldwide will be compiled and maintained. This bibliography will be made available through the Website.

Final Conference: The Final Conference aims at creating a broad forum where ICT professionals from relevant application domains meet with the researchers participating in the Action, as well as to disseminate results to researchers of related fields. A summary report will be made available on the Website.

Final Report: This report will summarise the key results and findings of the Action, illustrating showcases and highlighting success stories. The MC plans to sign a contract with an international publisher in order to ensure wider circulation of this Final Report as an edited book.