Intermediate Report

Project ANR/JST- 2010-JTIC-003

Framework and programming for Post Petascale Computing

FP3C

This report was compiled for the ANR-JST Joint workshop which was held in Kobe, Japan in March 2012. This report summarised FP3C group research activity of first half research period.
本報告書は、2012 年 3 月に神戸で開催された領域内中間ワークショップのために編纂されました。研究グループの研究期間前半の活動内容をまとめたものです。

A IDENTIFICATION

Project acronym	FP3C
Project title	Framework and Programming for Post Petascale Computing
Project coordinator (French side) (company/organization)	Serge Petiton, INRIA
Project coordinator (Japanese side) (company/organization)	Mitsuhisa Sato, University of Tsukuba
Project start date Project end date	September 1, 2010 August 31, 2013, We propose to **postpone the project end date to March 2014, cf. section C.4**
Project website if applicable	jfli.nii.ac.jp/medias/wordpress/?page_id=327

B DELIVERABLES AND MILESTONES

The deliverables and milestones are approximately seven month postponed (cf. section C.4 for explanations). A new virtual t0 is set such as t0'=t0'+7months. Then, t0' is set on April 1st, 2011.The milestones associated with the "first" year tasks were not explicitly concretized, as it was impossible to keep a strict timing between the teams during this period. We focus on the deliverables, which we sometime aggregated to larger ones. The milestones and deliverable scheduled for the "second" and "third" year are unchanged. Of course, the FP3C project begin on September 1st, 2010; we introduce the t0' to be able to keep the proposed schedule for the last two years of the project. Without future unexpected problems, we plan to complete the original project by the end of March 2014. Even if many people work worldwide in the HPC domains, all the proposed works are still at the forefront of the research on the domain and would be new results when they will deliver.

Deliverables of the "first" year (from September 1st, 2010, to March 31, 2012):

No.	Designation	Nature*	Date of supply			Partners (underline the responsible partners)
			Initially planned	Re-schedule d	Delivered	
TASK1	YML/XMP prototype, baptized "FP2C" for "Framework for Post-Petascale Computing".	Software	t0+12m	t0'+12m	March 2012	INRIA Saclay, TSUKUBA, CNRS PRISM
TASK2	New communication library for post-peta scale computing, PNewMadeline	Software	t0+12m	t0'+16m	July 2012	INRIA, U. of Tokyo, TITECH
	Communication and fault-aware middleware ¬ Extended XcalableMP for faultaware computing	Software	t0+12m	t0'+16m	July 2012	U. of Tokyo, TITECH, INRIA
	HGMDS/BlobSeer-WAN prototype	Software	t0+12m	t0'+16m	July 2012	Tsukuba INRIA
TAKS3	Performance-model driven auto-tuning framework for GPUs	Software	t0+12m	t0'+12m	t0'+12m	TITECH
	Fault tolerance algorithms and software for GPUs	Software	t0+12m	t0'+12m	t0'+12m	INRIA, TITECH
	XMP-dev: a prototype of Extended XcalableMP for Acceleration devices	Software	t0+12m	t0'+12m	t0'+12m	Tsukuba,
	XMP/StarPU", a prototype compiler for CPU/GPU hybrid workload	Software	t0+12m	t0'+12m	t0'+12m	Tsukuba INRIA

TASK4	MERAM with YML/XMP benchmark,	Software	t0+12m	t0'+12m	March 2012	PRiSM
	SS eigensolver with YML/XMP benchmark	Software	t0+12m	t0'+16m	July 2012	Tsukuba
	Use of MUMPS within SS eigensolver and XMP+MUMPS interface,	Software	t0+12m	t0'+16m	July 2012	Tsukuba , IRIT AICS
	Use of the GRID-TLSE web site to exchange/validate sparse matrices,	Data	t0+12m	t0'+12m	March 2012	IRIT and all the partners
	Preparation of the Block Cimmino code for benchmarking.	Software	t0+12m	t0'+16m	July 2012	IRIT
	Experiments on parallel algorithms for constrained based combinatorial optimization.	Software	t0+12m	t0'+16m	July 2016	Tokyo Univ.
	Parallel ILU(k) preconditioner for ill conditioned problems/scalable multigrid preconditioning + finite Element procedure on GPU + auto-tuning of GMRES	Software	t0+12m	t0'+16m	July 2012	Tokyo Univ. INRIA Saclay
	Framework for the development of scientific applications	Sotware	t0+12m	t0'+16m	July 2012	Kyoto Univ.

Scheduled deliverables and milestones of the "second" and "third" years of the FP3C project (from the accepted proposal). The dates are 7 months postponed from the initial schedule.

MODELS AND APIs FOR RUNTIME			
Task5	Task 5.1: - A one-sided communication API and implementation suitable for use by the XcalableMP runtime.	~~t0+24~~, t0'+24 March 2013	INRIA Bordeaux
	Task 5.2: - An YML engine adapted to use XMP encapsulated program (T0+18) - YML-XMP backend prototype for the runtime system to be proposed in Task 7. - Report on the end-user expertise which may be considered, synthesized and exploited through the different level of the runtime systems studied into the FP3C project.	~~t0+24~~, t0'+24 March 2013	INRIA Saclay
	Task 5.3: - A new storage approach for parallel execution checkpoints from non-volatile local storage devices (SSD, Phase Change Memory).	~~t0+24~~, t0'+24 March 2013	
	Task 5.4: - A prototype implementation for the storage architecture based on Gfarm and BlobSeer.	~~t0+24~~, t0'+24 March 2013	INRIA Rennes
MODELS AND APIs FOR ACCELERATORS			
Task 6	Task 6.1: - A prototype implementation of the port of XcalableMP over StarPU and scheduling heuristics adapted to it.	~~t0 + 24~~, t0'+24 March 2013	INRIA Bordeaux
BENCHMARK AND EVALUATION			
	Task 7.1: - Report on optimized implementation of numerical kernels on target architectures (GPU, muti-core architectures …) - Report on first implementation of some of numerical kernels using XMP on the target architectures.	~~t0+24~~, t0'+24 March 2013	CEA Saclay

	Task 7.2: - A few Codes of hybrid methods in YML/XMP, accompanied by a report on the study and the analyze of the realization. - Report on the first evaluation of hybrid methods on the runtime systems proposed on 2.	~~t0+24~~, t0'+24 March 2013	CNRS-PRISM
	Task 7.3: - Implementation of the algorithms defined in task 4.3, preliminary performance results	~~t0+24~~, t0'+24 March 2013	University of Tsukuba
INTEGRATION AND EVALUATION			
Task 8	Task 8.1: • A deployment tool to use XMP and YML for the different developed runtimes systems, including compliers, engine and component catalogues.	~~t0+36~~, t0'+36 March 2014	University of Tsukuba
	Task 8.2: • A evaluation of the BLOB-based file storage prototype on Aladdin-Grid'5000 and on the InTrigger platform.	~~t0+36~~, t0'+36 Macrh 2014	INRIA Rennes
LIBRARY PACKAGING, BENCHARK			
Task 9	Task 9.1 • A few kernels of a basic linear algebra library for Petaflops computers incorporating dense and sparse linear algebra solvers (direct, iterative and hybrid methods) and some eigensolvers.	~~t0+36~~, t0'+36 March 2014	CNRS IRIT
	Task 9.2 • Implementation description and performance report of Numerical Kernel suites using the different programming environment.	~~t0+36~~, t0'+36 March 2014	University of Kyoto
	Task 9.3 • A package containing the codes of proposed linear system solvers, the user manual and report on performances on application examples.	~~t0+36~~, t0+36 March 2014	CEA Saclay
DEMOS AND TRAINING			
Task10	Task 10.2: Final workshop (demos)	~~t0+36~~, t0+36 March 2014	CEA Saclay

C PROGRESS REPORT

C.1 INITIAL OBJECTIVES OF THE PROJECT

High performance computing systems used for cutting-edge of advanced computational science have reached to petaflops (a million billion calculations per second) performance, and will be targeted to the next generation of exascale systems as a post petascale system. Our goal is to contribute to establish software technologies, languages and programming models to explore extreme performance computing beyond petascale computing, on the road to exascale computing. The ability to program these future high performance systems efficiently is considered by all research national agencies all along the world as a strategic and important issue. France and Japan plan to host some of these systems by the end of the FP3C project. Post-petascale systems and future exascale computers are expected to have an ultra large-scale and highly hierarchical architecture with nodes of many-core processors and accelerators. That implies that existing systems, language, programming paradigms and parallel algorithms would have, at best, to be adapted, and often obsolete. To manage these ultra large-scale parallel systems, we require new adaptive runtime systems, allowing to manage huge distributed data, minimizing the energy consumption, and with fault resilient properties.

These criteria should be integrated to propose future systems for post-petascale computing. Moreover, accelerating technology, based actually on GPGPU and many-core processors, is a crucial domain for post petascale computing. They are the lower level of the computing process associated with

multithreading techniques, and their efficient programming in these large scale systems is an important challenge. New scientific methods and algorithms for these systems will have to be introduced. The different programming levels (from clusters of processors loosely connected to tightly connected manycore processors and/or accelerators) will generate new difficult algorithm issues. Language, framework and runtime systems should be defined, experimented and evaluated with respect to modern state-of-the-art scientific methods. Benchmarks and libraries adapted to the Post petascale systems will have to be proposed. Then, many issues related to the emergence of Post petascale computing are all linked together.

The overall structure of the FP3C project represents a vertical stack from a high level language for end users to low level architecture considerations, in addition to more horizontal runtime system researches. Programming models and APIs will be proposed as benchmarks and methods to evaluate each part of the project. The last year, integration of the different software, systems and tools will be proposed. Packaging of some libraries and benchmarks will be developed. Researches will focus on the relationship between different topics and how they may contribute to ameliorate the global post-petascale programming solution, which will be presented at the end of the project with tutorial and demos.

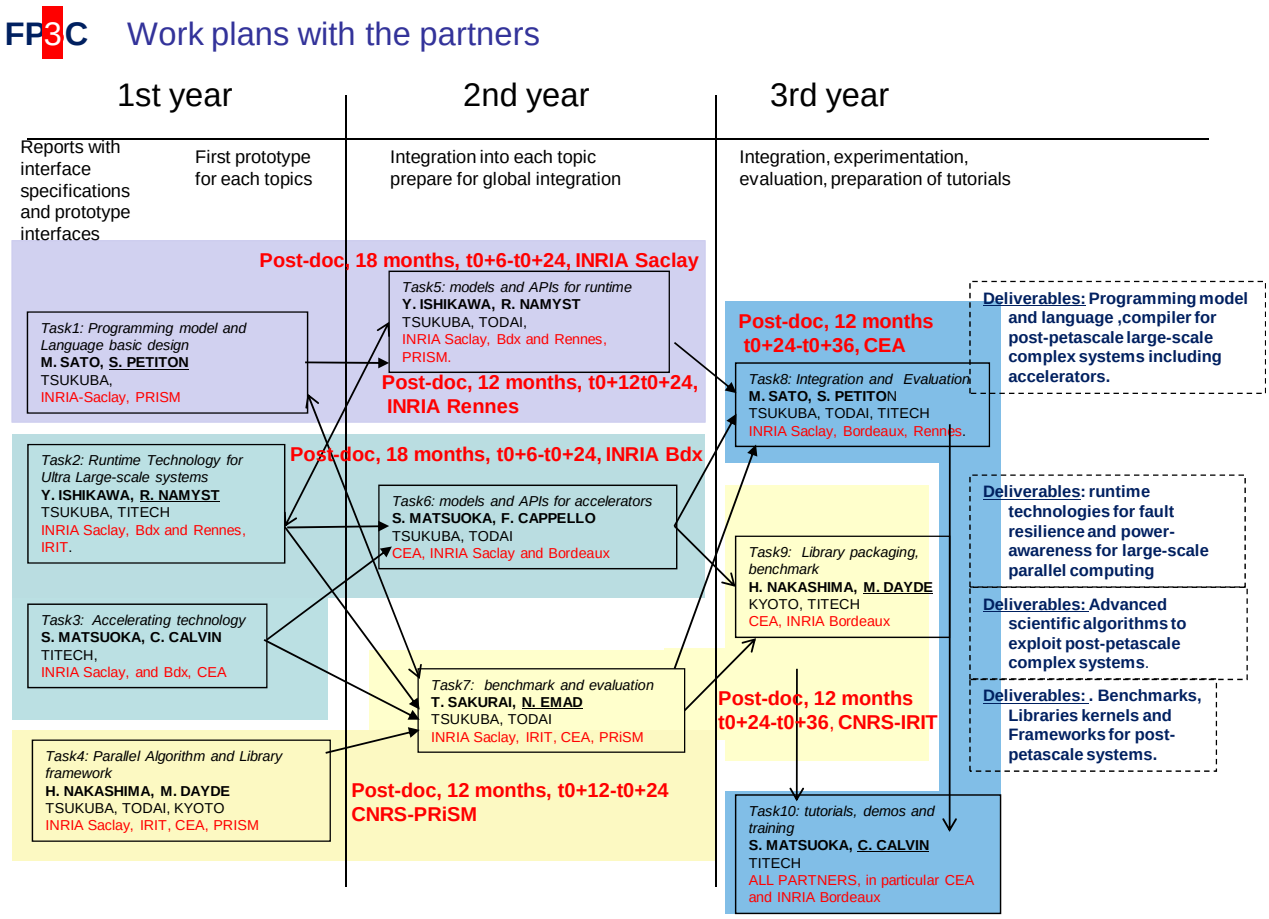

Figure 1 : the FP3C project organization into tracks and task

C.2 WORK PERFORMED AND RESULTS ACHIEVED IN THE FIRST HALF PERIOD

The FP3C project is organized into 4 tracks, corresponding to the main objectives. Each of them is organize into tasks, to specify works and dependencies into the project. During this first year we had 3 tracks distributed into four tasks, cf. Figure 1. The work performed and the results achieved since the beginning of the project represent approximately the third of the FP3C project. We consider a "new" initial schedule (cf. C.4), which is seven months postponed from the initial one.

TASK 1: the programming model and language basic design

After studying XcalableMP (XMP) language, YML has been extended in order to support XMP. New keywords

have been introduced in the implementation component of YML that allow generating automatically the corresponding data distribution in XMP. A new backend have been also developed in order to run XMP component. For this purpose, the already supported OmniRPC middleware have been extended to support MPI and XMP component execution. After some discussion on the running of YML/XMP programs, communication system between components has been also modified in order to get one executable binary. Now, the communication between components is made through the shared file system. Some functions in charge of the communication of distributed data are automatically generated from the data distribution description in YML. After some experiments, YML-XMP software has been released on February 2012. Experimentations on T2K in Japan and on GRID5000 in France have begun recently. These results are in concordance with the "new" initial schedule.

TASK 2: Runtime technology for ultra large scale systems

The first one year research efforts in task 2 are summarized as follows:

i) A novel interprocess optimization method and a network fault resilience mechanism have been designed and prototyped in NewMadeline developed at INRIA. ii) Using NewMadeline, non-blocking collective communication mechanisms have been designed and implemented by TITECH, U. of Tokyo, and INRIA. iii) Fault-tolerance mechanisms have been designed for XMP by U. of Tokyo, TITECH, and INRIA. iv) A scalable and reliable large-scale checkpointing library using SSDs and advanced encoding techniques have been developed by TITECH and INRIA. v) a storage system relying on HDMDS, a distributed metadata management system for global file systems, has been designed and implemented at U. of Tsukuba and INRIA.

TASK 3: Accelerating technology

This task is divided into two subtasks.

The 1st one is dedicated to the design and the implemented of a first version of a new parallel programming language to provide an easy programming environment for large scale heterogeneous parallel processing systems engaging high performance CPU and GPU. The language is named "XMP-dev" which is an extension of XMP (XcalableMP) parallel programming language for large scale distributed memory parallel systems. Although XMP-dev is basically designed to apply to any kind of heterogeneous node with any accelerating device, the first target of accelerating device is GPU. We have implemented the preliminary version of XMP-dev for PC clusters with NVIDIA GPU. Ongoing work concerns the utilization of StarPU environment to apply dynamic job scheduling system as a virtual accelerating device of XMP-dev. This work is undergoing as a collaborative work among University of Tsukuba and INRIA Bordeaux.

The 2nd task concerns the development and the evaluation of different technologies to improve scalability of petascale GPU-accelerated software. The evaluation was done on several testbeds small benchmarks were run on prototype test nodes with GPUs, whereas large runs involving 100s to thousands of GPUs were executed on TSUBAME2.0 supercomputer. The first optimization applied to these kernels is based on communication hiding techniques. Future work will concern the (semi-)automate load balance and tuning between GPU and CPU in cooperation with StarPU runtime team in Bordeaux. Also, we have designed a specific framework KASH (Krylov bAsed Solver for Hybrid architectures) which allow us fast design and implementation of various linear algebra kernels on both homogeneous multi-core architectures and hybrid GPU-CPU machines. Next step will to offer XMP and XMP-dev implementation of KASH framework in order to compare performances with basic MPI/OpenMP and CUDA implementation versus higher programming language environment like XMP and XMP-dev. Some of the kernels whose computational pattern is that of stencil computation was proposed to be handled and optimized with our new DSL (Domain Specific Language) called Physis, to separate description of computation of each point from complexity of GPU programming and details of communication method.

TASK 4:Parallel Algorithm and Library framework

Experiments are reported by PRISM, CEA and INRIA Saclay for ERAM eigensolvers on GPU and show high performance compare to multicore implementation both on dense and sparse matrices.

The Parallel ILU(k) Preconditioner for Ill-Conditioned Problems developed at Tokyo University demonstrates good scalability on the T2K Open Supercomputer. The MUMPS sparse solver from IRIT has been installed on this computer and will be use to accelerate the precondition phase. Experiments on some matrices have started. A multigrid preconditioner from Tokyo University using a OpenMP / MPI hybrid programming shows a good scalability on 3D problems using the T2K computer. A finite element method has been successfully implemented on GPU by Tokyo University. Auto-tuning approaches for GMRES(m) are developed both at Tokyo University and INRIA Saclay / PRISM and let expect good collaborations. Some results from the Japanese group are already available on the T2K.

The work on hybrid iterative restarted methods involves PRISM and TSUKUBA. Mixing the SSM method developed at TSUKUBA and MERAM from PRISM and implementation within YML is currently explored. The use of the MUMPS solver from IRIT within SSM is also considered. The Kyoto Group investigates a framework to describe application codes in the form of local-view kernels in close collaboration with TiTech (Tokyo Institute of Technology) in order to extend their DSL for GPU-oriented stencil computation to make it applicable to wider range of applications and platforms including multi-core/ multi-socket clusters. Some parts of the MUMPS sparse solvers are included in the set of computational kernels that are under consideration. Also, the Block Cimmino method is developed by IRIT for improving its scalabity.

The Tokyo Group from JFLI is investigating the design and implementation of local search algorithms for solving

constraint satisfaction problems on massively parallel computers. Speedups are very good for a few tens of processors, and good up to a few hundreds of processors

The use of GRID-TLSE site for exchanging sparse matrices The Kyoto Group investigates a framework to describe application codes in the form of local-view kernels in close collaboration with TiTech (Tokyo Institute of Technology) in order to extend their DSL for GPU-oriented stencil computation to make it applicable to wider range of applications and platforms including multi-core/ multi-socket clusters. Some parts of the MUMPS sparse solvers are included in the set of computational kernels that are under consideration.

C.3 WORK FORECAST IN THE SECOND HALF PERIOD

The second year, each track has one task. The global idea is

1. to integrate results from task 2 and 3 into the YML/XMP framework to allow XMP to target GPU based platforms (task5) ; and to begin to benchmark the tools developed into task 2 and 3 using classical library and numerical software studied in task 4 (task 7)
2. to continue developing models and APIs for accelerators (task 6)

The last year, a new track will be launch concerning the integration of all the results of the previous tasks (task8) and to prepare tutorials, demos and training for future users (task 10). A third task will integrate the API developed in task 6 and with respects to the results of the benchmarks done in task 7, will package studied library software and will define some stable benchmarks.

TASK 5, track 1: Models and APIs for runtime (April 2012-March 2013)

The main goal of this task is to integrate results on data management and accelerator runtime systems from task 2 into the software developed in task 1, to allow 3-level programming: YML, XMP and accelerators.

The NewMadeleine communication library developed in task 2.1 will be integrated into the XcalableMP runtime system developed in task 1.1. We will extend NewMadeleine in order to fit this communication paradigm, with asynchronous progression of one-sided communication thanks to NewMadeleine multithreaded communication engine.

TASK 6, track 2: Models and APIs for accelerators (April 2012-March 2013)

XMP-dev, proposed in task 3, will utilize StarPU to benefit from dynamic job scheduling on heterogeneous machines. Scheduling heuristics of StarPU will be rethought to improve the performance of the target applications.

We will further improve upon our frameworks to allow us to better describe more complex computational patterns, and to automate the execution on GPUs and possibly other accelerators, through both static and dynamic means. In particular we plan on collaborating with the team at Kyoto University on an extension of our Physis framework for particle-in-cell plasma simulations, which involves more flexible data access patterns than finite difference stencil code, resulting in load inbalance. Another scenario would be compute-bound problems that require both GPU/CPU usage at the same time; there load imbalance will occur naturally due to differences in GPU and CPU performances. To resolve this problem we are planning to (semi-) automate the load balancing and tuning process by cooperation with StarPU runtime team in Bordeaux.

TASK 7, track 3: benchmark and evaluation (April 2012-March 2013)

Languages and model proposed in Task 1 and accelerator programming techniques obtains in Task 3 and 4, during the first year, will be exploited with proposed/developed methods, algorithms and benchmarks for the targeted post-petascale computers. The evaluation of these will be made by making use of clusters of multicore and GPU in France, TSUBAME 2 and T2K computers in Japan. A first evaluation of the runtime systems developed on Task 2 will be proposed.

TASK 8, track 4: integration and evaluation (April 2013 – March 2014)

The integration of API developed in Task 6 will be integrated with all other software developed in former tasks, including data management systems. Benchmarks and library components will be adapted to use the

new developed framework to evaluate all the software and programming paradigms proposed into the project.

TASK 9, track 3, Library packaging, benchmark (April 2013 – March 2014)
TASK 10, track 4, Tutorials, demos and training (April 2013 – March 2014)

C.4 DIFFICULTIES ENCOUNTERED AND SOLUTIONS

We mainly face two major problems since the beginning of the project, which generate a seven months delay and postpone milestones and deliverables.

First, we confront a strong challenge to success to prepare the agreement consortium document. As this problem was discuss already with ANR and JST, we will not detail here the problem but we may just remind that the Japanese partners received their financial supports since the beginning of the project, contrary to the French partners which receive money just after the signature of the document. This difference of financial support of the teams generated some difficulties to start the project as we scheduled on the proposal, as French researcher was not able to travel in Japan before February –March 2011, i.e. 7 months after the beginning of the project. To be able to launch the project without financial support from ANR, we nevertheless were able to organize the kick-off meeting in September 6-7, 2010, at INRIA Saclay and at University of Versailles. Then, the project was launch as scheduled but only Japanese partners were able to hire people and travel in France.

Second, the majority of the French partners were not able to travel in Japan during a few months, as CNRS and CEA forbidden travels in Japan following the Fukushima accident. This interdiction was just a few weeks after the French teams get the financial supports and several planed trips in Japan had to be last-minute cancelled.

We just obtain at last a few weeks ago, approximately in October-November 2011, quasi-normal conditions to work on the project, i.e. more that one year after the beginning of the project.

These two major difficulties generate an important delay but we were able to work during visit of Japanese colleague in France and during international conferences. Hopefully, we success to organize the kick-off meeting and we already had a quite long experience to work all together, then, we always were able to keep the project working, but with an important delay concerning milestones and deliverable. We evaluate the delay at about 6-7 months.

Then, the schedule has to be updated. We'll keep the same scheduling of the milestones and deliverables, but with a new timeline. We propose to set the end of the project at March 31, 2014, i.e. seven months later that originally scheduled. Then, this report had to be considered as the report at the end of the first year approximately of the FP3C projects.

We plan to keep all the main milestones and deliverables as scheduled, but with this 7 months delay, as the ones originally scheduled for the end of the first year would be globally finished by the end of March 2012. Then, the project content has not to be reviewed. The proposed objectives of the FP3C project are still to be completed.

C.5 SIGNIFICANT EVENTS AND RESULTS

- *YML/XMP framework and two-level parallel programming, including first experiments on T2K and GRID5000.*

We finish to integrate YML and XMP as scheduled and the YML-XMP software has been released for users. Some experiments on T2K computer in University of Tsukuba in Japan and on GRID5000 in France have been made and based on the Block-based Gauss-Jordan method. At this time, we have inverted a matrix 8000x8000 on 64 nodes. Another result is the default to find the accurate grain of processing for each parallelism level.

- *XMP-dev: XcalableMP acceleration device extension*

We propose an extension of XcalableMP (XMP) PGAS language for a cluster with acceleration devices. In the XMP execution model, global data is distributed onto each node, and computation is mapped on nodes according to the data distribution. We extended the model to acceleration devices by allocating global data distributed on the memory of acceleration devices and offloading the codes working on it in each node. The communication between acceleration devices can be described directly in the PGAS

framework. We have implemented our extension for the NVIDIA GPUs. The performance evaluation for a GPU cluster demonstrated that not only does our model achieve scalable performance by the optimization of GPU-to-GPU communication, but it also improves the productivity since it only requires small modifications and additional codes of directives to the original serial counter-part.

- *Linear Algebra solvers*

Experiments are reported by PRISM, CEA and INRIA Saclay for ERAM eigensolvers on GPU and show high performance compare to multicore implementation both on dense and sparse matrices.
The Parallel ILU(k) Preconditioner for Ill-Conditioned Problems and the multigrid preconditioner developed at Tokyo University demonstrates good scalability on the T2K Open Supercomputer. A finite element method has been successfully implemented on GPU by Tokyo University.

C.6 WORK SPECIFIC TO THE COMPANIES (WHERE APPLICABLE)

No applicable

C.7 CONSORTIUM MEETINGS (COLLABORATIVE PROJECTS)

Cf. project website for details

Date	Place	Partners present	Subject of the meeting
September 2010, 6th	INRIA Saclay	All partners: 20 Frenchs, 23 Japanese, and 7 executive people from agencies or partner administration, or ANR.	Kick off meeting, first day. Opening from the director of INRIA, Michel Cosnard, ANR and JST executives attend the meeting.
September 2010, 7th	University of Versailles	All partners: 20 Frenchs, 23 Japanese.	Kick off meeting, second day. Plenary sessions, and 4 working sessions in parallel.
September 2011, 2d-3rd	Bordeaux	All partners: 16 French, 23 Japanese	Annual plenary meeting, Plenary sessions and 4 working sessions in parallel.
November 2011, 14th, 1-5pm.	Seattle, Hilton Hotel	All partners: 9 Frenchs, 13 Japanese.	Working meeting in conjunction of he conference SC11.

We often have informal FP3C meetings during international meetings or conferences, as all the partners participate to several program committees together and attend the major conference of the domains. The actual worldwide important activities on HPC generate a lot of such meetings. SC and SIAM conferences, for example, always generate several working meetings between French and Japanese colleagues participating to the FP3C project. As an example, during the next week SIAM conference “Parallel Programming” in USA, several discussions are planed; 5 Minisymposiums of this conference are organized by partners of the project (each includes participations from French and Japanese researchers of the project). Discussions during SIAM conference in “Linear Algebra”, June in Spain, or “PMAA’, June in London, are also scheduled.

See appendices for the new schedule of future FP3C meetings

C.8 FREE COMMENTS

Comments from the French/Japanese coordinators (PIs)

Over the scientific results, this collaboration explores several of the problems we have to face to establish stable official collaborations between French and Japanese teams. It would be much easier on the future to have French-Japanese research projects on HPC. Nevertheless, our project generates only open-source software. Otherwise, it would have been quite impossible to concretize the signature of a consortium agreement.

It would have been better if ANR and JST had harmonized the management of the project on both side, such as starting the financial support at the same time and having attributed the same amount of funding at each side.

Question(s) posed to the ANR/JST

French partners question to ANR: Due to the problems we faced during the first year which generate a seven months delay and oblige us to postpone milestones and deliverables, we would efficiently take advantage of an extra funding corresponding to the financial engagements we had during the first seven months of the project where we work without any support on the FP3C project. As we had to organize an unscheduled plenary meeting in Bordeaux, because we were forbidden to travel in Japan during a few months, we would like to have an extra financial support also to reimburse INRIA Bordeaux (cost 2275 euros).

Japanese partners question to JST: Due to the problems of seven months delay, we are afraid we will not have enough budget for the extended year (2013). We re-considered the budget plan for 2012 to extend some part to 2013, but it may be not enough for the out-reach activity scheduled in 2013. Especially, University of Tsukuba would request the budget to hire one Post-Doc for the last year.

D PROJECT VALORIZATION AND IMPACT SINCE BEGINNING

D.1 PUBLICATIONS AND COMMUNICATIONS

<Joint> Multinational Joint Papers, etc

List of the Multinational publications (resulting from jointly conducted work)		
International	**Peer-reviewed journals**	1. 2.
	Books or chapters in books	1.
	Communications (conferences)	1. Leonardo Bautista Gomez, Akira Nukada, Naoya Maruyama, Franck Cappello and Satoshi Matsuoka, Low-overhead Diskless Checkpoint for Hybrid Computing Systems, International Conference on High Performance Computing, Goa India, 2010 2. FTI: High performance Fault Tolerance Interface for hybrid systems, International Conference for High Performance Computing, Networking, Storage and Analysis (SC11), 2011, Seattle, USA.
France	**Peer-reviewed journals**	1.
	Books or chapters in books	1.
	Communications (conferences)	1.
Japanese	**Peer-reviewed journals**	
	Books or chapters in books	
	Communications (conferences)	1. Leonardo Bautista, Akira Nukada, Naoya Maruyama, Franck Cappello, Satoshi Matsuoka. Low-overhead checkpoint for large-scale GPU-accelerated systems, ARC192 HPC128-22, Jan. 2011.
Outreach initiatives	**Popularization articles**	1.

	Popularization conferences	1.
	Mini-symposium organizations	1. *Creating the Next Generation of High Performance Numerical Computing Capabilities*, Part I-IV, 7th International Congress on Industrial and Applied Mathematics (ICIAM 2011), Vancouver, Canada (2011) 2. *Parallel Programming Models, Algorithms and Frameworks for Scalable Manycore Systems.* Part I-IV, 15th SIAM Conference on Parallel Processing for Scientific Computing (PP12), Savannah, Georgia, USA (2012) 3. *Advanced methods for large eigenvalue problems and their applications*, organize by Tetsuya Sakurai and Nahid Emad for SIAM Conference on Applied Linear Algebra, June 18th-22nd, 2012, Valencia, Spain.

<French side> Single partner Papers, etc

List of single-partner publications (involving a single partner)		
International	**Peer-reviewed journals**	1. 2.
	Books or chapters in books	1. 2.
	Communications (conferences)	1. Alexandre DENIS. A High Performance Superpipeline Protocol for InfiniBand. In E. Jeannot, R. Namyst, and J. Roman, editors, Proceedings of the 17th International Euro-Par Conference, number 6853 of Lecture Notes in Computer Science, Bordeaux, France, pages 276-287, August 2011. Springer. 2. Elisabeth Brunet, Francois Trahay, Alexandre Denis, and Raymond Namyst. A sampling-based approach for communication libraries auto-tuning. In International Conference on Cluster Computing (IEEE Cluster), Austin, Texas, pages 299-307, September 2011. IEEE Computer Society Press.
France	**Peer-reviewed journals**	1. 2.
	Books or chapters in books	1. 2.
	Communications (conferences)	1. 2.
Outreach initiatives	**Popularization articles**	1. 2.
	Popularization conferences	1. 2.
	Others	1. 2.

<Japanese side> Single partner Papers, etc

List of single-partner publications (involving a single partner)		
International	**Peer-reviewed journals**	1. Yasunori Futamura, Hiroto Tadano, and Tetsuya Sakurai, Parallel stochastic estimation method of eigenvalue distribution, JSIAM Letters, Vol. 2, pp. 127–130, 2010. 2. T. Mifune, Y. Takahashi, and T. Iwashita, "New Preconditioning Technique to Avoid Convergence Deterioration due to the Zero-Tree Gauge Condition in Magnetostatic Analysis," IEEE Trans. Magn., Vol. 46, No. 7, pp. 2579-2584, 2010. 3. Y. Takahashi, T. Tokumasu, A. Kameari, H. Kaimori, M. Fujita, T. Iwashita, and S. Wakao, "Convergence Acceleration of Time-Periodic Electromagnetic Field Analysis by Singularity Decomposition-Explicit Error Correction Method," IEEE Trans. Magn., Vol. 46, No. 8, pp. 2947-2950, 2010 4. Nakajima, K.: Parallel Multistage Preconditioners by Extended Hierarchical Interface Decomposition for Ill-Conditioned Problems, Advances in Parallel Computing Vol.19 "From

		Multicores and GPU's to Petascale", IOS press, 99-106 (2010) 5. Yasuto Takahashi, Takeshi Iwashita, Hiroshi Nakashima, Shinj Wakao, Koji Fujiwara, and Yoshiyuki Ishihara. Performance Evaluation of Parallel Fast Multipole Accelerated Boundary Integral Equation Method in Electrostatic Field Analysis. IEEE Trans. Magnetics, Vol. 47, No. 5, pp. 1174-1177, May 2011. 6. Yasuto Takahashi, Takeshi Iwashita, Hiroshi Nakashima, Shinj Wakao, Koji Fujiwara, and Yoshiyuki Ishihara. Micromagnetic Simulations of Perpendicular Recording Head Using the Parallel Fast Multipole Method Specialized for Uniform Brick Elements. IEEE Trans. Magnetics, Vol. 47, No. 10, pp. 3805-3808, October 2011. 7. Y. Nakamura, K. -I. Ishikawa, Y. Kuramashi, T. Sakurai, H. Tadano, Modified block BiCGSTAB for lattice QCD, Comput. Phys. Commun. Vol. 183, pp. 34-37 (2011). 8. Y. Maeda, Y. Futamura, T. Sakurai, Stochastic estimation method of eigenvalue density for nonlinear eigenvalue problem on the complex plane, JSIAM Letters, Vol. 3, pp. 61-64 (2011). 9. I. Yamazaki, H. Tadano, T. Sakurai, K. Teranishi, A convergence improvement of the BSAIC preconditioner by deflation, JSIAM Letters,Vol.3 pp.5-8, (2011).
	Books or chapters in books	1. 2.
	Communications (conferences)	1. Akihiro Nomura and Yutaka Ishikawa, "Design of Kernel-level Asynchronous Collective Communication," Proceedings of Euro MPI, September 2010. . 2. Ali Cevahir, Akira Nukada, and Satoshi Matsuoka. “High Performance Conjugate Gradient Solver on Multi-GPU Clusters Using Hypergraph Partitioning”, In Proceedings of the 2010 International Supercomputing Conference (ISC’10), Hamburg, Germany, June 2010. 3. Koichi Shirahata, Hitoshi Sato, Satoshi Matsuoka. Hybrid Map Task Scheduling for GPU-based Heterogeneous Clusters, First International Workshop on Theory and Practice of MapReduce (MAPRED'2010), First International Workshop on Theory and Practice of MapReduce (MAPRED'2010), Jan. 2011. 4. Nakajima, K.: New Strategy for Coarse Grid Solvers in Parallel Multigrid Methods using OpenMP/MPI Hybrid Programming Models, ACM Proceedings of PPoPP/PMAM 2012, (2012) (in press) 5. Ohshima, S., Hayashi, M., Katagiri, T., Nakajima, K.: Implementation and evaluation of FEM code for 3D Solid Mechanics on CUDA, 2011 SIAM Conference on Computational Science and Engineering (CSE11): Parallel Programming Models and Algorithms for Scalable Manycore Systems Part-I of III, Reno, Nevada, USA (2011) 6. Nakajima, K., Hayashi, M., Ohshima, S.: Parallel Preconditioning Methods for Ill-Conditioned Problems on Multicore Clusters, 2011 SIAM Conference on Computational Science and Engineering (CSE11): Parallel Programming Models and Algorithms for Scalable Manycore Systems Part-II of III, Reno, Nevada, USA (2011) 7. Katagiri, T.: Towards Auto-tuning Language of Numerical Libraries in Heterogeneous Computing Era, 7th International Congress on Industrial and Applied Mathematics (ICIAM 2011): Creating the Next Generation of High Performance Numerical Computing Capabilities Part-I of IV, Vancouver, Canada (2011) 8. Nakajima, K., Hayashi, M., Ohshima, S.: Infrastructure for Application Development on Heterogeneous Parallel Computers, 7th International Congress on Industrial and Applied Mathematics (ICIAM 2011): Creating the Next Generation of High Performance Numerical Computing Capabilities Part-III of

		IV, Vancouver, Canada (2011) 9. Nakajima, K.: Coarse Grid Solvers in Parallel Multigrid Methods using OpenMP/MPI Hybrid Programming Models, 15th SIAM Conference on Parallel Processing for Scientific Computing (PP12): Parallel Programming Models, Algorithms and Frameworks for Scalable Manycore Systems Part-III of IV, Savannah, Georgia, USA (2012) 10. Ohshima, S., Hayashi, M., Katagiri, T., Nakajima, K.: Implementation of FEM Application on GPU, 15th SIAM Conference on Parallel Processing for Scientific Computing (PP12): Parallel Programming Models, Algorithms and Frameworks for Scalable Manycore Systems Part-IV of IV, Savannah, Georgia, USA (2012) 11. Hiroshi Nakashima. Generator for Library and Application --- Splitting What & How by Domain-Specific Local-View Programming---.In Proc. Intl. Cong. Industrial and Applied Mathematics, July 2011. 12. Hiroshi Nakashima. Local View Kernels: A New Programming Scheme for Plasma Simulation. In Proc. Plasma Conf., November 2011. 13. Takeshi Iwashita, Yu Hirotani, Takeshi Mifune, Toshio Murayama and Hideki Ohtani.Parallel Geometric Multigrid Solver for Fast Electromagnetic Wave Analysis'In Proc. Intl. Workshop on application of iterative methods to engineering and its mathematical element, pp. 166-173, October 2011 14. Shunsuke Mikami, Kazuki Ohta and Osamu Tatebe, "Using the Gfarm File System as a POSIX compatible storage platform for Hadoop MapReduce applications", Proceedings of 12th IEEE/ACM International Conference on Grid Computing (Grid 2011), 2011 15. Kenji Kobayashi, Shunsuke Mikami, Hiroki Kimura, Osamu Tatebe, "The Gfarm File System on Compute Clouds", Proceedings of 1st International Workshop on Data Intensive Computing in the Clouds (DataCloud 2011), 2011 16. Physis: An Implicitly Parallel Programming Model forStencil Computations on Large-Scale GPU-Accelerated Supercomputers, International Conference for High Performance Computing, Networking, Storage and Analysis (SC11), 2011, Seattle, USA. 17. Jinpil Lee, Minh Tuan Tran, Tetsuya Odajima, Taisuke Boku and Mitsuhisa Sato: An Extension of XcalableMP PGAS Language for Multi-node GPU Clusters, Ninth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar 2011), 2011.
Japanese	**Peer-reviewed journals**	1. Nakajima, K.: Parallel Preconditioning Methods for Iterative Solvers in Multi-Core Era, Mathematical foundation and development of algorithms for scientific computing, RIMS Kokyuroku 1773, Research Institute for Mathematical Science, Kyoto University, 1-10 (in Japanese) (2011) 2. Jinpil Lee, Minh Tuan Tran, Tetsuya Odajima, Taisuke Boku and Mitsuhisa Sato , An Extension of XcalableMP PGAS Language for a Cluster with Offloaded Acceleration Devices, IPSJ Transactions on Advanced computing systems(in Jpanaese),2012 (accepted)
	Books or chapters in books	3. 4.
	Communications (conferences)	1. Hayashi, M., Ohshima, S. and Nakajima, K.: Parallel ILU Preconditioner based on Extended Hierarchical Interface Decomposition for Heterogeneous Environments, IPSJ SIG Technical Reports (in Japanese), IPSJ-HPC11130005 (2011) 2. Ohshima, S., Hayashi, M., Katagiri, T. and Nakajima, K.: Implementation and Evaluation of 3D Finite Element Method for CUDA, IPSJ SIG Technical Reports (in Japanese), IPSJ-

		HPC11129020 (2011) 3. T. Odajima, T. T. Chan, J. Lee, T. Boku and M. Sato, "Extend to GPU for XcalableMP: A Parallel Programming Language",IPSJ SIGHPC Technical Report 2011-HPC-129 (on-line), May 2011.
Outreach initiatives	**Popularization articles**	3. 4.
	Popularization conferences	3. 4.
	Others	1.

D.2 OTHER VALORIZATION FACTORS

Publications are expected the next months. Papers are under writing and will be send soon to international conference and journal. The published paper and available software will be listed to the FP3C website.

List of factors. Indicate the titles, years and comments	
International patents obtained	No applicable
International patents pending	No applicable
French National patents obtained	No applicable
Japanese National patents obtained	No applicable
French National patents pending	No applicable
Japanese National patents pending	No applicable
Operating licences (obtained / transferred)	
Company creations or spin-offs	
New collaborative projects	1. A France-Berkeley-Fund project linked with FP3C is submitted in 2012, January on the theme “Sustainable Numerical Kernels for Extreme Scale Computing Environments” - PRiSM.
Scientific symposiums	1.
Others (specify)	1. A three months (2012, June-September) internship of a French student (University of Versailles) in Tsukuba University on the theme “MERAM/SSM programming with XMP”. 2. Talk of Tetsuya Sakurai on ”*Design of computational algorithms to utilize petascale supercomputers* in “Maison de la Simulation” 3. Talk of Mitsuhisa Sato on ”K computer and XcalableMP parallel language project-Toward programming environment for pet-scale computing” in « Maison de la Simulation”

E APPENDICES

E.1 NEW SCHEDULE OF THE FP3C PROJECT

Official beginning of the project
End of the « first year »: March 31, 2012
End of the « second year »: March 31, 2013
End of the project: March 31, 2014

We keep the work schedule of the second and third year of the project as described on the proposal, with the above new date for the end of the “first” year (i.e. t0+12 months is March 31, 2012).

E.2 NEW SCHEDULE OF THE FP3C FUTURE MEETINGS

September, 1st, 2010	**Official beginning of the FP3C project**	
September 6-7, 2010	Kick off meeting	INRIA Saclay (6th) and PRiSM, University of Versailles (7th)
September, 4-5, 2011	Plenary meeting	INRIA Bordeaux, France
November 14, 2011	Meeting during SC11	Seattle, WA, USA
March 13-14, 2012	Review meeting	AICS, Kobe, Japan
March 31, 2012	**End of the of the "first year", new T0+12 months**	
July 16, 2012	Plenary meeting	Univ. of Kyoto or AICS Kobe, Japan
November 12, 2012	Meeting during SC12	Salt Lake City, UT, USA
June 2013	Plenary Meeting	IRIT, Toulouse, France
November 18, 2013	Meeting and BOF at SC13	Denver, CO, USA
February 2014	Final plenary meeting	University of Tsukuba, Japan
March 2014	Demos and tutorials	CEA, Saclay, France
March 31, 2014	**End of the FP3C project**	