Updated proposal and workplan of the SUMA project

The INFN-SUMA team: R. Alfieri, R. Ammendola, S. Arezzini, D. Badoni, M. Brambilla, R. Benzi, L. Biferale, A. Carboni, A. Ciampa, L. Cosmai, R. Covati, G. De Divitiis M. D' Elia, R. De Pietri, C. Destri, F. Di Renzo, A. Feo, R. Frezzotti P. Giannozzi, L. Giusti, G. La Penna, V. Lubicz, A. Lonardo, E. Mazzoni, V. Minicozzi, S. Morante, P. S. Paolucci, A. Papa M. Papinutto, M. Pepe, R. Petronzio M. Pivanti, S. Piras, F. Rapuano, D. Rossetti, G. C. Rossi, A. Salamon, G. Salina, M. Sbragaglia, S. F. Schifano, S. Simula, N. Tantalo C. Tarantino, R. Tripiccione, E. Vicari, P. Vicini, A. Vladikas

January 14, 2013

Abstract

This document is a detailed proposal and workplan of the SUMA project. The main goals of this project are: i) support the currents need of the theoretical computational community at INFN; ii) perform early experiments on forthcoming computing architectures in order to prepare for a smooth and efficient transition towards exascale computing: iii) further support the technological developments in computing architecture made at INFN and investigate all ways in which these developments can be exploited for Exascale computing; iv) bring together several institutions and research groups, to help establish a national eco-system in the area of the computational sciences. This paper first considers the broad strategic ideas behind the project and then describes in details the structure of the endeavor, its organization, its break-up in workpackages and its cost estimate. The document ends with an appendix on the physics background to the project.

1 Overview

1.1 A bird's eye view on SUMA

This document is an updated and improved proposal and a detailed workplan of the INFN SUMA project. SUMA aims to provide a substantial support to the computing needs of the theoretical physics community within INFN, on a finite time horizon of 2 - 3 years.

Scientific case

The INFN theoretical community is active in several scientific areas that require significant computational support. These areas stretch over a wide spectrum, requiring in some cases fairly limited computing resources – for instance in nuclear physics, high-energy physics phenomenology, spin-system simulations – while, at the other end of the spectrum huge computing power, that can only be provided at the transnational level, is needed; examples in this class are LQCD, dynamical systems and classical and ab initio simulations of bio-systems.

At the same time, for most groups active in these areas, it is becoming more and more difficult to develop their computational strategies and algorithms in a way that allows to adapt to the increasingly fast changes happening in high performance computing architectures.

Last but not-least, several existing INFN projects have produced significant progress on technological developments that may be crucial building blocks for new generation HPC systems, if it can be shown that they are efficient solutions to (at least some) large scale computational problems.

SUMA plans to support all these physics goals, and at the same time aims to explore all suitable ways in which the technological developments made at INFN can be put to good use for the present and future needs of computational physics.

Methodology

SUMA plans to reach its objectives through a variety of coordinated activities:

1. Development of computational strategies and algorithms, and implementations of the corresponding codes for state-of-the-art platforms and for new architectures.

2. Optimization of the algorithms and codes of interest for several existing as well as experimental architectures, that include building blocks that may likely be key elements of future HPC architectures, and assessment of their performance.

3. Upgrade of the existing middle-size, cluster-based computational infrastructure available at INFN (Theophys).

4. Procurement and operation of a "large prototype" HPC system whose architecture comes as close as possible to that of a next generation supercomputer and, based on the assessment made in the work of the previous points, significantly enhances the computing resources available at INFN.

5. Further development at the technological level, of several important building blocks for future HPC, especially in the area of interconnections networks and their optimization for specific computational tasks.

The activities outlined above will be put in practice, as much as possible, in close collaboration with other institutions and research centers in Italy and abroad, in order to develop all possible synergies; we want to build a national ecosystem for HPC able to command increased support for computing for theory, both at the national and EU level and, as far as possible, to support Italian enterprises active in this area.

Project Timing

The natural time-frame of the project extends over three years. While points 1), 2) and 5) of the previous list stretch over the whole lifetime of the project, point 3) will start immediately, aiming to have an upgraded cluster operational in fall 2013. Point 4) will start in summer 2013, aiming to have a clear definition of the "large prototype" in late 2013 and an operational system up and running in early 2014.

Man-power

SUMA is based on a core team of some 30 people, who have a strong and complementary know-how in all areas in which the project is active. These senior members of the project will be able to work part-time on SUMA; for this reason successful development of the project requires a quick injection of fresh forces fully committed to the project: we propose to hire approximately 10 post-docs (*assegni di ricerca*); for most of these positions we envisage a two year contract. For some of this positions, co-financing by several Departments seems to be possible. On the other hand, the prsent members of the project are part of a large and established web of international collaborations with universities, research institutions and computer centers that provides rich opportunities that – we expect – will strongly support this new project.

Project costs

Our preliminary forecast for the cost structure of the project fits within the limits of the budget made available by MIUR for this "Progetto Premiale". The grand total costs of the project add up to 1925 K€, roughly split in i) ad hoc manpower: 600K€, ii) hardware procurement: 820K€, iii) advanced technological developments: 325 K€; iv) other: 180 K€.

1.2 Document structure

The rest of this document, that describes in details the structure and the workprogramme of the project, is organized as follows:

- Section 2 briefly outlines the current state of the art and its expected trends. It first covers the current status of some physics computational problems on which the SUMA team plans to work. This section then reviews the current situation and prospects of High Performance computing, the technological scenario in which new machines will be developed and the experimental computer systems that will be available to the project.

- Section 3 provides a global overview of the project workprogramme, that we plan to organize in 4 Workpackages. We provide here a short description of all Workpackages and outline their mutual dependence.
- Section 4 expands on the previous one, describing in details the goals and the structure of each workpackage.
- Section 5 outlines the organizational structure and the management of the project.
- Section 6 describes the web of collaborations linking SUMA with other institutions and research groups, both in Italy and abroad, explaining the synergies and mutual advantages to be expected by this collaborative structure.
- Section 7 presents a preliminary analysis of the cost structure of the project.
- An appendix ends the document, describing the physics programme behind the project.

2 State of the Art

This section describes the state-of-the-art in several areas that we regard as starting points for SUMA.

We start by briefly addressing the physics focus on which the INFN groups driving SUMA are active: this area is obviously the driving force behind the project, since it defines the computational requirements that must be addressed by the project. We then briefly review the main lines of evolution of High Performance Computing; this sets the boundary conditions – on which we have negligible control – on which SUMA has to build upon. Next, we discuss one specific area of development in computer architectures – interconnection networks – on which INFN has a long standing tradition of original and innovative research, that we would like to exploit in all possible directions. The next point is a description of two experimental HPC computer systems – both including some of the original developments made by INFN – that we will be able to use for the early phase of the project. Finally, we look at the current status of the middle-size INFN cluster, since part of the project will focus on improving – at the hardware and software level – this facility.

- **Physics Projects.** We present here in very concise form the physics work behind the SUMA project carried out by several groups; we provide a more detailed analysis in an appendix. The most critical area is obviously Lattice QCD (LQCD), but simulation of complex systems and quantitative biology also play an important role.

  Since several years the group in Roma Tre is involved in the initiatives of the European Twisted-Mass Collaboration (ETMC), which gathers more than 30 permanent researchers and young students belonging to many institutions all over the Europe. In the last few years the ETMC has produced a huge amount of high-quality ensembles of gauge configurations with two ($N_f = 2$ [1, 2]) as well as four ($N_f = 2 + 1 + 1$ [3, 4]) flavors of

dynamical fermions. In the latter case, besides two light mass-degenerate quarks, the strange and charm quarks with masses close to their physical values are included in the sea. Such a setup is the closest one to the real word, adopted till now only by ETMC in Europe and by MILC in the USA.

The long term physics goal at Milano Bicocca is to study the chiral regime of QCD, and to carry out non-perturbative computations needed in the flavour phenomenology of the Standard Model and beyond. Our most computational-intense activity is the generation and the analysis of lattices with dynamical flavours of non-perturbatively improved Wilson fermions. We simulate pions near or at the physical point with lattice spacing small enough and volumes large enough to be able to reach the infinite volume and continuum limit with confidence.

One of the main activities in Pisa is represented by the study of non-perturbative aspects of strong interactions by LQCD simulations, with a focus on color confinement and on the properties of the deconfinement transition at finite temperature and in presence of external parameters of phenomenological interest, such as a finite baryon density or external background fields. In order to approach the hard computational tasks involved, the group in Pisa will explore in details the merits of GPU-based accelerations for LQCD algorithms. GPUs were developed originally as fast co-processors for graphics tasks. GPUs evolved towards programmable many-core chips able to process in parallel massive amounts of data.

Numerical Stochastic Perturbation Theory has been for several years the main research interest in Parma in Lattice Gauge Theories, always with a focus on computing efficiency. In recent years, many experiments have been performed to chase efficiency in (much more) standard non-perturbative Lattice QCD simulation tools, with a focus on dealing with the multiple levels of parallelism which are in place on present computer systems. Taking advantage of these activities, a project aims at a better control on finite-lattice size and finite volume effects: models developed in NSPT are being extended to the non-perturbative framework.

Parma is also interested in the dynamics of the magnetic and gravitational fields, that play a major role on the physical properties of isolated Magnetized Neutron Stars and on the modeling of the gravitational signal coming form the merger of two Neutron Stars. From the INFN point of view, this kind of research is of particular interest in light of the LIGO/VIRGO effort to directly detect Gravitational Waves (GWs); accurate knowledge of the possible GWs signal of such phenomena is needed to have detection. The main application used by the INFN researcher involved in this kind of research is the Einstein Toolkit and, in particular, Cactus/Whisky. These applications are actual in production on the "Theophys" cluster where it is routinely used allocating 128 cores. They have already been tested on the Blue Gene/Q CINECA system up to 16384 cores and it proved to scale well up to this number of MPI processes.

Tor Vergata focuses its work on quantitative biology, fluid-dynamics and LQCD. Systems of biological interest are especially difficult to model and treat mathematically for two fundamental reasons. One is the large num-

ber of relevant degrees of freedom that are needed to describe the interesting features of the system and the second is their heterogeneity.

More specifically, even within a given model setting, a realistic description of the behavior of a biological system requires 1) dealing with long-range interactions and large volumes and 2) the ability of studying the effect of changing external boundary conditions and/or constraints (like temperature, pressure, pH, etc.).

All this makes it compulsory the recourse to numerical simulations for the study of the properties of any biological model and its comparison with experimental data. Numerical simulations are, however, particularly hard here since, owing to the lack of locality and homogeneity, parallelization is not as effective as in other situations (like, say, lattice QCD) and node-to-node communication becomes the key issue. New architectures, pointing to full node connectivity, will then be necessary in this fascinating and important area of investigation.

Tor Vergata, together with Ferrara, also works om Lattice Boltzmann methods (LBM). LBM represent an optimal formulation of the Boltzmann kinetic equation describing the dynamics of a fictitious ensemble of particles whose motion and interactions are confined to a regular space-time lattice. Major strengths and advantages with respect to other methods are the capability of handling boundary conditions associated with irregular geometries and the efficiency to describe non-ideal fluids with phase transitions/phase separation. Moreover, LBM allows for describing hydrodynamical fluctuations on scales of experimental interest. The investigation of the dynamical behavior of these systems requires very long time integrations, typically of the order of tens of million of time steps (as a reference, one LB time step can be taken of the order of 100 1000 molecular dynamics time-steps). As a result, the need of a very-long simulation span, sets a strong incentive for efficient implementations.

Coming now to LQCD, Tov Vergata has had for several years a strong activity on the study of nonperturbative QCD through lattice simulations. In the coming year, efforts will focus on the study of leptonic and semileptonic decay rates of heavy-light mesons, of precise estimates of isospin breaking effects and on the hadronic contribution to the anomalous magnetic moment. The challenge of dealing with heavy-light systems is to accomodate multi-scale problems and, beside the approach of finite volume recursive methods, this may lead to novel multiscale algorithms and have impact on computer architectures.

- **Current trends in HPC architectures.** In the last 5 years, processor manufacturers have adopted the *multi-core* design approach to overcome the de facto upper limit on processor frequency, that current technology sets at ≈ 3 GHz. A multi-core processor is a single chip integrating two or more independent CPUs. This approach allows processors to still scale according to Moore's law, but it bears a great impact on application design, further moving the challenge of sustaining performance from hardware to software.

  Currently two main classes of novel architectures start to become available on the market: *multi-core CPUs* and *many-core GPUs*.

Processors of the first class integrate 8-10 (and soon more) *x86-compatible* CPUs, including each 2 level of caches, and a large third level of cache shared among the cores. Devices of such class are the Intel processors based on the recent *Sandybridge* micro-architecture. Within the first class another approach is emerging called *Many integrated Cores* (MIC) designed by Intel. In this case the processor integrates around 60 lightweight core, each with a small L1 and L2 cache, and interconnected by a ring network which allows to share the contents of the L2 caches. A lightweight core is a simplified version of an x86-CPU not including, for example, dynamic scheduling of instructions. The first processor in this class is the *Knights Corner* processor that will be available in early 2013. These processors support the execution of SIMD instructions to further improve performances, and performance critically relies on the ability of programmers and compilers to exploit the execution of these instructions.

GPU-class processors are an evolution directly stemming from developments of graphic cards. Devices of this class are for instance the systems developed by NVIDIA using the Fermi or Kepler processor. The architecture of these devices includes many (e.g. more than 10) processors, each able to perform a 32-way SIMD instruction per clock cycle. Using NVIDIA terminology these chips run hundreds of cores and deliver a peak performance of approximately 1-4 TFlops in single precision, and 0.5-1 TFlops in double precision.

At the next hierarchy level, typical HPC systems for scientific computing today are based on large clusters of processing nodes integrating several commodity multi-core processor, and interconnected by a high-speed switched network, such as Infiniband. On each node typically two, four or eight CPUs can be installed; each CPU has its own memory-bank and is connected to the other sockets by dedicated buses. Moreover each node can includes also one, two or four GPUs, and soon also MICs boards. GPUs and MICs are used as accelerators to run the most compute intensive kernels of applications.

- **Network Developments at INFN.** One of the main research topics for "Exascale computing" is the development of novel network architectures enabling effective performance scaling as the number of computing nodes quickly grows to reach the one million threshold.

  INFN has been very active in this area in recent years; recently progress has focused on the APEnet+ architecture, that was developed at Roma 1 and Tor Vergata; these units will continue their collaboration within the SUMA project. APEnet+ is an FPGA-based PCIe board with 6 bi-directional links with 34 Gbps of raw bandwidth per direction, state-of-the-art signaling capabilities - up to X8 Gen2 lanes towards the host PC - and a simple Remote Direct Memory Access (RDMA) programming model that leverages upon peer-to-peer (P2P) capabilities of Fermi-class NVIDIA GPUs. In GPUs, the large ratio between computing performances (several TFlops) and I/O capabilities (few tens of GB/s) as well as the required staging on host buffers of GPU transmitted data, are the limiting factors to effective scaling of multi-node multi GPU system; APEnet+ implements the P2P autonomous mechanism allowing the board to read and write

the GPU private memory directly over the PCIe bus. In this way, real zero-copy, inter-node GPU-to-host, host-to-GPU or GPU-to-GPU transfers can be performed, with substantial reductions in latency. Further improvements of this mechanism are a key factor to scale application performances on hybrid systems.

- **Experimental machines.** SUMA will be able to use several experimental computing systems, using new processor architectures.

  For tests at the level of just one processor, several small GPU-based systems are available at several SUMA sites; we can also remotely access several MIC-based systems, available at CINECA and at Jülich.

  Two larger systems will also be shortly available for testing. Both machines use new processors and innovative interconnection structures.

  The first such system is QUonG (lattice Quantum chromodynamics ON GPU) system, a hybrid, GPU-accelerated x86 cluster with 64 nodes and a 3D toroidal mesh topology, with bandwidths and latencies balanced for the requirements of several scientific codes. The QUonG floating point engines are the Fermi-class GPUs from NVidia, while the 3D Torus network is built on the APEnet+ architecture, described above. The current QUonG system is a 25 TFlops platform; it will be upgraded to 50 TFlops before the end of 2012 and made available to the collaboration from the beginning of the project.

  SUMA will also be able - starting from early 2013 - to use the experimental Eurora machine, being currently installed at CINECA. This machine builds upon previous experience, largely based on the Aurora system developed by Eurotech in collaboration with the INFN/FBK AuroraScience project. Eurora has 64 processing nodes with 128 latest-generation INTEL processors (Sandy-Bridge) and 128 MIC processors. The MIC-based peak performance is $\simeq$ 120 TFlops. The processors are connected not only by a standard Infiniband harness but also by a reconfigurable toroidal-network, that can be easily modified and upgraded;

- **Available middle-size computing facilities.** In Pisa there are now two clusters dedicated to theoretical Physics. Theophys, born in late 2009, has 1024 computing cores connected by an Infiniband network (Cisco SFS DDR IB switch 7024, 144 copper ports); the peak performance of this system is $\simeq$ 8.5 TFlops. The Cluster can be used serially (normal Grid usage) or in parallel (also through Grid). The cluster is intended as centralized service for serial and parallel computation. This service is available to all "gruppo IV" members. A priority "fair-share" access is granted on request based on specific scientific proposal. GRID is the method used to access Theophys.

  In early 2011 the CSN4 funded another cluster (THEONUC) directly aiming at the computing requirements of the nuclear physics community, for which large memory is a key feature: THEONUC has 240 computing cores and more than 800 Gbyte of memory. It uses a "memory driven" queue manager: users define their needs in term of RAM and the LSF scheduler organizes jobs considernig these requirements. THEONUC is accessible at

the national level using batch queues, by users that can use their INFN-AAI credentials. This is the first use of the INFN-AAI project for scientific computing.

3 The overall structure of the project

SUMA plans to organize a number of activities that should optimally respond – within the limits of the resources available to the project– to the strategic challenges outlined in the previous section.

The projects has a estimated time frame of three years and involves a fairly large initial team – necessarily working part time – scattered across several Sezioni, that we plan to enlarge opening several dedicated post-doc positions; this is described in details in Section 5. SUMA relies also on the collaboration of several research groups belonging to other institutions and working on strongly correlated research areas, in order to build up a nation-wide eco-system in the area of the computational sciences.

The project is organized in several interacting and coordinated workpackages, that are briefly summarized here and described in details in the following section:

- **WP1 - Computing strategies, Algorithms and Code Optimization for existing and new architectures.** This is one of the most significant WPs of the project. Its goal is to help the theoretical community to use as efficiently as possible the new computational building blocks that are becoming more and more relevant in high-performance systems. In this WP, we plan to select a number of relevant large size computing projects directly relevant to the INFN theoretical community, in areas such as LGT, fluid dynamics, complex systems, quantitative-biology, numerical relativity, and explore at all appropriate levels – from programming strategies to algorithms to coding style – the best ways to use new architectures and networks. This work will be supported by several experimental machines either directly available at INFN or made available by collaborating institutions. We also plan to identify computational problems studied in collaborating institutions and work together with their researchers to help optimize their algorithms and programs.

- **WP2 - Upgrade of the TheoPhys Cluster.** The goal of this WP – to be activated in the early phases of the project – is a substantial upgrade of the TheoPhys cluster currently operational at Sezione di Pisa. This machine, currently used by approximately 100 users, is rapidly becoming obsolete; we plan a complete overhaul of the system, replacing the processors and the network infrastructure, and improving access to all users. A key element here is the integration of the new cluster with a similar system to be installed soon at SISSA, so that INFN and SISSA users share both machines.

- **WP3 - Assembly of a large prototype of a state-of-the-art parallel system.** This WP builds on the assessment of the results made by the WPs in the first year of work, and selects and procure a reasonably large experimental system that may be considered as an intermediate step

towards Exascale level machines. This system should combine at best commercially available building blocks, as well as any element directly or indirectly stemming from the development activities made by INFN in recent years. The large prototype will be a crucial resource tool for all physics groups (in LQCD and elsewhere) to conceive and test new computing strategies for new probes which are particularly well suited for numerical computation and are also easily connected with interesting physical observables. We consider it very important that this procurement involves an industrial vendor and that the system be installed at a mayor national computer center and be made available for tests also to external collaborators, in order to provide the greatest possible visibility to the results of our project.

- **WP4 - Advanced developments.** This WP provides continued support to the computer development activities that have been carried out for years at INFN, aiming to develop and test fresh ideas that may evolve into key building blocks for scientific computing at the Exascale level.

4 Detailed description of the workpackages

This section describes in details the structure of the workpackages in which we plan to structure the project.

4.1 WP1 - Computing strategies, Algorithms and Code Optimization for existing and new architectures

WP1 focuses on a large variety of activities playing the dual roles of enhancing the quality of the physics results that the SUMA community is able to deliver in computing intensive areas, and *at the same time* strongly developing our knowledge of the new computing architectures that will be considered at the Exascale level and enabling us to exploit in the most timely and efficient ways these new architectures. Otherwise states, this WP supports the short term goal of improving the quality and impact of our physics results and the longer term goal of making sure that we can still grow in this direction by using at best and as early as possible the new computing tools that will become available.

Several sub-groups of the SUMA team will work on different physics projects, but we will make sure that all ideas and developments relevant across diverse physics area are quickly shared and exploited. In the following we list our proposed work under a few broad headings, and – within each heading – we describe the work done by each subgroup.

Lattice QCD

The work at Milano Bicocca will focus on generating lattice configurations with dynamical non-perturbatively improved Wilson fermions using the Lüscher's publicly available openQCD HMC code [7, 8]. It profits from several highly efficient numerical techniques such as the locally deflated solver, twisted-mass re-weighting, frequency splitting of the quark determinant through mass shifts, etc. The code is written in standard C, and in the next few months the new

AVX instructions (for Intel processors) and the QPX ones (for Blue-Gene/Q) will be included to optimize the most expensive parts of the program.

We will port and optimize this code for the new architectures which will be suited for running efficiently this code (task N. 1).

Once the optimization is available, we will conduct a scaling test of the openQCD code (task N. 2).

In the next step will will run the code on the new machines for producing the gauge configurations (task N. 3).

The computation of the correlation functions and the physics analysis will be carried out with our codes which use the same optimization strategy of the production ones described above (task N. 4).

For the next years the ETMC group, based mainly in Rome III, plans to generate new ensembles of gauge configurations with $N_f = 2 + 1 + 1$ and pions near or at the physical point. With these ensembles we want to perform the calculations of all the correlation functions relevant for light and heavy meson physics removing the need for a long extrapolation of the light-quark mass to the physical point. We expect to carry out such a program using the HPC resources made available through PRACE including the new BG/Q installation Fermi at CINECA. Nevertheless the search for new powerful algorithms will continue and we plan to test ideas and tools to address faster computations using the experimental machines provided by the SUMA project.

Work in Pisa will focus on the development of efficient algorithms for the numerical simulation of LQCD at finite temperature and in presence of chemical potentials or external background fields, with a particular interest in the physics and the critical properties of the deconfinement transition. On one hand, we will work at an efficient porting of existing LQCD codes to the new BG/Q architecture available at CINECA. On the other hand, our activity will aim at achieving further progress on LQCD codes running on many core architectures, starting from the single GPU implementation already developed by our group: we will focus on the portability to different kinds of many-core architecture (e.g. by making use of an OpenCL implementation) and on an implementation for multiGPU architectures, involving different level of parallelism. Such activity will be strictly interrelated with the development of experimental machines within the SUMA project.

The group in Parma will investigate general requisites for good performances of a Lattice QCD code on nowadays architectures.

- Task 1. The basic building block of a Lattice QCD simulation is the inversion of the Dirac operator by iterative methods, which require several applications of the Dirac operator itself. Since communications are the weak part, a natural algorithmic strategy tries to reduce the communications amount that is needed for the inversion. Domain decomposition techniques have been shown to be quite effective. A few publicly available codes and/or libraries exist (*i.e.* Luescher's DD-HMC, QUDA) which are natural starting points to compare the effectiveness of this algorithmic solution on the MIC and on GPU-based computing nodes.

- The quest for optimal balance between computation and communications on nowadays parallel architectures has to deal with the multiple layers of

parallelism which are in place: what we call nodes are actually parallel systems on their own. A few basic ingredients have been well established for a few years in the context of multicore architectures, *e.g.* the tools enabling core affinity and memory binding; solutions based on the simultaneous usage of MPI and threads-based coding have proven to be quite effective for Lattice QCD codes. Effectiveness of these choices on the MIC architecture will be assessed, always trying to compare with the progress going on on GPUs.

- Numerical Stochastic Perturbation Theory has been a playground for experiments that make use of both an Infiniband and a Torus networks at the same time, if they are both available on the same machine. Further investigation of this issue will continue on the architectures SUMA is going to test. This will be in the general framework of testing new network solutions that are expected to come from the work of WP4.

Specific computational issues of the LQCD group in Tor Vergata are the use of the unquenced Schroedinger functional scheme for matrix element evaluation and of suitable preconditioning techniques for heavy quark propagator inversion. Furthermore the medium size computing resources from SUMA are well suited for the Tor Vergata group program; indeed in order to apply the step-scaling method [9] for the calculation of the heavy-light meson observables, one needs to carry out simulations of small physical volumes: these simulations are needed in order to separate the high energy scale characterizing a heavy-light meson system (the heavy quark mass) from the low energy scala (Λ_{QCD}); they can be easily and efficiently performend on medium-size computer resources. The calculation of the QED isospin breaking corrections to the hadron spectrum requires the calculation of correlators involving a few sequential inversions of lattice quark and photon propagators[10]: the computational cost is comparable to the one required to extract momemtum dependent form factors on the lattice and fits well the computer resources available within SUMA. For the calculation of the muon anomalous magnetic moment we shall need, according to a recent proposal[11], observables very similar to the ones required in QED+QCD simulations and, consequently, well suited to the computer resources of the SUMA project.

Numerical Relativity

The Parma group will focus on porting their codes to GPU kernel and the MIC architecture and the possible optimization needed to take advance of these promising computational architectures. Main attention will go in assessing the effective performance of the existing "code base" on the new proposed large prototype of a state-of-the-art parallel system. This kind of application are very different from the one used in LGT, even if they are quite suitable to be parallelized (basically they involve the integration in time of Hyperbolic Equations on 3D grids); to achieve the desired accuracy one needs to use grid as big as 1000x1000x1000 points and likely needs to allocate order of TBytes of RAM.

Fluid Dynamics

Work in the improvement of Lattice Boltzmann algorithms and codes will be carried out at Tor Vergata and Ferrara, who have a long standing experience

on Lattice Boltzmann codes, in 2 and 3 dimensions. Recently the group developed an extremely sophisticated D2Q37 code that correctly handles the thermohydrodynamics of a 2D fluid that follows the equation of state of a perfect gas; in 3D a conceptually simpler code that deals with multiphase flows is also available.

These codes are already strongly optimized for massively parallel architectures. For instance, very high scalability up to 64 Kcores has been demonstrated on BG/P and BG/Q. The codes have been also ported to GPU-clusters and to the experimental QPACE machine; in both cases the codes rely on very careful overlap between communication and computation. This work has been done in close collaboration with the Ferrara group.

In this area, Task1 will focus on adapting the algorithm and the core to the available MIC architectures, assessing the overall performance and the scaling properties as a function of the number of cores.

Task 2 will explore the parallelization of the 2D and 3D codes on a large number of processors, first using a standard MPI-based communication infrastructure and then considering the torus network available on our experimental systems; the latter case may be very appropriate in this case, considering that the LBM algorithm is associated to nearest-neighbor communication patterns.

Molecular Dynamics and Quantitative Biology

A typical problem in biophysics is the investigation of the structural properties of bio-molecules (proteins, DNA, RNA, ...) as a starting point for the study and the modelization of their cross-interactions.

We want to concentrate on the problem of studying the physical-chemical properties of the atomic micro-environment where the first initial steps of the process leading to the formation of beta-amyloid neural fibrils takes place. The latter are peptide aggregates that are typically found in the cerebral tissue of people affected by the Alzheimer disease (AD). They are the result of yet not well understood processes of misfolding and successive aggregation of certain peptides, called beta-amyloids, in turn produced in the proteolytic cleavage of the APP (amyloid precursor protein). It is an experimental fact that the initial steps of the degenerative process are influenced by the presence of metallic ions (like Zn, Cu, Al) that among other effects have the ability of cross-linking pairs of peptides [12, 13, 14]. The presence of metals makes the numerical investigation of model systems describing beta-peptides assembly in interaction with them (possibly in an aqueous ambient) extremely demanding as quantum-mechanical ab initio calculations are necessary to detect and study the formation of chemical bonds between the metal ion and the peptide [15].

Simulations at the quantum level are performed in the framework of Density-Functional Theory, using first-principles molecular dynamics of the Car-Parrinello (CP) type. The code employed is the CP module of the Quantum ESPRESSO distribution, based on ultra-soft pseudo-potentials and a plane-wave basis set. This technique requires a large amount of RAM memory for complex systems like those under study, plus a sizable I/O, and does not lend itself to easy parallelization. The CP module is however parallelized (RAM memory is also distributed) using both MPI and (partially) OpenMP and scales well on large parallel machines like the Fermi BG/Q at CINECA.

In this situation it is one of the main goals of this part of the SUMA project

to build and test the performances of a middle-size but fully connected machine and compare them to those of the massively parallel platforms of the BlueGene family. If the nodes are taken as the latest available GPUs and are connected with the fast APEnet+ network, one can come up with a machine (provisionally code-named "Tamburo") having a limited number of nodes (from 7 to 15), but a peak performance of about 15 to 30 TFlops and a very reasonable price (of the order of 50-100 K€). An architecture of this kind looks perfectly adequate for the simulation of systems of biological relevance as neither a much too large band-width nor a too big local node memory is required. Naturally before launching a massive simulation program one has to port and optimize classical MD codes (like Gromacs) or ab-initio packages (like Quantum Espresso) on the new architecture. More information on this point is also given in WP4.

4.2 WP2 - Upgrade of the TheoPhys Cluster

The aim of this Work Package is to design, develop and deploy a new computational facility ("TheoCluster") that will provide resources for all the intermediate level HPC computing needs of the theoretical INFN community. The new facility will also provide an early "fast lane" to develop and test new algorithms and new codes before they can be moved lo large scale machines.

"TheoCluster" is seen as a successor to "TheoPhys" that has supported the "not high-end" computational demands (serial and medium-sized parallel applications) of the whole community for the period 2010-2012. In TheoPhys, the choice was made of using Grid services (provided and supported by the Italian Grid Infrastructure), for both serial and parallel applications This choice de facto integrated the Theophys cluster with other Grid resources used by the community for serial jobs, but raised several not yet resolved usability issues, especially for large parallel jobs.

WP2 plans to start from (and improve upon) the experience of the "TheoPhys" cluster, along the following lines.

- supporting the use of the cluster also for non-Grid users. This objective can be achieved through the exploitation of the new INFN Infrastructure for Authentication and Authorization (INFN-AAI). INFN-AAI was developed as a centralized access method mainly for Web and network resources, but it can be extended to meet our requirements.

- developing and maintaining the infrastructure needed to access the system according to the needs of its large users base; one aims at a horizontal use of the system through: (i) serial jobs submitted via Grid; (ii) serial jobs submitted by local queues; (iii) parallel jobs submitted via Grid and local queues; (iv) serial or parallel memory-bound jobs; (v) serial and possibly parallel jobs requiring special hardware resources.

- assembling and operating a new hardware platform, that uses state-of-the-art nodes based on commodity multi-core processors connected by a state-of-the-art new general purpose interconnection system. The structure of the computing nodes will satisfy the requirements of a large base of users (including for instance the nuclear physics community, for which a large amount of memory per core is mandatory).

- cooperating with other scientific communities with similar needs and similar computing infrastructures (starting with SISSA), developing and maintaining an interoperable infrastructure so that researchers of both institutions are able to access in a reasonably transparent mode either systems.

In the long term, we also plan to explore ways to integrate into a wider structure (beyond IGI and SISSA) , e.g., integrating the cluster into a larger interoperable network (including e.g. the systems available at CINECA) or within the PRACE infrastructure.

Workpackage description
The project will consist of the following tasks (each one with its own time line)

- Task 1 (month 0 – 6): implementation of additional new access methods (integration within the INFN AAI).
- Task 2 (month 0 – 24) : Software and queue optimization of the new cluster. This task will set up a sophisticated queuing system, based on some commercial package – e.g. LSF – able to control job submission by a fairly large set of authenticated users, assigning to them the appropriate resources in terms of computing nodes, memory and – at a later stage – specific hardware resources. An important part of this development is the support of an interoperability environment with the cluster at SISSA.
- Task 3 (month 1 – 12): Plan, procurement and deployment of the new cluster (TheoCluster)

Task 1: In order to provide more flexible access to the facility, we plan to use the INFN Authentication and Authorization Infrastructure (AAI) in order to give an additional access method – alongside with the Grid. This approach has already been tested for the TheoNuc facility.

INFN-AAI has been used so far by INFN to allow administrative access to employees and associates. TheoNuc has used this system in a scientific environment for the first time, allowing selected users to access the cluster from any location via `ssh`.

INFN-AAI has an IDP (Identity Provider) allowing web-based connections. Part of this project will be the creation of a Web based portal to facilitate access to the community. In addition the IDP will allow also federated access. In this mode users from other institutions, with appropriate agreements, will be authorized to use the facility. This is the case, for example, of the collaboration with SISSA.

As this task is completed, user will be able to access the computing resources through: (i) Grid services; (ii) direct interactive login; (iii) a federated prototype portal, providing Web access to users of at least two institutions (INFN and SISSA). Finally, user's support will be improved through the implementation of a Wiki manual and a WEB based ticketing system.

Task 2 This task aims to optimize the particular queuing system adopted for the cluster. Activities are addressed to study and realize the contemporary usage via Grid and via "local queue" in order to permit a complete usage of the resource pool. The coexistence of parallel jobs and high memory job on the same infrastructure can be complex to manage because there are problems due to the

integration and coexistence within the same system queue of different selection mechanisms. In particular we have to combine a queue where the selection is made by number of job slots with a queue characterized by memory bound. The final result can offer great flexibility of use, but implies a complex architecture to be studied in detail.

On a longer time frame, there is in also the need to support machines with GPU. The new LSF version (LSF 8) will include GPU support and GPU cores direct addressing will be possible. We intend to add this functionality to the new cluster facility so that GPU can be added to the cluster at a later stage.

Finally, we intend to achieve a synergy between our cluster and the other that will be shortly installed at the SISSA computing center, in Trieste. The latter will have 5000/10000 computing cores with 2 GB RAM each, and a peak performance of at least 100 TFlops. We are planning to make the queues of the two installations mutually visible via federated accesses. In this way INFN users can submit jobs to the SISSA cluster while SISSA users can submit jobs to TheoCluster. Since the two facilities will have the same Job Scheduler, we intend to explore the possibility to configure common queues and guarantee a balanced use of the resources of the two facilities. Providing a mechanism to make storage areas available to each other will be possible a complete share between the two sites. This would be a relatively simple task if both the sites are adopting the same file-systems (for instance GPFS). In case of different File System architectures, the integration will require more complex solutions to implement.

Task 3 The new cluster will be the result of the evolution of the existing cluster (Theophys) and the of TheoNuc facility. The current cluster, born in late 2009, consists of 1024 computing cores connected via an Infiniband network (Cisco SFS DDR IB switch 7024, 144 copper ports). The Cluster can be used serially (normal Grid usage) or in parallel (also through Grid). In early 2011, TheoNuc was added, dedicated to the computing requirements of the nuclear physics community, which needs a large amount of RAM per core.

Starting from these two computing resources we intend to acquire a new cluster that supports both features: fast network for parallel jobs and large amount of memory for CPU-bound jobs.

The project consists in the introduction of machines with 4 socket processors (32/64 cores) with large shared memory (128 GB per processor, 512 per node) with IB QDR / FDR and the associated network infrastructure. We plan to assemble 30-40 nodes, that is 2000-2500 cores, for a total peak performance of 20-25 TFlops. The estimated cost of each node is of the order of 10 K€. The estimated cost of the Infiniband network switch is of the order of 10 K€. This corresponds to a total cost of the infrastructure at the level of 300-350 K€, partly covered by funds made available by the CSN4.

4.3 WP3 - Assembly of a large prototype of a state-of-the-art parallel system

This workpackage aims to make available to the INFN computational community a state-of-the-art large computing systems that incorporates all the new building blocks for Exascale computing that a preliminary analysis – performed in the first year of the project – will have shown to be useful for performance for our set of programs.

We plan to have the new system available for physics-driven experimentation early in 2014, in order to have a competitive platform that the INFN groups may use in parallel with the resources available through PRACE. This time frame forces us to define the detailed structure of the system not later than fall 2013.

We have in mind a system with a peak performance well in excess of 200 TFlops, based on a set of standard high-end processors, combined with the most appropriate accelerators. Based on current wisdom, a first choice will have to be made between GPUs and MIC processors (or, in principle, a combination thereof). A second key choice will have to be made on the interconnection structure that should be present alongside a standard Infiniband structure; preliminary studies – at the hardware and software level – made in the first year of the project will be used to assess the performance advantages made possible by the use of toroidal interconnections.

We insist that our large prototype in assembled and installed by an industrial partner (to be chosen appropriately) that will handle all procurement phases, for which the project has no available man-power.

We plan to install the large prototype at CINECA, in the framework of a collaboration agreement between INFN and CINECA. The machine will be used for a large fraction of its available time by users directly associated to SUMA. We plan however to make the machine available to potentially interested CINECA users from other scientific communities, to help assess the impact of the new technologies. This policy, in our opinion, will also greatly add to the visibility of the SUMA project at the European level, and will be an important point in the establishment of a national ecosystem for the computational sciences.

4.4 WP4 - Technology developments

Work in this WP focuses on supporting an efficient use the APEnet+ network structure in conventional and unconventional setting and in further developing its architecture and its actual implementation.

In this framework, we plan to focus on the following tasks:

- Test and validation of the 3D torus in medium size system. Eurotech is currently installing a a medium size system that integrates a FPGA-based 3D Torus network, similar to the APEnet+ hardware but offering a dual PCIe Gen2 X8 interface to host processors. We will collaborate to port our APEnet+ hardware on the Eurotech system, having two goals in mind:

  1) Development of an APEnet+ based medium/large size test bed to verify RDMA programming model and effectiveness as well as robustness of selected routing algorithms for larger scale systems. Moreover, this large test-bed will be also used to test and verify future architectural improvements of the 3D Torus network;

  2) Evaluate the effect of host I/O bandwidth doubling on overall performances of the network.

- Experiment in the use of APEnet+ for unconventional interconnect topologies. For a non-homogeneous and/or long-range interacting system the domain decomposition technique – which is now standard practice for homogeneous and/or short-range problems – is almost useless. In this case,

the parallelization of a code consists in a balanced load of the degree of freedom among the node. The efficiency is strongly limited by communication latency and, overall, by network topology. The scaling with number of nodes is neither linear nor always monotonic (often bounded). In an attempt to overcome these problems, We propose to experiment with a compact, dedicated cluster (the "Tamburo machine", briefly referred to in a previous section) with a limited number of high-performance processors and the largest possible communication network connectivity, that uses GPUs for processing and the APEnet+ communication infrastructure. Using a structure of a regular graph for the interconnection structure, APEnet+ – having 6 links – fixes the degree of the graph. What is of utmost importance is to keep low the number of hops needed to move data across two non-adjacent nodes. One can easily see that one can build a system with 7 nodes if one wants to keep the number of hops to 1. If one accepts up to two hops, there are several ways to assemble up to 15 nodes. If each node uses a GPU with a peak performance of 2 TFlops, the complete system is a 30 TFlops machine that we plan to test with several programs in the realm of quantitative biology.

- Technological improvements of current 3D Torus architecture exploiting next generation 28nm FPGAs The SUMA Torus network will be an evolution of the APEnet+ architecture leveraging on technological advances of the 28nm next generation FPGA. These new components, recently introduced on the market, can be considered as a real alternative to ASIC development; they have several high speed (up to 28 Gbps) serial lines and multiple PCIe Gen3 hardware blocks, allows to reach an integrated device peak bandwidth of the order of Terabit/s. We will improve the PCIe interface of the network component, adopting the Gen3 protocol for the next device release. PCIe Gen3 is the adopted I/O standard protocol of the current GPU and INTEL MIC processors; it doubles the current PCIe performances adopting a more efficient data encoding scheme on the serial links (128b/130b instead of 8b/10b) and pushing the switching rate to 8 GT/s. We will test our new logic using hardware development kits and we will merge it, in a later stage, on dedicated custom board and/or on programmable components integrated in our partners next generation HPC system.

- Architectural improvements of Torus network using feedback from large-scale grand-challenge applications. Feedback from preliminary porting of INFN traditional and "non-traditional" applications on QUonG and the SUMA experimental computing platform will be used to design specific hardware/software blocks to improve applications sustained performances of SUMA machine. The huge amount of resources available in the 28nm FPGA components will be used to develop specific computational task accelerators either in the form of an ASIP (Application Specific IP) or as a custom hardware blocks. In the attempt to speed-up the RDMA and P2P protocol handling, we will explore the use of the dual core ARM9 microprocessor and we will develop a DMA gather-scatter engine able to autonomously gather data into packets. In addition we will provide the torus network controller with dedicated logic to improve the level of system

fault tolerance and to raise the reliability of large-scale systems. In particular, we will implement innovative "fault awareness" and "fault recovery" hardware/software mechanisms, distributed and hierarchical, proposed in the framework of the EURETILE project.

5 Project organization and management

SUMA starts with a core team of approximately 30 researchers coming from 10 Sezioni (or Gruppi Collegati). The team has an outstanding and track-record in several complementary relevant areas such ad theoretical computational physics, algorithm development, program development, computer architecture and electronics design. Even if the core team is scattered on a relatively large number of sites, many team members have long-standing collaborations among them, so the project critical mass is not weakened by geographic dispersion.

Equally important, the team enjoys long-standing collaborations at the national and international level with several large groups active in areas such as LQCD, fluid dynamics, complex systems, quantitative biology, advanced programming, architecture development and so on.

Each member of the core team will be able to work on SUMA for a fraction of time in the $20 \cdots 50\%$ range. The team will be enlarged by recruiting approximately ten post-docs ("assegni di ricerca") that we want to hire typically for a two-year term. These post-docs will be of course working 100% of their time on the project.

In the current preliminary phase, the project has a provisional steering-committee with one member for each participating Sezione. This will be changed as soon as the project structure has consolidated with a steering-committee whose members are the leaders of all Workpackages of the project. The steering-committee is in charge of managing the overall organization of the project, making sure that appropriate exchange of information and coordination is in place among the various lines. The steering-committee will typically meet once per week (via tele-conferences) and will have frequent contacts with the referees. We also plan to have several global collaboration meetings (e.g. one every three months) open to all members of the project.

We plan to collect all information available on the project on a WEB site (http://web2.infn.it/SUMA/), that we plan to have on line before the end of the current year.

6 Collaborative actions with other institutions

SUMA plans to work in close collaboration with many different institutions and research groups, in order to capitalize on all possible synergies and to favor the dissemination of all SUMA results.

At the formal level, SUMA will be the appropriate operational arm of a number of collaboration agreements already established between INFN and other institutions:

- A collaboration agreement between INFN and CINECA for common work in HPC computing. In the framework of this agreement INFN will have early access to the Eurora machine, currently being installed at CINECA

and to a small number of prototype systems available at CINECA and adopting MIC processors. In the same framework, CINECA will host and help operate the Large Prototype that SUMA plans to assemble.

- A collaboration agreement between INFN and SISSA. In the framework of this agreement INFN and SISSA will share the use of their middle-size clusters. SUMA will work to develop the system and queuing tools needed to support this federated access to both cluster.

- An agreement between INFN (through the COKA project) and the Supercomputing Center in Júlich will allow early access to MIC-based systems installed at Jülich for test and performance measurements.

At this point we are also discussing common work also with Fondazione Bruno Kessler (FBK) in Trento, as well as considering common work with the Interdisciplinary Center for Advanced Computing in Science and Tecnology (CAST) at the University of Tor Vergata.

On top of these formal agreement, SUMA will take advantage from close ties with several research projects working on similar directions; just to name a few, members of the SUMA team belong to the European Twisted Mass Collaboration (ETMC), the Coordinated Lattice Simulations (CLS) initiative, the QPACE2 project the FP7 STRONGnet network and the FP7 EURETILE project. Strong links are already in place with the DEEP project at Jülich and with the NEPAL project at Orsay.

7 Project costs

Table 1 is a very preliminary attempt to suggest a breakup of the SUMA budget. We expect that this breakup will have to be adjusted as the structure of the project becomes more stable, and costs for specific items are more accurately estimated.

Short comments to each item of the SUMA budget are as follows:

- Personnel: We plan to open 10 - 12 post-doc positions ("assegni di ricerca"). Most of these positions will be for two years – starting at appropriate time points during the first year of the project; two years is a long enough appointment to make sure that the new collaborator is able to give a significant contribution to the project and to make the position interesting enough at the European level. For the same reason, we anticipate a typical budget at $\simeq$ 30000 €/ year for each position. Several collaborating institutions may be available to co-finance at least some of these positions.

- Running costs: The project will build on the work of a team of approximately 25 people, half of them working part time on the project and half of them working full time. We envisage a conservative figure for running costs (travels, conference fees, ...) at the level of $\simeq$ 2500 €per person per year, corresponding to $\simeq$ 180000 €for the whole project.

- Cluster improvements: A preliminary estimate of the costs associated to the cluster upgrade in Pisa is around 300000 €. CSN4 has agreed to support this development with 80000 € from the 2013 budget, so we earmark

Personnel	600K€
Large Proto	600K€
Cluster upgrade	220K€
Adv'd developments	325K€
Running costs	180K€
Grand Total	**1925K€**

Table 1: Preliminary breakup of project costs

a contribution of 220000 € for this item; Note that CSN4 will support the operation costs of the cluster with an additional yearly budget of 20000 €.

- Large prototype: Cost estimates for this and the following item are still very preliminary. As a general rule, we would like to earmark as much resources as possible for these two areas of development. At a preliminary level, based on reasonable extrapolation of current costs, we expect that an investment of $\simeq 600000$ € will be sufficient for a large prototype system with a peak performance of at least 200 TFlops.
- Advanced Developments: At this point in time, we anticipate an early investment of $50000-60000$ € for the development of the Drum prototype, leaving of the order of 250000 € for further advanced developments, to be discussed in details at a later stage.

Appendix A: Physics motivations

SUMA gathers groups active in different fields of theoretical physics for which significant computational support is required. Distinguished expertise is ensured on a wide spectrum of interests which we sum up in the following. These application areas represent a significant subset of the applications that use a major fraction of the computational resources of the large European computing centers. Being this spectrum so wide, completeness is by no means our goal. Ideally, we would like the reader to take a glance at some windows of opportunities we can see open for our researches if we can effectively profit from current progress in High Performance Computing. We are confident that this will offer a wide assessment of the potential of the new architectures we are going to test not only for the INFN theoretical community but also for a broader audience.

Lattice QCD

The final goal of Particle Physics is to cast into a coherent fundamental theory the existing experimental data and the forthcoming results expected from high energy experiments (at the moment, in primis from the LHC). This opens a theoretical scenario in which strongly-interacting theories play a major role, both in the Standard Model and in any model beyond it. Numerical simulations represent a unique tool to systematically investigate features which, being

of non-perturbative nature, are not computable by presently known analytic techniques. Lattice gauge theory is a successful proof of the viability of such a program: the dynamics of strongly-interacting field theories can be studied by discretizing them on a four-dimensional lattice with finite lattice- spacing and volume. The functional integrals are computed non-perturbatively via Monte Carlo simulations, and the results are extrapolated to the infinite-volume and continuum limits. This has to be done while physical quantities are constrained to their actual values during this renormalization process (the quest for physical, light values of the pion mass is the prototypical example of such a programme). If at all stages of the computation one can keep track of the statistical and systematic errors, fundamental questions in Quantum Chromodynamics (QCD, the Quantum Gauge Theory describing strong interactions) can be addressed from first principles. Lattice QCD has been able to provide insight into issues like color confinement or the QCD phase diagram; it has also provided values for quantities like quark masses, decay constants or form factors which in turn constrain our understanding of the Standard Model on its own and possibly as a part of a broader, more fundamental theory.

In the picture we sketched, numerical effort is clearly a big issue: namely, Lattice QCD has been over the years one of the major research interest motivating massively parallel computing, and it is widely recognized that the Lattice QCD community has been quite successful in the field. All in all, Lattice QCD practitioners have been among the best in getting high efficiency. This is to a large extent due to the problem itself. Given the local nature of QCD as a quantum field theory, updating the local (sub)lattice residing on each computing node is a computational task which only requires communications with nearest neighbors nodes. On top of this bonus, much more has been achieved out of constant, coordinated efforts and innovative ideas. By now many generations of parallel computers have been developed having the inversion of the Dirac operator as one of the major benchmarks. On the other hand, algorithmic developments have been constant: the progress of the field would have been much slower if a constantly increasing computational power had been the only resource at disposal. We also remark that in the near future a specific effort will be necessary to make it possible to evaluate hadronic matrix elements with order percent accuracy, in order to spot possible signals of physics beyond the Standard Model at forthcoming flavour factories.

SUMA and large international Lattice projects

There are reasons for Lattice QCD practitioners to claim the field has become a mature one. One of this is the presence of well established, large international collaborations. Not only these address specific aspects of LQCD, thus providing new insights on particular subjects. More noticeably, they build on theoretical approaches which are different in terms of the formulation of the theory itself, *e.g.* different regularizations are used of both fermion and gluon degrees of freedom, aiming at clear understanding and control on systematics. Also, these collaborations have over the years developed different algorithmic solutions. As a result, this is a time in which we can have many cross-checks of results coming from independent sources, all taking a fairly good control on different sources

of systematics. Each collaboration has developed a large set of codes for the generation of gauge configurations and for physics analysis; these codes can be typically adjusted to suit several available computing platform. We regard as a valuable opportunity the fact that members of the SUMA team belong to two major collaborations, both in a position to adapt, optimize and test their own suites for the new available architectures.

SUMA and CLS

The group in Milano Bicocca belongs to the Coordinated Lattice Simulations (CLS) initiative [16], in which major lattice groups in Europe use the same code and coordinate the generation of the gauge configurations, to be shared for a variety of lattice projects. Several years ago, CLS embarked on a large-scale program aimed at producing a set of high-quality ensembles which enable full control over (i) discretisation errors, (ii) extrapolations in the light quark mass to the physical pion mass, and, eventually, (iii) the influence of dynamical strange and charm quarks on hadron observables.

All simulations performed so far are based on two flavours of non-perturbatively improved Wilson fermions [17, 18, 19, 20, 21]. The large set of lattices already simulated allows CLS groups to compute several interesting quantities, to extrapolate them to the chiral, infinite-volume and continuum limits. More simulations will be carried out when needed to control the systematics in the various observables considered.

In the next few years the most computational-intense activity will be the generation and the analysis of a new set of lattices with a larger number of dynamical flavours of non perturbatively improved Wilson fermions. The aim is to simulate pions near or at the physical point with lattice spacing small enough and volumes large enough to be able to reach the infinite volume and continuum limit with confidence.

CLS algorithmic environment has several distinguished features. Fig. 1 enlights the very good scaling properties of the mass-preconditioned Hybrid Monte Carlo (MP-HMC) algorithm (results are depicted of a detailed scaling analysis of the code on 4096 to 16384 cores on a Blue-Gene/P). The particular feature of this implementation is the use of a locally deflated solver for the Dirac equation [22, 23, 24], which has proven to be very efficient in the region of the small quark masses one is aiming at.

The generation of the ensembles with a larger number of flavours will make use of Lüscher's publicly available openQCD HMC code [7, 8]. Numerical efficiency is ensured by techniques such as the a locally deflated solver, twisted-mass re-weighting, frequency splitting of the quark determinant through mass shifts. Analogous good scaling behavior is expected for the openQCD HMC code

SUMA and ETMC

Since several years the group in Roma Tre has been deeply involved in the initiatives of the European Twisted-Mass Collaboration (ETMC), which gathers

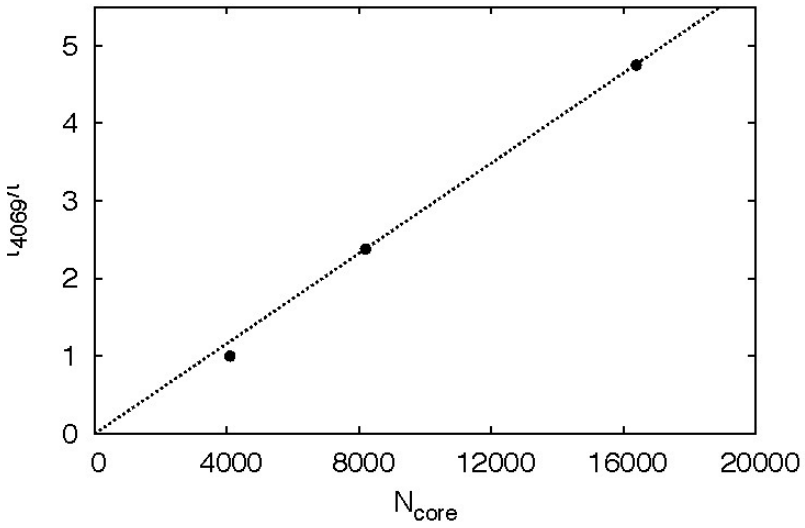

Figure 1: Scaling behavior of the MP-HMC algorithm with the number of cores on Blue-Gene/P (from [25])

more than 30 permanent researchers and young students belonging to many institutions all over the Europe.

In the last few years the ETMC has produced a huge amount of high-quality ensembles of gauge configurations with two ($N_f = 2$ [1, 2]) as well as four ($N_f = 2+1+1$ [3, 4]) flavors of dynamical fermions. In the latter case, besides two light mass-degenerate quarks, the strange and charm quarks with masses close to their physical values are included in the sea. Such a setup is the closest one to the real word, adopted till now only by ETMC in Europe and by MILC in the USA.

Other important features of the gauge ensembles produced by ETMC are that several lattice spacings and volumes have been used and that the simulated pion mass ranges from ~ 250 MeV to ~ 600 MeV. This has allowed to achieve a nice control over the uncertainties induced by finite size effects and by the extrapolation to the physical pion mass (~ 140 MeV) as well as to perform the continuum limit using at least three or more lattice spacings. Moreover, adopting maximally twisted Wilson fermions [5, 6] the renormalization properties of several observables are greatly simplified by the automatic $\mathcal{O}(a)$ improvement of bi-linear operators as well as by the continuum-like pattern of the mixing among more complicated operators, like the four-fermion operators relevant in the oscillation of the neutral $K - \overline{K}$ system.

The group in Roma Tre has participated actively to the calculations and to the analysis of 2-point and 3-point correlation functions of pseudo-scalar mesons from which several important studies of light-flavor physics [27]-[28] as well as of charm and beauty physics [29, 30, 31, 32] have been completed. Most of the ETMC results are included in the recent lattice averages performed by the Flavor Lattice Averaging Group [34]. The current work on isospin break-

ing by the Roma123 collaboration has also made extensive use of the ETMC configurations[33].

Such results could not have been accomplished without the computing time made available on several HPC systems in Europe by the PRACE infrastructure and at the national level by ISCRA. In this respect the group in Roma Tre has been the principal investigator of two PRACE projects, namely the project PRA027 entitled "QCD simulations for flavor physics in the Standard Model and beyond" (35 millions of core-hours at the BG/P system in Jülich from December 2010 to March 2011) and the project PRA067 entitled "First Lattice QCD study of B-physics with four flavors of dynamical quarks" (30 millions of core-hours at the BG/Q system in Jülich and 30 millions of core-hours at the BG/Q system at CINECA from May 2012 to April 2013), and of two ISCRA projects, namely the class A project HP10A7IBG7 "A New Approach to B-Physics on Current Lattices" (10 millions of core-hours at the BG/P system in CINECA) and the class C project HP10CJTSNF "Lattice QCD Study of B-Physics" (5 millions of core-hours at the BG/Q system in CINECA).

Thanks to the experience made during the last years through the PRACE and ISCRA projects, ETMC software suite [35] has been largely optimized for the BG/P system and is currently under optimization for the BG/Q system. The scaling of the code with the number of cores is close to be optimal, as shown in Figs. 2-3.

For the next years the ETMC plans to generate new ensembles of gauge configurations with $N_f = 2 + 1 + 1$ and pions near or at the physical point. The aim is to perform with these ensembles the calculations of all the correlation functions relevant for light and heavy meson physics removing the need for a long extrapolation of the light-quark mass to the physical point. They expect to carry out such a program using the HPC resources made available through PRACE including the new BG/Q installation Fermi at CINECA. Nevertheless the search for new powerful algorithms will continue. The experimental machines provided by the SUMA project will be yet another chance to to test ideas and tools to address faster computations.

Finite temperature and density Lattice QCD

The understanding of the phase diagram of QCD in the temperature - baryon density plane is an open issue. At a critical temperature of the order of 200 MeV a transition takes place between the ordinary hadron phase and a deconfined phase of quark gluon plasma; in the same range of temperatures the weakly broken chiral symmetry gets restored. This rich scenario has certainly been in place (with a non null barion density) in the early stages of the universe and heavy ions collisions experiments aim at recreating this primordial plasma, to address fundamental questions like baryon asymmetry or the seeding for structure formation. On another hand, the astrophysics of compact objects is *per se* a convincing proof of the relevance of understanding the phase diagram over a wide region of the temperature-density plane. Finite Temperature QCD has extensively been a playground for Lattice QCD for a while, but severe limitations to finite density simulations come from the so-called "sign problem": when a finite chemical potential is switched on, one can no longer read a positive measure in the functional integral, so that a straightforward path to Monte Carlo simulations is disrupted. One way out is the analytic continuation from

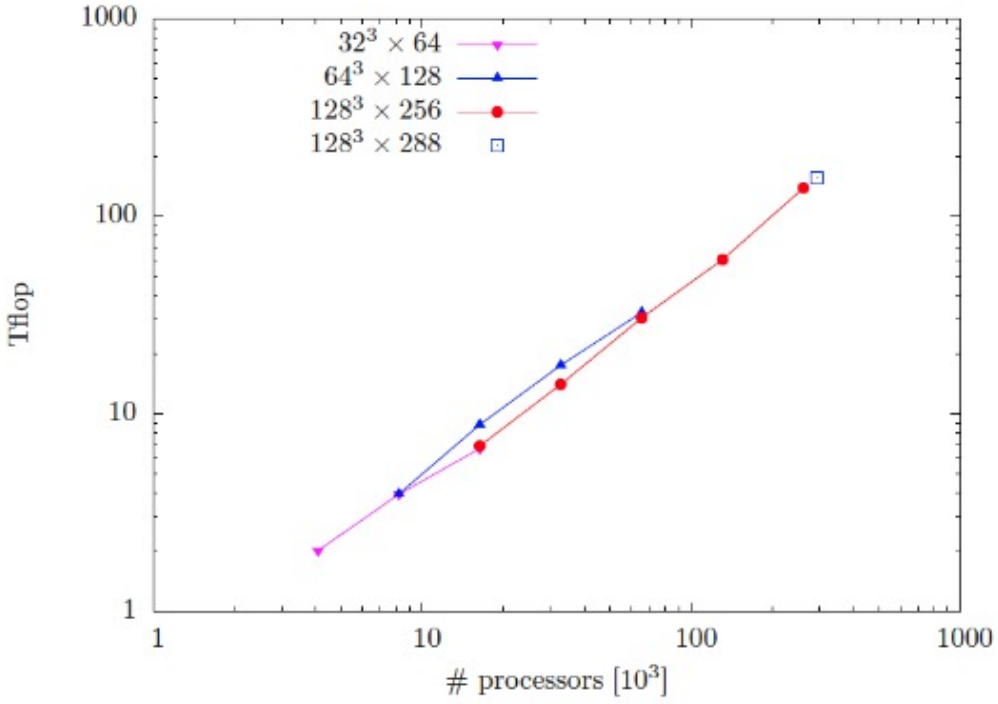

Figure 2: Performance in TFlops of the tmLQCD code [35] achieved on the BG/P versus the number of processors. The open square corresponds to the use of the full 72-rack installation at Jülich.

imaginary chemical potential, studied extensively also by INFN collaborations. Although there is substantial consensus on the position of the transition line at small baryon density, several open questions remain: the nature of the deconfinement transition, its relationship with the chiral transition, the presence of critical points at real or imaginary chemical potential and their influence on the high-temperature deconfined state, the role of an external electro-magnetic or chromo-magnetic field. On top of these, all the well-known (technical and fundamental) issues which are in place in zero temperature Lattice QCD are more and more entering the scene to pursue a higher level of accuracy: continuum and infinite volume limits, physical values for the quark masses, relevance of extra flavors on top of the two light ones, pros and contras of different fermion regularizations.

Because of technical reasons (mainly, the lattice sizes which are in place), finite temperature QCD has been a (maybe even the) major playground for GPU applications. The Lattice QCD community pioneered the GPGPU paradigm (General Purpose computations on GPU), which eventually resulted in a quite stable programming environment following the introduction of CUDA (Compute Unified Device Architecture [36]) by NVIDIA. The introduction of GPUs in lattice QCD calculations started with the seminal work of Ref. [37], in which

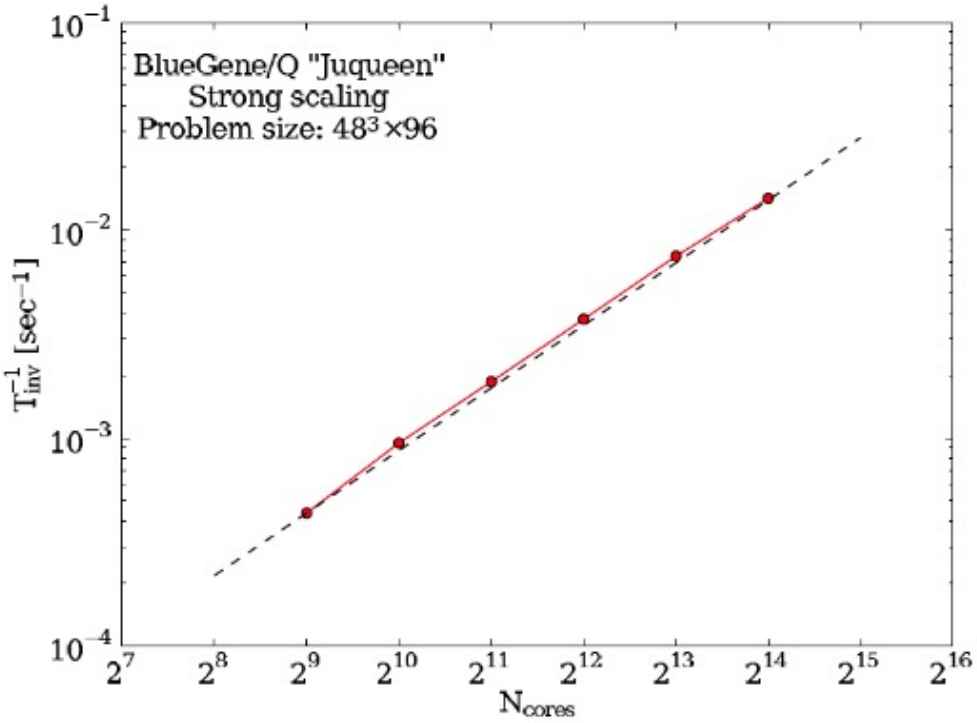

Figure 3: The inverse time to solve a problem of size $48^3 \times 96$ on the BG/Q machine versus the number of cores used. The nice scaling behavior demonstrates that the tmLQCD package [35] can use the BG/Q machine in a massively parallelized mode.

native graphics APIs were used. By now multiGPU simulations are the real challenge. This yet another example of a common trend: the big technical issue is how to deal with the multiple levels of parallelism which are in place on nowadays computer architectures. Solutions will come out of coordinating computer architecture and algorithmic progresses. It is interesting to note that a solution like domain decomposition has proved quite effective both on more traditional architecture [22] and on multiGPU systems [38], even of the order of 100 GPUs (roughly speaking: communications are the weak part of the computation, but there are frameworks in which these are *per se* reduced).

Finite temperature QCD has been a playground for GPUs within INFN as well, the Pisa group having gained remarkable insight and experience. In Ref. [39] one of the first examples of Lattice QCD code was presented where most computations are actually performed on the GPU. In particular, in the implementation of the Rational Hybrid Monte Carlo algorithm [40], the whole molecular dynamics evolution of gauge fields and momenta, which is the most costly part of the Hybrid Monte Carlo algorithm, runs completely on the GPU, thus reducing the costly CPU-GPU communications (going through the PCI-express channel) at the minimum. One still needs a CPU to run the main program, but mostly in the role of a mere controller of the GPU instruction flow.

The present single GPU implementation makes use of NVIDIA's CUDA platform together with a standard C++ serial control program running on the CPU. It has been used in production runs on C1060 and Fermi architectures, to explore the properties of the QCD vacuum and of its phase diagram [41, 42]; preliminary test have been performed on Kepler cards and specific improvements to optimize the code for such architectures are on the way. Sustained performances are close to 10 % of the peak card performances; while this figure could seem rather modest, it results in quite relevant boost factors with respect to traditional C++ codes running on a single CPU core.

A step further will focus on portability to different many-core architectures experimenting with an OpenCL (instead of CUDA) implementation of the GPU kernels, and assessing its relative performance first on GPUs and them – if appropriate tools are available – on other many core systems. The other compelling direction to go is that of LQCD simulations on multiGPU systems.

Lattice QCD Miscellanea

There are Lattice QCD research lines in the SUMA team besides the ones we have mentioned above and that are in a sense a bit out of the mainstream, either being a bit more topical or being freshly new ideas: we expect valuable add-on contributions from both.

One example of the first kind is Numerical Stochastic Perturbation Theory (NSPT), which has been the main research interest in Lattice Gauge Theories in Parma for many years [43]. The lattice regularization is a hard one when one has to tackle Perturbation Theory (PT), which is of course a valuable tool also on the lattice, like in any formulation of a field theory. NSPT is a numerical implementation of the Stochastic Perturbation Theory which stems from Stochastic Quantization, thus avoiding the cumbersome nature of lattice PT diagramatics. NSPT enables high orders computations, and the main efforts in recent years have been dedicated to the three-loops computation of renormalization constants, which can be compared to non-perturbative results for a better control of systematics [44], in particular in the extraction of continuum and infinite-volume limits. Numerical efficiency has always been a goal in this context, resulting in expertize which has been in recent years spent in the quest for efficiency for (more standard) non-perturbative Lattice QCD [45].

Developing new ideas and tools to address theoretical questions which may unveil a deeper understanding of the strong interactions can often allow for faster computations as well. Taming lattice artifacts and effectively extracting continuum physics is and example of this type. In recent years, an alternative to the more standard approaches of Symanzik improvement or perfect actions is being investigated in Milano Bicocca. The formulation provided by topological lattice actions has potentially the appealing advantages of simplicity and easy-to-use applicability [46]. The results obtained for the two-dimensional O(N) spin models show that the lattice discretization effects have been removed within the statistical error bars. This clearly motivates the investigation of the reliability and the efficiency of this new approach to non-Abelian gauge theories like the Yang-Mills theory or QCD.

Lattice Boltzmann Methods

Computational techniques are absolutely necessary to compute reliable solutions to the highly non- linear equations of motion of Fluid-dynamics in the regimes that are interesting for physics or engineering. Over the years, many different numerical approaches have been theoretically developed and their implementation has been a useful testbed to benchmark many state-of-the-art massively parallel computers. Lattice Boltzmann methods (LBM) [47] are a one of these techniques, on which there is distinguished expertize both in TorVergata and in Ferrara. LBM are a flexible and powerful approach, able to cope with many different fluid equations (e.g., multiphase, multicomponent and thermal fluids) and to handle complex geometries or boundary conditions. LBM builds on the fact that the details of the interaction among the fluid components at microscopic level do not change the structure of the equations of motion at the macroscopic level, but only modulate the values of their parameters. What is actually described (and put on the computer) is some simple synthetic dynamics of fictitious particles that evolve explicitly in time and, appropriately averaged, provide the correct values of the macroscopic quantities of the flow. Members of the SUMA team have a strong background in the development [48, 49] and computational optimization of these approaches [50]

LBM formalism results in a dramatic reduction of the degrees of freedom associated with the velocity space, with a consequent computational boost. In particular, the low communication/computation ratio makes LBM nearly ideally amenable to parallel computing. Besides the natural realm of macroscopic fluid dynamics, the LBM portfolio of applications keeps expanding across scales of motion, particularly towards micro- and nanofluidics with active thermal dynamics. Realizing the full potential of the method in the latter case, however, still raises many challenges.

The rheology of flowing soft systems, such as emulsions, foams, gels, colloidal glasses and related complex fluids, is a subject of increasing importance in modern non- equilibrium thermodynamics, with a broad range of applications in fluid dynamics, chemistry and biology. From the theoretical standpoint, flowing soft systems are challenging because they do not fall within any of three conventional states of matter, gas-liquid-solid, but live rather on a moving border among them. Foams are typically a mixture of gas and liquids, whose properties can change dramatically with the changing proportion of the two; wet-foams can flow almost like a liquid, whereas dry-foams may conform to regular patterns, exhibiting a solid-like behavior. Emulsions can be paralleled to bi-liquid foams, with the minority species dispersed in the dominant (continuous) one. The behavior and, to same extent, the existence itself of both foams and emulsions are vitally dependent on surface tension, Living, as they do, out of equilibrium, these materials exhibit a number of distinctive features, such as long-time relaxation, anomalous viscosity, aging behavior, which necessitate profound extensions of non-equilibrium thermodynamics. Owing to this complexity, and particularly the slow relaxation properties, the investigation of the dynamical behavior of these systems requires very long time integrations, typically of the order of tens of million of time steps (as a reference, one LB time step can be taken of the

order of 100 1000 molecular dynamics time-steps). As a result, the need of a very-long simulation span, sets a strong incentive for efficient implementations.

Quantitative Biology

The group in TorVergata focuses its work on quantitative biology. The investigation of the structural properties of bio-molecules (proteins, DNA, RNA) is one of the big challenges in Computational Physics. Quantitative Biology has by now settled as a discipline on its own; it has also shown to be one of most prominent example of a research area interconnected with a variety of other. The level of complexity of the typical problem is huge: huge number of degrees of freedom, hugely different time scales involved in a given problem, huge sensitivity to external/environmental conditions.

The overall numerical apparatus is mainly that of numerical condensed matter, which is worth to briefly recall. The scientific community in condensed matter simulations is currently relying on 5 different approaches, based on 5 kinds of approximations, each of which is roughly as popular as the others.

- Coarse-grain models for molecular assemblies

  Simple analytical models have been designed for pair-interactions among sites representing nano-scale molecular aggregates. The prediction of static properties of these models are the object of investigation, with properties ranging from crystallization to chromatin structure in cell nuclei. As an example in this category, a typical task is the construction of conformational maps in chromatin models.

- Empirical models for molecular assemblies

  There are two popular packages for Molecular Dynamics (MD) of large assemblies: GROMACS and NAMD. They use two different strategies of parallelization. GROMACS (http://www.gromacs.org/) uses domain decomposition of the system and eight shell method for assigning the forces. The replica-exchange MD can be simulated with Gromacs within the standard Gromacs package.

  The NAMD approach is based on domain decomposition of pair-interactions. The algorithm allows the MD simulation of samples containing up to 1 M atoms and sizes in the range of the μm. Almost independent replica, that are the basis of replica-exchange MD and metadynamics, can be simulated in the same generalized statistical ensemble (use is made of scripting languages like TCL, PERL, PYTHON).

- Semi-empirical models for molecular assemblies

  In the tight-binding approximation one can perform the computation of atomic forces in the ground electronic state of systems composed of up to 10K atoms. The most stable method to perform this task is the self-consistent charge iterative method which is implemented in the Density-functional tight-binding package (http://www.dftb-plus.info). The code uses different levels of parallel distribution for the diagonalization of large sparse matrices. The distribution of tasks among non-synchronous and

partially independent replica can be performed using scripting languages (PERL or PYTHON).

- Density functional theory for molecular assemblies Among the most popular codes in this context, the Quantum-Espresso (http://www.quantum-espresso.org) package is the one which best fits parallel computing; in particular, it can be highly efficient under different parallelization strategies. Depending on the HPC architecture, these can be used all together, allowing the access to statistical properties of large molecular assemblies, which can be accurately described. Parallel tasks, with different levels of communications within tasks, are used for different computations:

  1. Diagonalization of large and sparse matrices.

  2. Matrix and array construction (distribution of K-points and Kohn-Sham states).

  3. Fourier transforms.

  4. Three-dimensional grids of different densities and spatial resolutions.

  5. Statistical algorithms making use of largely independent replica in the configuration and momentum spaces.

- All-electrons models for molecular assemblies Quantum Monte Carlo methods can be used to compute the statistics of electrons in confined systems, like molecules and nano-scale samples of liquids and solids. They are by now as accurate as post-Hartree Fock methods in computational quantum chemistry, or Coupled-cluster (CC) methods. Two implementations that have been shown to scale up to 1M cores are the QMC packages distributed by the University of Toulouse (http://qmcchem.ups-tlse.fr) and by the University of Cambridge (http://www.tcm.phy.cam.ac.uk/ mdt26/casino2.html). The computational architectures are in both cases well suited for almost independent tasks.

We finally mention in this area Large Scale Cortical Simulations. In the last decade, a sequel of experimental, theoretical and modeling novelties converged to create, for the first time, the possibility to exploit HPC to create large scale simulation of cortical activities and compare the result with non-invasive experiments [51, 52].

Since 2011 the FP7 EURETILE PROJECT [53] has been working on the development of the DPSNN benchmark suite, a natively Distributed Polychronous Spiking Neural Network code. A first version of this neural benchmark is foreseen for December 2012.

The development of the cortical simulation benchmark is a grand-challenge of parallel/distributed programming, an essential tool for Computational Neuroscience and for Cognitive Computing, as well as a source of requirements and architectural inspiration toward extreme parallel and distributed architectures and – last but not least – a tool to deliver simulation result of great scientific interest. In the framework of SUMA, the DPSNN code will be further extended by the group in Roma I to run efficiently on the SUMA architecture(s) and to identify areas of improvement for HPC architectures.

Numerical Relativity

The direct detection of Gravitational Waves (GWs) is one of the great physics challenges of the next years, an effort in which INFN is playing a major role through the (LIGO/)VIRGO research programs. A successful detection can only take place if for the candidate events resulting in the emission of GWs one has accurate knowledge of the signal. This is a typical result that has to come from Numerical Relativity.

There is a powerful community computational environment to support Computational Relativistic Astrophysics, and INFN researchers in the field (in particular those in Parma) are deeply concerned with the Einstein Toolkit [54] and, in particular, Cactus/Whisky [55]. These applications have proved to scale well up to 16K MPI processes on the Blue Gene/Q CINECA Fermi system.

A typical candidate for the emission of GWs is the merger of two Neutron Stars, for whose modeling the dynamics of the magnetic and gravitational fields plays a major role, as it also holds true with respect to the physical properties of isolated Magnetized Neutron Stars. There is a general consensus that strong magnetic fields and a fast rotating compact object are at the origin of many high-energy phenomena, most notably Gamma-Ray Bursts (GRBs). In particular the recently developed millisecond-magnetar model [56] has suggested that such violent events might be the signature of a rapidly rotating and highly magnetized neutron star (NS). However the origin of this strong and ordered magnetic field remains poorly understood. In particular, the environment in which these magnetic fields are supposed to arise is often characterized by turbulence, instabilities and convection. All these are topics Numerical Relativistic Magneto Hydro-Dynamics is concerned with.

It is clear that the intrinsic multidimensional and non-linear nature of the problem prevents simple analytical treatments from going beyond a general idea or proof of principle. Testable predictions, in view also of the quality of present and upcoming observational data, demand a higher level of accuracy in the modeling. That demands to perform high-resolution 3D simulations in full General Relativity to study the effects that magnetic fields may have on the dynamics of magnetized differentially rotating neutron stars. Such objects are likely to form during core collapse of massive stars [57], or as the result of binary neutron stars coalescence [58], or the accretion induced collapse of a white dwarf [59].

To do meaningful simulations one has to evolve in time (a usual Methods of Line integration on time of the variables requires to save 3 time levels) 10 metric variables, 6 extrinsic curvature variables, 5 hydro-dynamical variables (+5 conservative counterparts) and 3 magnetic field variables that on the 1000^3 grid (needed to cover the star interior at the desired accuracy) amount to 0.65 TBytes of allocated physical variables.

References

[1] P. Boucaud, *et al.* [ETM Collaboration], Phys. Lett. B **650** (2007) 304 [hep-lat/0701012].

[2] P. Boucaud, *et al.* [ETM Collaboration], Comput. Phys. Commun. **179** (2008) 695 [arXiv:0803.0224 [hep-lat]].

[3] R. Baron, P. .Boucaud, J. Carbonell, A. Deuzeman, V. Drach, F. Farchioni, V. Gimenez and G. Herdoiza *et al.*, JHEP **1006** (2010) 111 [arXiv:1004.5284 [hep-lat]].

[4] R. Baron *et al.* [ETM Collaboration], Comput. Phys. Commun. **182** (2011) 299 [arXiv:1005.2042 [hep-lat]].

[5] R. Frezzotti and G. C. Rossi, JHEP **0408** (2004) 007 [hep-lat/0306014].

[6] R. Frezzotti and G. C. Rossi, JHEP **0410** (2004) 070 [hep-lat/0407002].

[7] `http://luscher.web.cern.ch/luscher/openQCD/index.html`

[8] M. Luscher and S. Schaefer, [arXiv:1206.2809 [hep-lat]].

[9] M. Guagnelli, F. Palombi, R. Petronzio and N. Tantalo, Phys. Lett. B **546** (2002) 237 [hep-lat/0206023].

[10] G. M. de Divitiis, P. Dimopoulos, R. Frezzotti, V. Lubicz, G. Martinelli, R. Petronzio, G. C. Rossi and F. Sanfilippo *et al.*, JHEP **1204** (2012) 124 [arXiv:1110.6294 [hep-lat]].

[11] G. M. de Divitiis, R. Petronzio and N. Tantalo, Phys. Lett. B **718** (2012) 589 [arXiv:1208.5914 [hep-lat]].

[12] S. Morante, Curr. Alz. Res. **5** (2008) 508.

[13] V. Minicozzi, F. Stellato, M. Comai, M. Dalla Serra, C. Potrich, W. Meyer-Klaucke, S. Morante, J Biol Chem 283 (2008) 10784.

[14] P. Giannozzi , K. Jansen, G. La Penna, V. Minicozzi, S. Morante, G.C. Rossi, F. Stellato, Metallomics **4** (2012) 156.

[15] F. Guerrieri, V. Minicozzi, S. Morante, G.C. Rossi, S. Furlan, G. La Penna, J. Biol. Inorg. Chem. **14** (2009) 361.

[16] `https://twiki.cern.ch/twiki/bin/view/CLS/WebHome`

[17] M. Lüscher, Comput. Phys. Commun. **165** (2005) 199 [hep-lat/0409106].

[18] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio and N. Tantalo, JHEP **0702** (2007) 056 [hep-lat/0610059].

[19] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio and N. Tantalo, JHEP **0702** (2007) 082 [hep-lat/0701009].

[20] M. Marinkovic and S. Schaefer, PoS LATTICE **2010** (2010) 031 [arXiv:1011.0911 [hep-lat]].

[21] P. Fritzsch, F. Knechtli, B. Leder, M. Marinkovic, S. Schaefer, R. Sommer and F. Virotta, Nucl. Phys. B **865** (2012) 397 [arXiv:1205.5380 [hep-lat]].

[22] M. Lüscher, Comput. Phys. Commun. **156** (2004) 209 [hep-lat/0310048].

[23] M. Lüscher, JHEP **0707** (2007) 081 [arXiv:0706.2298 [hep-lat]].

[24] M. Lüscher, JHEP **0712** (2007) 011 [arXiv:0710.5417 [hep-lat]].

[25] M. Marinkovic, S. Schaefer, R. Sommer and F. Virotta, PoS LATTICE **2011** (2011) 232, and M. Marinkovic, PhD Thesis.

[26] `http://luscher.web.cern.ch/luscher/DD-HMC/index.html`

[27] R. Frezzotti *et al.* [ETM Collaboration], Phys. Rev. D **79** (2009) 074506 [arXiv:0812.4042 [hep-lat]].

[28] V. Bertone, N. Carrasco, M. Ciuchini, P. Dimopoulos, R. Frezzotti, V. Gimenez, V. Lubicz and G. Martinelli *et al.*, arXiv:1207.1287 [hep-lat].

[29] B. Blossier *et al.* [ETM Collaboration], JHEP **0907** (2009) 043 [arXiv:0904.0954 [hep-lat]].

[30] B. Blossier *et al.* [ETM Collaboration], Phys. Rev. D **82** (2010) 114513 [arXiv:1010.3659 [hep-lat]].

[31] B. Blossier *et al.* [ETM Collaboration], JHEP **1004** (2010) 049 [arXiv:0909.3187 [hep-lat]].

[32] P. Dimopoulos *et al.* [ETM Collaboration], JHEP **1201** (2012) 046 [arXiv:1107.1441 [hep-lat]].

[33] http://web.infn.it/CSN4/index.php/fenomenologia-delle-particelle-elementari

[34] G. Colangelo, S. Durr, A. Juttner, L. Lellouch, H. Leutwyler, V. Lubicz, S. Necco and C. T. Sachrajda *et al.*, Eur. Phys. J. C **71** (2011) 1695 [arXiv:1011.4408 [hep-lat]].

[35] K. Jansen, C. Urbach, Comput. Phys. Commun. **180** (2009) 2717 [arXiv:0905.3331].

[36] NVIDIA Corporation, `http://developer.nvidia.com/object/gpucomputing.html`

[37] G. I. Egri, Z. Fodor, C. Hoelbling, S. D. Katz, D. Nogradi, K. K. Szabo Comput. Phys. Commun. **177** (2007) 631.

[38] R. Babich, M. A. Clark, B. Joo, G. Shi, R. C. Brower and S. Gottlieb, arXiv:1109.2935 [hep-lat].

[39] C. Bonati, G. Cossu, M. D'Elia and P. Incardona, Comput. Phys. Commun. **183** (2012) 853.

[40] A. D. Kennedy, I. Horvath, S. Sint *A New exact method for dynamical fermion computations with nonlocal actions.* Nucl. Phys. Proc. Suppl. **73** (1999) 834.

[41] C. Bonati, G. Cossu, M. D'Elia, F. Sanfilippo *The Roberge-Weiss endpoint in $N_f = 2$ QCD.* Phys. Rev. D **83** (2011) 054505.

[42] M. D'Elia, M. Mariti and F. Negro, arXiv:1209.0722 [hep-lat].

[43] F. Di Renzo and L. Scorzato, JHEP **0410** (2004) 073 [hep-lat/0410010].

[44] M. Hasegawa, M. Brambilla and F. Di Renzo PoS LATTICE **2012** 240 (2012).

[45] M. Brambilla and F. Di Renzo PoS LATTICE **2012** 192 (2012).

[46] M. Bogli, F. Niedermayer, M. Pepe, U.J. Wiese (Bern U.), JHEP 1204 (2012) 117.

[47] S. Succi, *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond*, Oxford University Press (2001).

[48] M. Sbragaglia et al., J. Fluid Mech., **628** (2009) 299.

[49] A. Scagliarini et al., Phys. Fluids, **22** (2010) 055101.

[50] L. Biferale, F. Mantovani, M. Pivanti, F. Pozzati, M. Sbragaglia, A. Scagliarini, S. F. Schifano, F. Toschi, R. Tripiccione, Computers and Fluids (2012) in press.

[51] Modha, S. Dharmendra et al., Commun. ACM **54** (2011), 62.

[52] Izhikevich EM, Edelman GM (2008) Proc. Nat. Acad. Sci. USA **105** (2008) 3593.

[53] `http://euretile.roma1.infn.it/`

[54] F. Loeffler, J. Faber, E. Bentivegna, T. Bode, P. Diener, R. Haas, I. Hinder, B.C. Mundim, C.D. Ott, E. Schnetter, G. Allen, M. Campanelli, P. Laguna., " Classical and Quantum Gravity, **29** (2012) 115001.

[55] B. Giacomazzo, L. Rezzolla, L. Baiotti, "Accurate evolutions of inspiralling and magnetized neutron-stars: equal-mass binaries", Phys.Rev. D **83** (2011) 044014

[56] N. Bucciantini, E. Quataert, B.D. Metzger, T.A. Thompson, J. Arons, L. Del Zanna, "Magnetized relativistic jets and long-duration GRBs from magnetar spin-down during core-collapse supernovae", MNRAS **396** (2009) 2038.

[57] A. Burrows and J. M. Lattimer, "The birth of neutron stars", Astrophys. J. **307** (1986) 178.

[58] M. Shibata and K. Taniguchi, "Merger of binary neutron stars to a black hole: Disk mass, short gamma-ray bursts, and quasinormal mode ringing", Phys. Rev. D **73** (2006) 064027.

[59] L. Dessart, A. Burrows, C.D. Ott, E. Livne, S.C. Yoon, N. Langer, "Multidimensional Simulations of the Accretion-induced Collapse of White Dwarfs to Neutron Stars", Astrophys. J. **644** (2006) 1063.