PHYSICS PROJECTS ON APEmille

September 20, 1994

C.R. Allton[1], S. Antonelli[1], L.M. Barone[1], M. Bellacci[1], R. Benzi[2], N. Cabibbo[2], M. Crisafulli[1], G.M. Divitiis[2], V. Gimenez[1], G. Curci[3], A. Donini[1], R. Frezzotti[2], M. Guagnelli[2], D. Lancaster[1],V.Lubicz[4], G. Marchesini[5], E. Marinari[2], G. Martinelli[1], E. Onofri[6], P.S. Paolucci[1], G.Parisi[1], R. Petronzio[2], F. Rapuano[1], J.J. Ruiz-Lorenzo[1], R. Sarno[1], M. Serone[1], R. Simonazzi[1], N. Stella[1], A. Tenenbaum[1], R. Tripiccione[7], A. Vladikas[2], F. Zuliani[1].

[1] Dip. di Fisica "G. Marconi", Università degli Studi di Roma "La Sapienza" and INFN, Sezione di Roma, P.le A. Moro 2, 00185 Roma, Italy,
[2] Dip. di Fisica, Università di Roma " Tor Vergata" and INFN, Sezione di Roma II, Via della Ricerca Scientifica 1, 00133 Rome, Italy,
[3] Dip. di Fisica Università degli Studi di Pisa and INFN, Sezione di Pisa, Piazza Torricelli, 56100 Pisa, Italy,
[4] Dept. of Physics, Boston University, Boston, MA 02215, USA,
[5] Dip. di Fisica, Università degli Studi di Milano and INFN, Sezione di Milano, via Celoria 16, 20133 Milano Italy,
[6] Dip. di Fisica, Università degli Studi di Parma and INFN, Gr. Coll. di Parma, Sez. di Milano, v. delle Scienze, 43100, Parma Italy,
[7] INFN, Sezione di Pisa, Via Livornese 582, 56010, San Piero a Grado (Pisa) Italy.

1 INTRODUCTION

This document represents the effort of several physicists, presently engaged in numerical simulations on computers of the APE100 family. The document intends to present a wide spectrum of projects, that will become possible with the realization of a parallel computer in the one TeraFlops range. Based on previous experience, several common demands have emerged. Besides computer power, good performances can be achieved if the new machine will have:

1. a compact structure, with a small number of nodes per direction and the minimum possible total number of nodes;
2. more than 120 data registers per node;
3. a local memory of 20–30 GigaWords;
4. I/O performances of at least 60 MegaBytes/sec;
5. availability of double precision floating point operations;
6. integer, arithmetic and logical operations on MADmille;
7. local addressing capability, i.e. the processing nodes may access their own memories using different local addresses.

The reasons for the above requirements will be clearer from the discussion in secs. 2–9 and will be summarized in the conclusion.

Here, we only want to add a few words about 7., which will not be discussed any more in the following. Many physical studies require processing of large matrices both by rows and columns. A typical example is a multi-dimensional Fourier transform, which finds many applications, such as Fourier acceleration in lattice gauge theories, proposed by R.K. Wilson and collaborators, metereology, climatology, signal analysis, etc. In these cases, one requires a "data transposition phase", that reshuffles the data among the nodes of the machine from a "cut along the row" distribution toward a "cut along the column" one. A local addressing capability allows a high efficiency implementation for matrix transposition on a 3D interconnection topology

[1]. Moreover, the availability of local addressing is of high importance for complex systems simulations, as shown in sec. 7. Molecular dynamics studies may also benefit from local addressing features.

In what follows, we have denoted by APE100-"Tubo" the 128 node machine, with a peak speed of about 6 GigaFlops, by APE100-"Torre" the 24 GigaFlops machine, and simply by APE100 the final hundred GigaFlops computer, which is now being assembled.

The document has been divided in sections, related to different sectors of activity. It is clear that many of the subjects, treated separately, could have been discussed together. For example, the calculation of weak matrix elements, discussed among the "quenched" calculations, will certainly be repeated in full QCD. It is also evident that the panorama exposed in the following does not exhaust the possible, interesting problems which can be studied with a one TeraFlops computer. The problems discussed here represent the main subjects on which the authors of this document are presently engaged, or plan to be engaged in the foreseeable future.

2 HADRON PHENOMENOLOGY ON THE LATTICE

2.1 Generalities

In the last few years, numerical simulations were extensively used in hadron phenomenology, with relevant results in the domain of strong and weak interactions. With the increasing computer power available, lattice QCD has become an essential and unique approach for obtaining accurate determinations of several quantities such as decay constants, weak transition amplitudes, form factors, structure functions, etc. In some cases, theoretical predictions from the lattice can be confronted to new experimental measurements (e.g. the prediction of the leptonic decay constant of the charmed meson D_s [2]), in other cases, they had an important impact in weak interactions and CP violation physics (e.g. the calculation of the kaon B-parameter B_K, relevant for K^0–$\bar{K}^0$ mixing [3]).

The advent of a one TeraFlops computer, with an adequate local memory and I/O, would open extraordinary perspectives in the calculation of matrix elements relevant to hadronic processes, allowing a substantial reduction of statistical and systematic errors. A large group of physicists, in Rome, has started its activity on the APE-100-"Tubo", about three years ago. During this time, the group has acquired the necessary experience and software on the APE100 architecture. We also recall here that much of the research, currently done in the field, has been initiated in Rome by this group, also in collaboration with the main international groups (e.g. UKQCD) [4]–[12]. In the following, we give a brief description of the possible phenomenological calculations, for which the advent of APEmille would allow extraordinary progress.

2.2 Kaon decays

"Improved" actions a lá Symanzik [13]–[14] and non-perturbative renormalization of the of operators in the effective weak Hamiltonian [15] remove the main obstacles which have jeopardized, in the past, the calculations of the $\Delta I = 1/2$ transition amplitude, which is essential for the understanding of the octet enhancement. Moreover, it will become possible to compute the matrix elements of the operators contributing to ϵ'/ϵ.

From previous studies [3, 6], it is clear that a similar program can be carried out only if it is possible to approach values of quark masses close to the chiral limit, with a very large statistical sample of gauge field configurations. A "quenched" calculation of the relevant amplitudes, with an "improved" action, low values of quark masses (corresponding to $M_\pi \sim 200$ MeV), at a moderate value of the lattice strong coupling constant, $\beta = 6/g_0^2 = 6.0$ and 6.2, would require the following parameters and resources

$\beta = 6.0$–6.2

$V = 48^4$–72^4

$N_{conf} = 500$

Number of light quark masses= 5–6

Local Memory $\sim$ 5.0 GigaWords

I/O= at least 60 Megabytes/sec

CPU time $\sim$ 6 Months

2.3 Electro-magnetic form factors and structure functions

An important sector, which deserves a renewed interest, is constituted by the study of form factors in the time-like region $q^2 < 0$ and of structure functions. After the first semi-quantitative attempts [7], which showed the feasibility of these calculations, even in the presence of large non-perturbative subtractions, no real progress has been done [16, 17]. Also in this case, the use of "improved" actions and non-perturbative renormalization of the relevant operators allow a reduction of systematic errors. This recent theoretical progress, combined with the power of APEmille, would make the results very accurate.

Let us mention some of the most interesting quantities that could be computed with high accuracy: electro-magnetic form factors for π, ρ and nucleons, at large momentum transfer , $q^2 \gg \Lambda_{QCD}$, which are very important for the understanding of QCD dynamics; form factors relevant in hyperon decays, their dependence on quark masses, F and D terms and $SU(3)$ breaking effects; the first four or five moments of the valence quark structure functions for π, K and nucleons; proton decay amplitudes, relevant for GUT and SUSY-GUT model predictions. In all these cases, besides running at values of β, lattice sizes, etc., used for the study of kaon decays (see previous section), it would be useful to perform a study, at a larger value of β, in order to have access to values of $q^2 = 4$–8 GeV^2, without large discretization errors. Assuming that all quantities have a smooth dependence on the quark mass, as preliminary results seem to indicate, it would not be necessary to use extremely large lattices. A possible choice of parameters could be

$\beta = 6.6$–6.7

$V = 72^4$

$N_{conf} = 500$

Number of light quark masses= 4

Local Memory $\sim$ 5 GigaWords

I/O at least 60 Megabytes/sec

CPU time $\sim$ 4 Months

2.4 The η' mass and polarized structure functions

An intense experimental and theoretical activity has been triggered by the discovery of the EMC collaboration, confirmed at SLAC, of a large violation of the Ellis-Jaffe sum-rule. The experimental results indicate a large component of sea quarks (gluons), present in the polarized structure functions, as is also the case for the nucleon σ-term. Several experiments (SMC and HERMES for example) have been planned to improve the accuracy of the EMC measurements and new results have already been presented or are under way.

The ideal approach would be to perform a lattice study of the polarized structure functions in full ("unquenched") QCD. It is not clear, however, that the present state of the art would allow results accurate enough in this domain, even though APEmille would be of great help in this direction. An alternative approach, valid also in the calculation of the η' mass, is to try to compute the polarized structure functions, in the "quenched" approximation, following the Witten-Veneziano method. Previous essays have shown that this requires an extraordinary number of configurations ($\sim$ 10000) [18]. A study on a small volume, at $\beta = 6.0$ and with a statistics of about 10000 configurations would be a significant progress in this field. A possible choice of the parameters is

$\beta = 6.0$

$V = 24^3 \times 64$

$N_{conf} = 10000$

Number of light quark masses= 2–3

Local Memory negligible

I/O not required

CPU time $\sim$ 2 Months

2.5 Heavy flavour physics

A sector where lattice QCD has already given an important contribution is the physics of D and B mesons. Besides f_{D_s}, we recall the large value of f_B, predicted by lattice calculations, and then confirmed by QCD sum rules, and its implications for CP violation in B decays [2]; the calculation of the form factors relevant for $B \to D$, D^* and the Isgur-Wise function [19]; the form factor entering the radiative decay $B \to K^*\gamma$ [20], and the calculation of the transition amplitudes $B \to \pi$, ρ, essential for the experimental determination of $|V_{ub}|$ [9, 21]. Lattice calculations can give important information on inclusive B decays, through the determination of the binding and kinetic energy of the heavy quark inside a beauty hadron. Finally, one can obtain, without any extra effort, results for quarkonia spectra, their wave functions, spin-spin and spin-orbit effects and decay constants.

The main limitation to accurate determinations of the quantities described above comes from the extrapolation of results, obtained at values of the quark mass around the charm, to the b quark mass, and from the limited range in momentum transfer in the study of the q^2-dependence of the form factors. These limitations can be strongly reduced by working at larger values of β and on larger lattice volumes. Possible choices, that would allow a sensible reduction of the extrapolation in the heavy quark mass and in momentum transfer are

$\beta = 6.6$–6.8

$V = (72^3 \times 144)$–$(96^3 \times 192)$

$N_{conf} = 500$

Number of light quark masses= 3–4

Number of heavy quark masses= 5–6

Local Memory $\sim$ 10–30 GigaWords

I/O at least 60 Megabytes/sec

CPU time $\sim$ 10–20 Months

3 FULL QCD SIMULATIONS

Full QCD simulations, using Wilson Fermions with two flavours, are at an early stage, with respect to quenched simulations. First of all, statistical and systematic errors are not fully understood. There is evidence of long range statistical correlations among the gauge field configurations, generated by standard Monte Carlo algorithms [22, 23] and, concerning possible systematic errors, only finite size effects have been studied [24].

Quantitative estimates of full QCD corrections to many physical observables, already studied in the quenched approximation, such as the hadron mass spectrum, weak matrix elements and meson decay constants, are not available yet. This is mainly due to the values of the physical parameters β and $m_{\rm sea}$, the mass of the "sea" quarks, that are outside the interesting region (β should be in the scaling region, and $m_{\rm sea}$ near the chiral limit) and to the smallness of current lattice volumes. All the simulations have been done, so far, with $(M_\pi/M_\rho)^2 \geq 0.35$, $\beta \leq 5.6$ and volumes up to $16^3 \times 32$ in size, corresponding, for current values of β, to $1.5 - 2.0$ Fm. Moreover, with Wilson Fermions, there are no calculations either of the glueball mass spectrum or of the $\pi - \eta'$ mass splitting. Finally, most of the results, reported in the literature, have been obtained using sets of no more than 1000-2500 trajectories (after ~ 500 thermalization trajectories), with a typical autocorrelation length, of $\sim 80 - 120$ trajectories [25]. Apart from the corrections to physical observables, already computed in the quenched approximation, fermionic loops cause new phenomena, which have been until recently unexplored, such as the breaking of the string tension [26]–[28] and strong decays of hadronic excited states. An important field of application of full QCD is also the study of the deconfinement phase transition (QCD thermodynamics, see below).

Computing power is a severe limit to obtaining realistic results in full QCD. For Wilson Fermions, the CPU time needed to perform a calculation grows as fast as $(V)^{5/4}/(m_{\rm sea})^{7/2}$. This means that, in order to make calculations in the interesting region (high β values, large lattice volumes and low quark masses) CPU time is prohibitive. A simulation, performed on a box of length 2 Fm, with β at the beginning of the scaling region, $\beta = 5.6$ [22], should be done on a lattice of at least $24^3 \times 32$. In order to approach the point where $m_{\rm sea} \leq m_{\rm strange}$, using the scaling law given before, would require

20 years of CPU time, on a one GigaFlops machine, to obtain statistics of 1000 full QCD trajectories. A further problem is the generation of statistically independent gauge field configurations. With present algorithms, on a volume $16^3 \times 32$ and with $m_{\rm sea} \sim m_{\rm strange}$, the autocorrelation time for the plaquette energy is about 100 trajectories. It is very difficult to extrapolate the previous autocorrelation time to larger lattices and consequently to make estimates of the statistical errors for the different observables. One can assume that it is possible to have ten independent field configurations, out of 1000 trajectories.

From the above considerations, it is clear that only with a computer in the 100-1000 GigaFlops range (and $70 - 100\%$ efficiency) can we hope to reach the interesting region of coupling constants and quark masses in full QCD, for example in order to compute the $\rho \to \pi\pi$ decay width.

3.1 Speed requirements

The APE group is presently working on a lattice of size $16^3 \times 32$, with $m_{\rm sea} \sim m_{\rm strange}$, using 6.4 GigaFlops and 25 GigaFlops Machines. Other groups make similar studies, using computers in the 1–15 GigaFlops range. With a $16^3 \times 32$ lattice, one can have a physical size ranging from 3.0 Fm, at $\beta = 5.0$, down to 1.2 Fm, at $\beta = 5.6$. One year of a 100 GigaFlops machine would be necessary for a run with $m_{\rm sea}$ lower than the strange quark mass, on a $16^3 \times 32$ lattice, or for a run with $m_{\rm sea} \sim m_{\rm strange}$, on a $24^3 \times 32$ lattice. In order to work in the scaling region, $\beta \geq 5.6$, it would be necessary, however, to perform simulations with lattice sizes of at least $24^3 \times 64$, which require computers in the 100-400 GigaFlops range. With the one TeraFlops machine one can also perform simulations on a $32^3 \times 64$ lattice volume and $m_{\rm sea}$ in the strange quark mass range.

3.2 Memory requirements

The memory occupation for the $16^3 \times 32$ lattice is 64 MegaWords, which represent 50% of the APE100-"Tubo" memory and the 12.5% of the APE100-"Torre" memory. On APE100, we can arrange a $24^3 \times 64$ lattice with an occupation of $20 - 25\%$ of the memory capacity.

On APEmille, as the CPU speed puts an upper limit of $32^3 \times 64$ to the lattice size, a memory capacity of 1.5–2.0 GigaWords would be enough. More efficient inversion algorithms may need the use of many auxiliary fields. For this reason, a large margin of memory capacity is welcome.

3.3 I/O requirements

It is important to backup the gauge configurations, at least once a day, to be ready to restart the machine after any possible stop. We have to backup:

$V = 16^3 \times 32 = 40$ MegaBytes,

$V = 24^3 \times 64 = 256$ MegaBytes,

$V = 32^3 \times 64 = 640$ MegaBytes.

Using the APE100-“Tubo”-“Torre”, a gauge field configuration, on a $16^3 \times 32$ lattice, can be written in 1-2 minutes; the full APE100, on a $24^3 \times 64$ lattice, in 15 minutes. In this way we use about 1% of the CPU time for I/O. It is important to maintain the same I/O to CPU time ratio on the one TeraFlops machine. In this case, the relevant figure is a I/O speed of 1.0-1.5 MegaBytes/sec.

3.4 Architecture requirements

Since the typical volumes will be in the range $V = (16^3 \times 32)$–$(32^3 \times 64)$, it is more convenient to work on a machine with a small number of nodes in each direction. A possible choice is to have a three dimensional array with 8 nodes in two directions and 32 in the third one ($8 \times 8 \times 32$), or a 4-dimensional array ($8 \times 8 \times 2 \times 32$).

The typical efficiency in the inversion of the pseudo-fermionic matrix, which is the most CPU consuming part of our programmes, is about 70%. In the future, it will be important to obtain similar or greater performances. A machine with many nodes can reduce the efficency, because of the great ratio between remote transfers versus internal ones (with a $16^3 \times 32$ lattice, on a APE100-“Torre”, we have only 256 points per node, while with a $24^3 \times 64$ lattice, on APE100, we will have 432 points per node).

4 QCD AT FINITE TEMPERATURE

One of the most interesting consequences of QCD is the existence of a new phase, the so called quark-gluon plasma phase. We believe that the quark-gluon plasma can exist at high temperature and/or high density. The conditions for a phase transition, from ordinary hadronic matter to the quark-gluon plasma, could be realized in heavy ion collisions, thus opening a new frontier of investigation. Lattice simulations, performed in the absence of matter fields, have established the existence of a "weak" first order phase transition and have allowed a measurement of the critical temperature , with great accuracy, using finite size effects. Large violations to the asymptotic scaling, predicted in perturbation theory, have been observed [32]–[34]. It is desireable to decide whether these violations persist on larger lattices and improve the measurement of the critical temperature (T_c): the energy density, necessary for the formation of the plasma increases as T_c^4, and it is at present the main limitation in projects of heavy ion experiments.

The inclusion of matter fields has allowed a partial exploration of the phase diagram in the T, m plane (in the $SU(N_f)$ flavour symmetry approximation) and there are increasing indications that the value of T_c is smaller than in the pure gauge case. The high CPU speed requirements of "full" QCD has severely limited the sizes of the lattices used so far, and we do not have satisfactory measurements of the scaling behaviour of the physical quantities or measurements at realistic values of quark masses [30]. In particular, it has not been decided yet whether the transition line in the T_c, m plane is continous, or whether there is a break in passing from heavy quark masses ($m \to \infty$ corresponds to the pure gauge theory) to masses close to the chiral limit. It would be of preminent phenomenological interest to establish if the break point is due to lattice artefacts or persists in the continuum limit [31].

Finally, it would be very interesting to explore the phase diagram at finite baryon density. These simulations are even more demanding, because the chemical potential makes the fermionic determinant complex.

In order to evaluate the impact of APEmille in this field, we observe that the measurement of the peak of the specific heat, on a lattice of temporal extension $N_t = 6$ and with a quark mass $m\,a = 0.025$, would require a CPU power of 200 MegaFlops, for a few months run; by assuming that the computational effort grows as a^{-8}, the simulation would require 500 GigaFlops, for several months, if one works at $N_t = 16$. The three dimensional structure of APEmille fits perfectly with finite temperature lattice simulations. The availability of double precision is desireable[1], whereas the I/O requirements are probably smaller than for fluid-dynamics, see below.

5 MEASUREMENT OF α_s

A further field of application consists in the non-perturbative calculation of the effective strong coupling constant α_s. The method is based on the study of scaling laws in the lattice size. With a CPU power of about two GigaFlops, one can obtain a few per cent accuracy in the calculation of α_s, for $SU(2)$, without dynamical fermions, in a period of four months. This is the precision required to be competitive with the LEP measurements of the strong coupling constant, at high energies. The extension to full QCD would require the treatment of fermion loops which, with the present algorithms, would imply an increase of a factor 100, with respect to the $SU(2)$ case, i.e. a CPU power of about one Teraflops. New algorithms, actually under test, need a high number of auxiliary boson fields, thus demanding a local memory/processor ratio larger than the APE100 one. An important issue, for some of the calculations needed with these new methods, is the availability of double precision in the floating point operations, acceptable even at the price of some degradation of CPU performances.

The calculation of α_s does not require, in general, an asymmetric extension of the lattice in one direction and it could profit of a more symmetric topology of the machine. Usually, for $SU(2)$, it has been possible to obtain efficiencies of about 60%; this could be improved, expecially in those cases which have a poor F/L ratio, by increasing the number of register files, beyond the number available for APE100 (120).

[1]Notice that double precision is an essential requirement for metereology and other important fluid-dynamical studies.

6 PERTURBATIVE CALCULATIONS ON THE LATTICE

A project of weak coupling perturbative calculations, at high orders, has been started on the APE100-"Tubo" in Parma [37]. The first studies have shown that the stochastic method, based on the ideas developed by Parisi and Wu, offers a computational technique able to reach high orders in perturbation theory. At present, results have been obtained, up to four loops, for observables in the pure gauge theory. Using this technique, one may hope to study the problem of the resummation of the perturbative series. The method is also applicable to the calculation of β-functions at high orders, to chiral models and to low temperature series in statistical mechanics. The actual configuration of the APE100-"Tubo" is adequate for the first few orders. However, the complexity of the calculations increases factorially with the number of loops, and the method would become enormously powerful with a computer like APEmille.

7 COMPLEX SYSTEMS

7.1 Generalities

In the following, we will describe a series of problems, in complex and disordered systems [38, 39], which can be implemented on APE computers. The set of subjects given below is certainly not exhaustive, since to cover all aspects of complexity and disorder, which include so many different biological, economic and social systems, is impossible. For this reason, we will focus on subjects that are already under study and whose numerical simulations run on the APE100 architecture, with good efficiency, ranging between 20% to 70%.

We find interesting, and worthy to stress, the interplay between problems of exclusive theoretical interest (lower critical dimension, ultrametricity, breaking of dimensional reduction) and practical applications (protein folding, optimization methods, etc.). APEmille will be the ideal computer to achieve the goals of the research program described below.

7.2 Spin-Glasses

The original seed of a large part of disordered systems is given by materials commonly denoted as *Spin-Glasses* [38, 40]. The Parisi solution shows characteristics which are completely new for physical systems. It describes a situation where the existence of many (stable and metastable) states is the crucial ingredient. The main question is whether the main features of the mean-field solution do survive in a finite number of dimensions. Whereas in 4 dimensions ($4d$) one obtains a picture very similar to the mean field solution, it is very likely that $d = 3$ is the lower critical dimension. This makes the interpretation of numerical results difficult. Only simulations on very large spin systems can give an answer to these questions.

At present, numerical simulations with real coupling constants J, distributed according to a Gaussian with zero mean, are running on APE100. The efficiency is of order 20% and the main limitations come from the poor I/O performance. On the contrary, APEmille, with high I/O performance, will allow the (pseudo-)broken phase to be studied in detail. Local memory requirements would not be a problem. Two different kinds of numerical simulations are usually performed. In the first case, we have a different physical system on each processor (typically with a different realization of the noise, in order to average over a quenched noise). We use this approach to study the broken phase. In the second case, we simulate a large lattice, lowering the temperature T from the warm phase.

The addition of an integer processor, to the MADmille, will allow the simulation of models with integer couplings with spectacular efficiency (thanks also to a look-up table, allowed by local addressing). It will also be interesting to study the behavior of the $2d$ model (both with real and integer couplings) since there are interesting open questions about the behavior of physical quantities as $T \to 0$.

7.3 Dynamics and Aging

Many interesting features of the dynamical behavior of complex systems (and also of standard spin systems like the ferromagnetic Ising model) are emerging after recent studies [41]. APEmille will be a perfect device to study dynamics and aging, in the mean field spin-glass Sherrington–Kirkpatrick model. A computer programme has already been implemented for running on APE100, with high efficiency [42].

Aging phenomena demand the simulation of many time decades of evolution for a system, subject to an instantaneous perturbation. APE100, and in the future APEmille, are crucial tools to get accurate information about the long time region.

7.4 Heteropolymers

Simulations of heteropolymers (in a simple approximation where different components of the polymeric chain have different quenched random interaction) run with high efficiency on APE [43]. One stores one disordered configuration (which specifies the heteropolymer) per processor, and studies the dynamics of a few copies of chains. The efficiency is of the order of 60% already for quite short chains (of the order of 30 sites). Heteropolymers behave in a way very similar to the so called REM model, i.e. with a pattern reminiscent of a classical statistical system undergoing a first order phase transition.

These simulations, important to obtain information about the mysterious phenomenon of protein folding, are extremely CPU demanding. Thus a one Teraflops machine would have a decisive impact in the field.

7.5 Random Field Ising Model

The Random Field Ising Model is, on the one hand, a fascinating theoretical problem, and, on the other hand, it is an interesting way to study interface dynamics. A code for solving the mean field equations [44], in a finite number of dimensions already exists on APE runs with an efficiency close to 50%. We are trying now to determine critical exponents, at the transition point. Only with machine like APEmille, will it be possible to keep under control *all* the solutions in the broken phase (where the number of solutions grows exponentially).

7.6 Optimization

The *tempering* Monte Carlo method (in which the system rearranges its temperature, while remaining at Boltzmann equilibrium) has proved to be of terrific effectiveness for systems like spin-glasses and random field Ising models[2]. Tempering is more than an optimization method at zero temperature, but can also act as an optimization method as well. It can be superior to classical annealing (since it contains a built–in scheduling, which is the big problem of annealing-like methods). Tempering is easily parallelized.

7.7 Interface Dynamics and Surface Growth

Dynamics of interface in presence of disorder (which can be different in nature and behavior) is the subject of a large number of studies. Simulations devoted to a better understanding of these phenomena are running on APE100. The efficiency is close to 20% and would improve dramatically on the APEmille architecture. We have in mind here to continue studies of the KPZ equations, of charge density wave systems, and of the pinning by impurities. A relevant application is the behavior of flux lines in high T_c superconductors.

7.8 Neural Networks

Much effort has been devoted to implement neural networks on APE. The availability of integer arithmetics on the MADmille will allow large flexibility in planning and training neural networks (also in connection to applications of neural networks useful to high energy physics experiments).

7.9 Random Surfaces

A computer with the facilities of APE, plus local addressing, will make possible studies of random surface dynamics. It will also be possible to study adaptive grids. For this problem too, there is a (twofold) theoretical interest: quantum gravity and string theory, on the one hand, and liquid interfaces on

[2]On the contrary it seems to be quite poor for systems like heteropolymers and in general REM-like systems.

the other. There is also a more practical interest, adaptive grids, which are of crucial importance in many engineering problems.

7.10 Architectural Requirements from Complex Systems

Integer, arithmetic and logical operations should be added on MADmille, together with local addressing. For many physical systems, the use of integer arithmetics can largely simplify the problem, and make its numerical solution far easier to obtain.

As compared to the quantity of silicon, needed for a floating point device, the room needed for integer and logical basic operations is very small. A large gain in the potentialities of the machine is paid only at a small price in device complexity.

We propose to implement the basic integer and logical operations: $*$, $+$, *AND*, *OR*, *NOT* and *XOR*. The shift operation will also be crucial. We propose to implement periodic (i.e. cyclic) shifts and the normal shift. It would be wise to have hardware built-in shift operations of 1, 2, 4, 8 and 16 bits. This feature will complement local addressing.

The MADmille based machine will be a computer of excellence in computational physics, suited for a wide spectrum of applications.

8 FLUID-DYNAMICS

Viscous fluids, described by the Navier-Stokes equations, are a typical example of a non-linear classical system which exhibits, under appropriate conditions, a chaotic behaviour. The study of the statistical properties of turbulent systems is therefore an extremely interesting area of classical field theory. The techniques that can be applied in this area bear striking resemblances with those developed in LGT. Several different techniques have been developed to solve numerically the Navier-Stokes equations. One such technique, the "Lattice Boltzmann Equation" (LBE), is based on statistical averages of quantities defined on a cellular automata. The "macroscopic" averages can be shown, under appropriate conditions, to follow the Navier-Stokes equation for incompressible fluids. The LBE approach is very well suited for a parallel processor and has been succesfully implemented on APE100 [46]. Efficiencies of the order of 50 % of the peak speed of the machine have been reached.

An understanding of the physical mechanisms underlying the dynamics, obtained through LBE simulations, and comparison with experimental data have recently led to argue that an unexpected property of the correlation functions of the velocity fields, called "Extended Self Similarity" (ESS), does hold. This property consists in the possibility of absorbing deviations from asymptotic scaling of the correlation functions of the velocity field, by re-writing them in terms of just one correlation function, taken for reference [47, 48].

This expectation has been verified in a systematic study of the statistical properties of convective turbulence (described by the Rayleigh-Benard cell) that has been recently completed on APE100-Tubo. The degree of turbulence of convective systems is measured by the so called Rayleigh number, Ra. A APE100-Tubo unit has been used to simulate a (memory limited) mesh of 160^3 points. Typical runs require about one month of APE100-Tubo time. In these conditions, Ra is limited to a value of 10^7 [49]–[51]. Taking into account "critical slowing down", we may expect that a full APE100 machine could accomodate a mesh of 512^3 points, allowing to reach Ra of a few 10^8, in three or four months of running time.

From the physical point of view, the region $Ra = 10^9$–10^{12} is extremely important, since it covers the transition from soft to hard turbulence, which has recently been observed experimentally. This region cannot be studied with present versions of APE100 due to limitations of both peak performance and memory, but might become accessible with a machine in the TeraFlops range of performance. For instance, a system with $Ra \simeq 10^{10}$ can be simulated under this conditions:

lattice size 1024^3

Local Memory 20 GigaWords

I/O at least 25 Megabytes/sec

CPU time $\simeq 10 - 15$ months

9 MOLECULAR DYNAMICS

Molecular Dynamics (MD) is widely used in condensed matter physics to simulate the dynamics of many different systems, ranging from atomic liquids to molecular solids, from interfaces to macromolecules. The integration of the equations of motion gives detailed knowledge of all mechanical variables and, through time averages, yields estimates of the thermodynamical properties of the system. In the last two years the portability and efficiency of simple MD programs on APE100 have been checked, with various machine configurations ranging up to 128 nodes. The results have shown that one can simulate on APE100 very large systems (up to 11 million particles with 128 nodes). On the other hand, the present structure of the machine diminishes the capability to simulate a generic system. The main limitations come from the available memory per node, the maximal size of integer vectors and the lack of local addressing capability. These limitations become less severe if one studies lattices without atomic diffusion. Finally, for some specific problems, it may be necessary to have double precision. The list of desiderata for APE1000 depends on the kind of problem one wishes to handle on this computer.

1. In APE100, each node is endowed with a memory of 1 MegaWords (floating point numbers). In a realistic simulation of a liquid or a solid, with a standard range of interaction, each node could host about 70,000 particles if it were endowed with an additional memory of about 3 MegaWords (integers) for the "list of neighbours". With this additional capacity APE could be used to simulate any kind of large system, solid or liquid.

2. The standard technique based on the "list of neighbours" requires local addressing capability of each node on its own memory.

3. If the simulation of the system aims at a precise knowledge of some or many dynamical variables, that is, if mechanics rather than statistics is important, double precision seems unavoidable. Double precision would also allow to exploit the volume available to simulate large systems for statistical (thermodynamical) purposes; otherwise, the accumulation of round-off errors due to the simple precision can hamper this exploitation.

4. Transferring data from and to APE is relatively slow; this bottle-neck should be improved if one needs frequent transfers of large amount of data (e.g., if one wants to follow the dynamics of the simulated system on-line, on a graphic station) or frequent access to auxiliary files on an external disk (e.g., access to force parameters for a highly inhomogeneous macromolecule).

Many of the requiremets of MD are common with other fields discussed in this proposal, and will be implemented on the APE1000 project.

10 CONCLUSION

From the above discussion, a fascinating physics program, in fundamental sectors of particle physics and fluid-dynamics, emerges. An extraordinary improvement of our present resources, can be obtained, by using a computer with the charateristics of APEmille. As discussed in the sections devoted to "quenched" calculations, a fundamental point, for the success of the research program, is the possibility of working on very large lattices, corresponding to volumes in the range $V = 48^4$–$(96^3 \times 192)$. The CPU speed is, from this point of view, less critical. Our previous experiences have shown that the main limitations to the accuracy of the results were due to local memory requirements, which have forced us, in the past, to use the "thinning" approximation. It is also clear that the problem can find an efficient solution only if it is possible to combine a large local memory (20–30 GigaWords) with a rapid I/O, in order to be able to read/write, with an acceptable overhead in time, quark propagators on external disks. With a 60 MegaBytes/sec, the time for I/O would be equal to the CPU time. Any improvement on the I/O performances is consequently welcome. A large local memory and good I/O performances are also demanded by calculations of the strong coupling constant and by fluido-dynamics. An increase of register files with respect to APE100 is thus also needed. Finally, full QCD calculations do not pose special problems with memory or I/O (1000-1500 Kbytes/sec), but require the highest CPU performances. A low number of nodes in each machine direction and an overall low number of machine processors is appreciated. Finally, the floating point unit has to be built in such a way to mantain the actual efficiency in the inversion of the pseudo-fermionic matrix. An important issue, for some of the calculations considered here, is the availability of the double precision in floating point operations, even though this can be obtained only at the price of some degradation of CPU performances. Integer, arithmetic and logical operations should be added on MADmille, together with local addressing. Finally to implement the basic integer and logical operations $*$, $+$, AND, OR, NOT and XOR, and the shift operation, is also very important.

References

[1] N. Cabibbo, P.S. Paolucci, Nota Interna, Dip. di Fisica Roma, 93/1018.

[2] As. Abada et al., Nucl. Phys. B376 (1992) 172; K. Bitar et al., Phys. Rev. D48 (1993) 370; C. Bernard, J. Labrenz and A. Soni, Phys. Rev. D49 (1994) 2536; R.M. Baxter et al., Phys. Rev. D49 (1994) 1594; C.R. Allton et al. Nucl. Phys. B(Proc. Suppl.) 34. (1994) 456; C. Alexandrou et al., DESY-93-179/PSI-PR-93-2.

[3] M.B. Gavela, L. Maiani, G. Martinelli, O. Pene and S. Petrarca, Phys. Lett. B211 (1988) 139; F. Rapuano, Nucl. Phys. B(Proc. Suppl.)16 (1990) 381; C. Bernard et al., Nucl. Phys. B (Proc. Suppl.) 17 (1990) 495; N. Ishizuka et al., Phys. Rev. Lett. 71 (1993) 24; R. Gupta et al. Phys. Rev. D47 (1993) 5113; S. Sharpe, Nucl. Phys. B (Proc. Suppl.) 34 (1994) 403.

[4] N. Cabibbo, G. Martinelli and R. Petronzio, Nucl. Phys. B244 (1984) 381.

[5] M. Bochicchio, L. Maiani, G. Martinelli, G.C. Rossi and M. Testa, Nucl. Phys. B262 (1985) 331.

[6] M.B. Gavela, L. Maiani, G. Martinelli, O. Pene, S. Petrarca and F. Rapuano, Nucl. Phys. B306 (1988) 677.

[7] G. Martinelli and C.T. Sachrajda, Phys. Lett. B190 (1987) 15; Phys. Lett. B196 (1987) 184; Nucl. Phys. B306 (1988) 865; Nucl. Phys. B316 (1989) 305; Phys. Lett. B217 (1989) 319.

[8] M.B. Gavela, L. Maiani, O. Pene and S. Petrarca, Phys. Lett. B206 (1988) 113.

[9] M. Crisafulli, V.J. Hill, G. Martinelli and C.T. Sachrajda, Phys. Lett. B223 (1989) 90; V. Lubicz, G. Martinelli and C.T. Sachrajda, Nucl. Phys. B356 (1991) 301; Phys. Lett. B274 (1992) 415; As. Abada et al., Nucl. Phys. B416 (1994) 675.

[10] C.R. Allton et al., Nucl. Phys. B349 (1991) 301.

[11] C.R. Allton et al., Nucl. Phys. B413 (1994) 461; Rome prep. 94/981 to be published in Phys. Lett. B.

[12] U. Aglietti, G. Martinelli and C.T. Sachrajda, Phys. Lett. B324 (1994)85.

[13] B. Sheikholeslami and R. Wohlert, Nucl. Phys. B259 (1985) 572.

[14] G. Heatlie, G. Pittori, G. Martinelli, C.T. Sachrajda and A. Vladikas, Nucl. Phys. B352 (1991) 60; B358 (1991) 212.

[15] G. Martinelli, C. Pittori, C.T. Sachrajda, M. Testa and A. Vladikas, Rome prep.1022, in preparation.

[16] T. Draper et al., Nucl. Phys. B318 (1989) 319; T. Draper, R.M. Woloshyn and K.F. Liu, Phys. Lett. B234 (1990) 121; T. Draper, D.B. Leinweber and R.M. Woloshyn, Phys. Rev. D43 (1991) 1659; D46 (1992)3067.

[17] B. Anderson-Pugh, W. Wilcox, U. Baylor, BU-HEP-93-01.

[18] G. Martinelli, M. Testa and C. Ungarelli, unpublished.

[19] C. Bernard, Y. Shen and A. Soni, Phys. Lett. B317 (1993) 164; S.P. Booth et al., Phys. Rev. Lett. 72 (1994) 462.

[20] C. Bernard, P. Hsieh and A. Soni, Nucl. Phys. (Proc. Suppl.) B26 (1992) 347 (note that there is a factor 2 missing in eq. (4) of this paper); WASH-U-HEP-93-35; K.C. Bowler et al. (CN-UKQCD Collaboration), SHEP 93-94-01 (November 1993).

[21] C. Bernard, A. El-Khadra and A. Soni, Phys. Rev. D43 (1991) 2140 and D45 (1992) 869; C.R. Allton et al. Nucl. Phys. B(Proc. Suppl.) 34 (1994) 477; the UKQCD collaboration, presented by D. Richards at the Glasgow Conference, to appear in the Proceedings; T. Bhattacharya, D. Daniel and R. Gupta, LA-UR-93-3580.

[22] R. Gupta et al., Phys. Rev. D44 (1991) 3272.

[23] K.M. Bitar et al., Nucl. Phys. B(Proc. Suppl.) 30 (1993) 401.; COLO-HEP-321.

[24] S. Antonelli et al., Rome prep.94/1013, submitted to Phys. Lett. B.; M. Fukugita et al. Phys. Lett. B294 (1992) 380.

[25] D. Weingarten, Nucl. Phys. B(Proc. Suppl.) 34 (1994) 29.

[26] R. Altmeyer et al,. Nucl.Phys. B389 (1993) 445; K. Bitar et al., Phys. Rev. D42 (1990) 3794; preprint FSU-SCRI-92-152; F.R. Brown at al., Phys. Rev. Lett. 67 (1991) 1062; M. Fukugita et al., preprint KEK TH-340 (1992); C. Bernard et al., Nucl. Phys. B(Proc.Suppl.) 26 (1992) 262.

[27] U. M. Heller et al., FSU-SCRI-94-09.

[28] C. Bernard et al. Phys. Rev. D46 (1992) 4741; T. Blum et al., prepr. AZPH-TH/94-06.

[29] S. Gupta et al. Phys. Lett. B242 (1990) 437.

[30] F. Karsch, Nucl. Phys. B(Proc. Suppl) 34(1994) 63.

[31] A.A.V.V, Int. Jour. Mod. Phys. C, 2 (1991) 829.

[32] G. Cella et al., Phys. Rev. D49 (1994) 511.

[33] G. Cella et al., Nucl. Phys. B (Proc. Suppl.) 34 (1994) 283.

[34] G.Cella et al., prep. IFUP-TH 12/94, to be published in Phys. Lett. B.

[35] G.M. de Divitiis et al., to appear in Nucl. Phys. B; Rome II prep. (1994).

[36] M.Luescher et al., Nucl. Phys. B389 (1993) 247; CERN-TH.6996/93 (1993).

[37] F. Di Renzo et al., Lattice93, Nucl.Phys. B (Proc.Suppl.) 34 (1994) 795 and paper submitted to Nucl. Phys. (hep-lat/940519).

[38] M. Mézard, G. Parisi and M. A. Virasoro, *Spin-Glass Theory and Beyond* (World Scientific, Singapore 1987).

[39] G. Parisi, Nucl. Phys. **B** (Proc. Suppl.) **26** (1992) 181.

[40] G. Parisi, Physica **A194** (1993) 28.

[41] See for example E. Marinari and G. Parisi, J. Phys. **A**: Math Gen. **26** (1993) L1149.

[42] P.S. Paolucci, Nota Interna 93/947.

[43] G. Iori, E. Marinari and G. Parisi, Int. J. Mod. Phys. **C** (1993) 1333.

[44] M. Guagnelli, E. Marinari and G. Parisi, J. Phys. **A**: Math Gen. **26** (1993) 5675.

[45] A. Bartoloni et al., Int. J. of Mod. Phys. C4 (1993) 993.

[46] A. Bartoloni et al., Int. J. of Mod. Phys. C4 (1993) 993.

[47] R. Benzi et al., Phys. Rev. E48 (1993) 29.

[48] R. Benzi et al., Eur. Lett. 24 (1993) 275.

[49] R. Benzi et al., Eur. Lett. (1994) in press.

[50] R. Benzi et al., submitted to Eur. Lett.

[51] R. Tripiccione, Mem. Soc. Astr.It. (1994) in press.