



# *Dedicated Computers for (High-Energy) Physics*

*R. Tripiccione*

*Physics Department, Universita' di Ferrara*

*tripiccione@fe.infn.it*

*Arcetri, September 20th, 2005*

## *Outline of this talk*

- *Why did it happen that quite a few theoretical physicists turned into computer builders?*
- *Why was it possible?*
- *Three examples:*
  - *Stellar dynamics (GRAPE)*
  - *Lattice QCD (mostly APE)*
  - *Spin-glass systems (SUE/IANUS)*
- *Conclusions: What did we learn? What are we going to learn next?*

*The main point of this talk*

*In many cases, computing accurate predictions of the behaviour of a physical system is hopeless, unless numerical techniques are used*

*The main point of this talk*

*In many cases, computing accurate predictions of the behaviour of a physical system is hopeless, unless numerical techniques are used*

*However Nature has been friendly to us*

*The main point of this talk*

*In many cases, computing accurate predictions of the behaviour of a physical system is hopeless, unless numerical techniques are used*

*However Nature has been friendly to us*

*so the (simple) physics laws behind the behaviour of computers make it relatively easy to build machine that simulate complex physics!!!!*

• *Some problems are really computationally challenging:*

• *Example 1: Stellar dynamics*

*Just try to integrate the equation of motion of  $\mathcal{N}$  stars in gravitational interaction among them*


$$\forall i, \mathbf{V}_i = \sum_j^N \alpha_j / r_{ij}$$

*While a time step is*  $O(N)$

*Computing the forces is*  $O(N^2)$

*What if you have a simple object (a globular cluster  $\sim 10^6$  bodies)?*

Example 2: Lattice Quantum Chromo-Dynamics



*Unfortunately it is not yet known whether the quarks in Quantum Chromodynamics actually form the required bound states. To establish whether these bound states exist one must solve a strong coupling problem and present methods for solving field theories don't work for strong coupling.*

*K. Wilson, Cargese Lectures, 1976*



## Example 2: Lattice Quantum Chromo-Dynamics

The “obvious” solution is to make the path-integral approach numerically manageable by going to a discrete and finite lattice and using Monte Carlo techniques


$$P(\mathbf{U}) \sim e^{-S_g(\mathbf{U})} \cdot \int D[\psi] e^{-\bar{\psi} M[\mathbf{U}] \psi}$$

which boils down to finding the determinant of the sparse but hugely large Dirac op.  $M$

$$M_{xy}(\mathbf{U}) \psi_y = (\delta_{xy} - \kappa H_{xy}(\mathbf{U})) \psi_y$$

The lattice spacing has to be small enough that no UV problems are present, HOWEVER

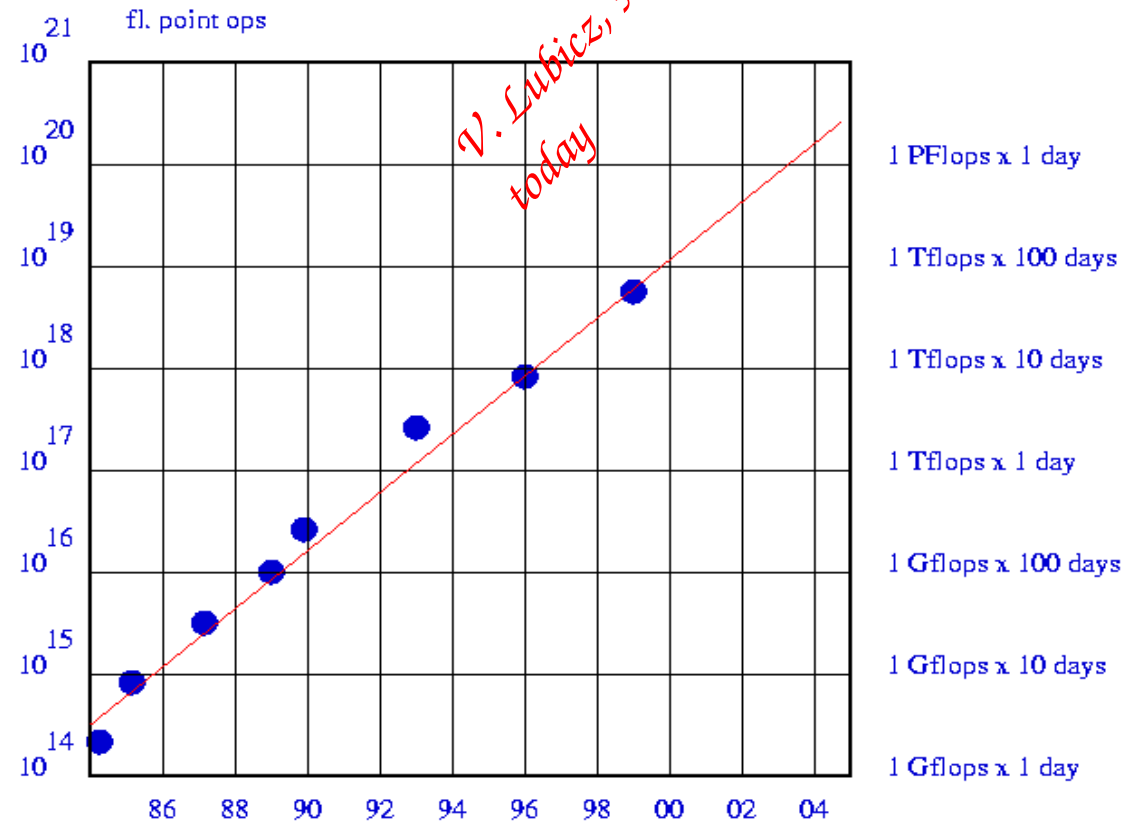
The Lattice size has to be much larger than the Compton length of your lightest particle



*There have been several attempt at parametrizing phenomenologically the computational complexity of the problem:*

$$C \propto \left(\frac{m_\rho}{m_\pi}\right)^6 \cdot (L/L_0)^5 \cdot (a_0/a)^7$$

*The experimentally measured values are challenging indeed:*



### Example 3: Spin Glasses


*A very simple Hamiltonian (defined on e.g. a discrete 3-D lattice)*

$$H = - \sum_{\text{NB}(ij)} \sigma_i J_{ij} \sigma_j, \quad \sigma = \{+1, -1\} \quad J = \{+1, -1\}$$

*G. Parisi's talk, yesterday*

*may hide a tremendously complex dynamics, due to the extremely irregular energy landscape in the configuration space*

*Here one may want to study e.g., the phase structure of the model.*



## *Better computers than those you can buy?*

*Fine: you need a lot of computing power, BUT ....*

*... Why on earth do you think you can do better than an established computer industry?*

*Two answers to this question:*

*1) What we need is not exactly what traditional computers have been good at*

*AND*

*2) What we need is very simple to achieve in terms of computer architecture ...*

*... if we proceed in the direction that basic physics laws point to us*

## *Better computers than those you can buy?*

*1) What we need is not exactly what traditional computers have been good at*

*Either we need long straight sequences of complex mathematical (f.p.) operations*

*e.g.:  $1/\sqrt{r^2}$  (stellar dynamics)*

*or:  $a \times b + c$  among complex numbers (LQCD)*

*or conversely, we need long straight sequences of extremely simple boolean operations*

$$\sum_{\text{NB}(ij)} \sigma_i J_{ij} \sigma_j$$

## *Better computers than those you can buy?*

*2) What we need is very simple to achieve in terms of computer architecture*

*Basic physics help us in many ways:*

*1) Parallel computing is trivially possible in all cases ...*

*... and parallel computing is the physics sponsored way to compute:*

*The basic object is the transistor*

*Industry learns to build smaller and smaller transistors. As  $\lambda \rightarrow 0$*

*obviously  $N \propto 1/\lambda^2$  but speed scales more weakly  $\tau \propto \lambda$*

*that is: it will increasingly be convenient to do more and more things in parallel rather than to do a fixed number of things faster and faster*

# Better computers than those you can buy?


2) What we need is very simple to achieve in terms of computer architecture

Basic physics helps us in two ways:

2) We are interested in modeling local theories:

*This goes (must go) over to the computer structure ->*

*Keep data close in space to where it is processed*




*Failure to do so will asymptotically bring a data bottleneck:*

$$P(L) \propto L^2$$

$$B(L) \propto L$$

*does not work in stellar dynamics, but there*

$$B \simeq \sqrt{P}$$



An historical remark:

- *Doing things one after the other (serially)*
- *Keeping data storage and data processing separated (in principle and practice)*

*are the cornerstones of the famous von Neumann model of computing*

*Q: So was Von Neumann wrong?*

*A: No, simply he was interested in the  $P \rightarrow 0$  regime  
while today we are approaching the  $P \rightarrow \infty$  regime .....*



## *Better computers than those you can buy?*

*These advantages have been exploited in big projects and small projects:*

*Stellar dynamics and LQCD are examples of big projects*


*Spin Glass simulation engines are much more back yard attempts.*

## *Important dates in LQCD computing*

- *1979: The early pioneers: the Caltech Ising machine (D. Toussant, G. Fox, C. Seitz)*
- *circa 1985:*
  - APE (16 nodes, 1 Gflops)*
  - Columbia (~ 1 Gflops)*
  - GF11 (IBM/Yorktown)*
- *1990 - 1995:*
  - APE100 (500 – 1000 nodes, 50 – 100 Gflops)*
  - Columbia2 (also about 100 Gflops)*
- *1995 – 2000:*
  - APEmille (1.8 Tflops installed)*
  - QCDSF (1 + 1 Tflops at Columbia & Brookhaven)*
  - CP-PACS (Tsukuba + Hitachi, 600 Gflops)*
- *2000 – 2005:*
  - apeNEXT*
  - QCDOC (Columbia + Brookhaven + IBM/Yorktown)*




## *APE gallery*




*“It is easy! It may as well turn out to be possible.....”*

*G. Parisi, Spring 1985*






## *ApeNEXT picture Gallery*



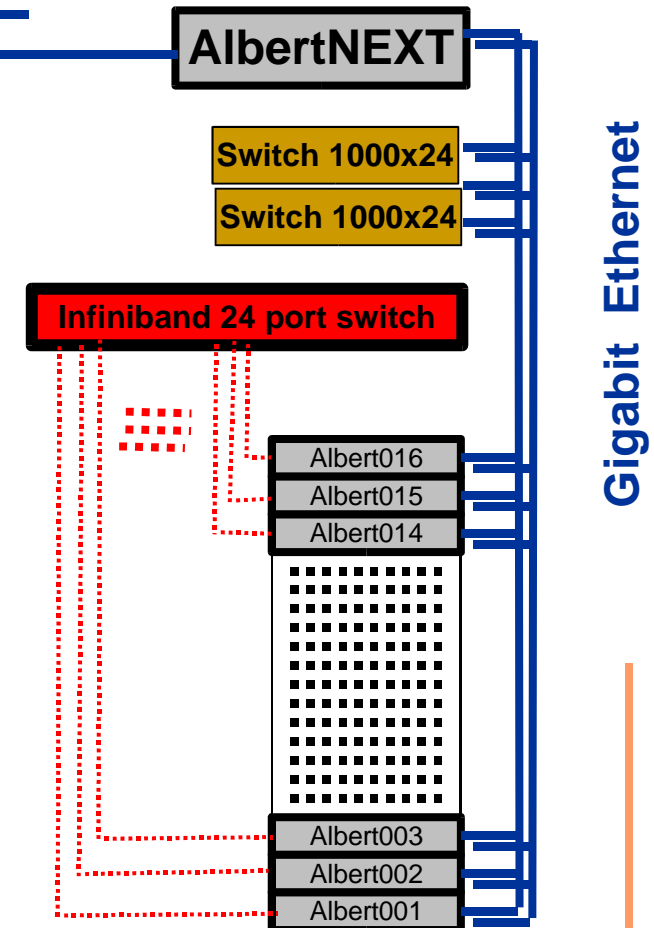
## *A less obvious example: numerical relativity*

*In numerical relativity things are less stable --->  
better to use (at best) traditional machines*



- ♦ *L. Baiotti, et al., Phys.Rev.Lett.94:131101,2005*
- ♦ *A. Nagar et al., Phys.Rev.D72:024007,2005*

University Network






# GRAPE

*GRAPE concentrates on computing the potential, the force, and its derivative.*

*It has been consistently the fastest computer on Earth*

*~ 70 Tflops equivalent (with some reasonable normalization factor)*

*Each element of the machine computes forces for  $N/P$  stars*



# SUE

*Monte Carlo simulation of an  
Ising spin glass is just a few lines  
of C code*

*SUE perform the complete  
computation in hardware*

*At the same time on many replicas*

*The original machine (~1998)  
updates on average one spin in  
~ 200 ps*

```
//-----  
void update(J)  
double J;  
{  
int i, j, tp;  
double deltaE, rndValue;  
double prob;  
  
for(i=1; i<=SIZE; i++) {  
  for(j=1; j<=SIZE; j++) {  
    deltaE =  
      (double)(2*spin[i][j]*(spin[i+1][j]+spin[i-1][j]+spin[i][j+1]+spin[i]  
[j-1]));  
    deltaE = J * deltaE;  
    if(e deltaE < 0.0) spin[i][j] = -spin[i][j];  
    else {  
      rndValue = (double) random() / ((double) RAND_MAX) ;  
      spin[i][j] = rndValue < exp(-deltaE) ? -spin[i][j] : spin[i][j];  
    }  
  } // end of for_j  
} // end of for_i  
}
```

$$H = - \sum_{\text{NB}(ij)} \sigma_i J_{ij} \sigma_j$$




# SUE

*Typical physics result: make sure you have brought to equilibrium a 203 lattice, and explore its phase structure (Ballesteros et al., cond-mat/0006211)*

*At the age of 7,  
SUE is roughly equivalent to  
a small-size high-end PC cluster  
(say 32 PC's) +  
state-of-the-art program writing*

*however ....*



## *superSUE / IANUS*

*A new scaled up version, to be ready in summer next year, will improve by a factor*

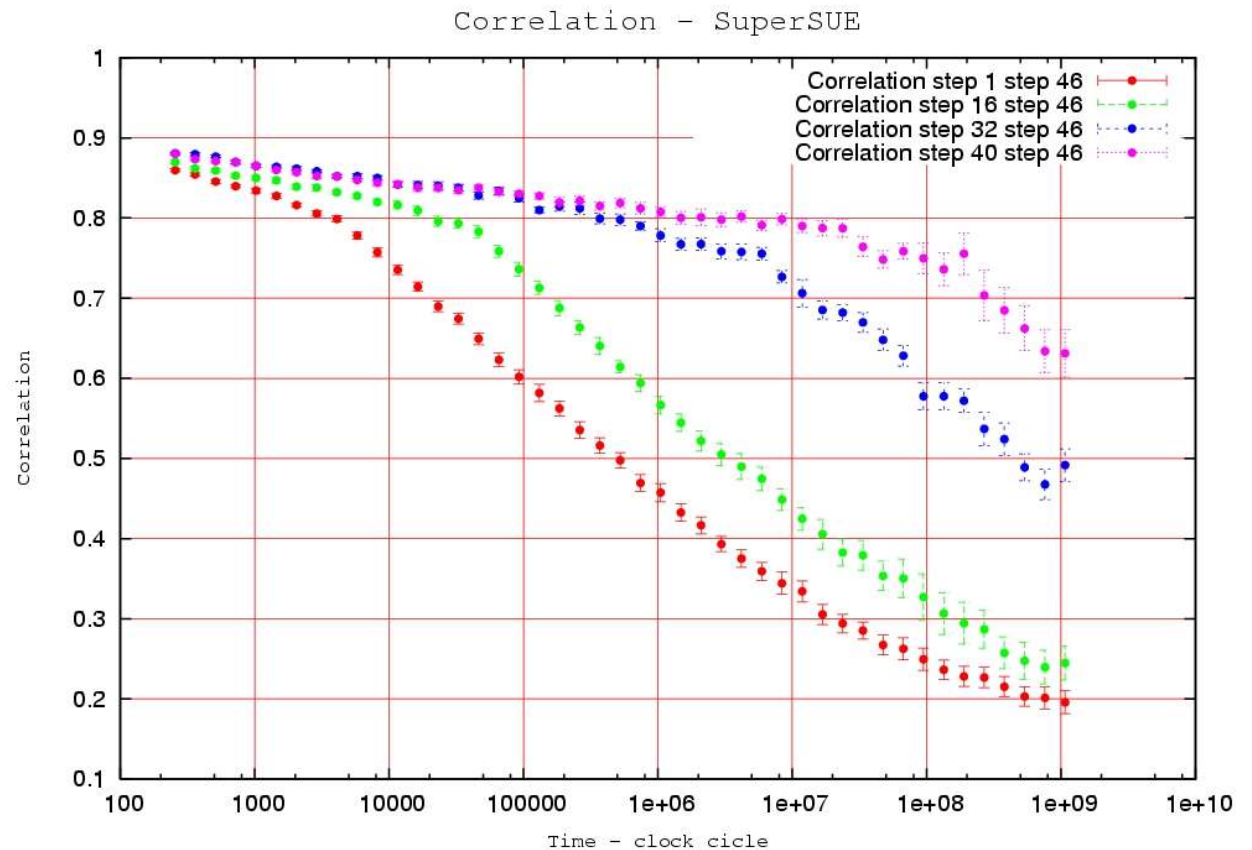
*~ 500,*

*breaking the 1 ps / spin-update barrier*

*that is ~ 10000 PC's working  
all together on the same problem*

*at the cost of ~ 300 PC's*

*preliminary tests already  
available ----->*



## *Conclusions*

*Quite often traditional computers do not do exactly what physics simulation needs most*

*And more often than not they do it the “wrong way” (from the physics point of view)*

*This provides an opportunity for physicists to build their own computing tools*

*Recent developments seem to signal that industry is slowly learning that it is better to follow physics than fighting against it (IBM Blue Gene/L)*

*so we may hope that computers for physics will be readily available soon ...*

*“... and I'll be able to do something else at last”*

## *Conclusions*

*Quite often traditional computers do not do exactly what physics simulation needs most*

*And more often than not they do it the “wrong way” (from the physics point of view)*

*This provides an opportunity for physicists to build their own computing tools*

*Recent developments seem to signal that industry is slowly learning that it is better to follow physics than fighting against it (IBM Blue Gene/L)*

*so we may hope that computers for physics will be readily available soon ...*

*“... and I'll be able to do something else at last” (RT, Lattice conference, Capri 1989)*