Sublattice Parallel Replica Dynamics (SLPRD)

Enrique Martinez, Arthur F. Voter and

Blas P. Uberuaga

Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

Exascale Computing

From

To

Exascale Computing

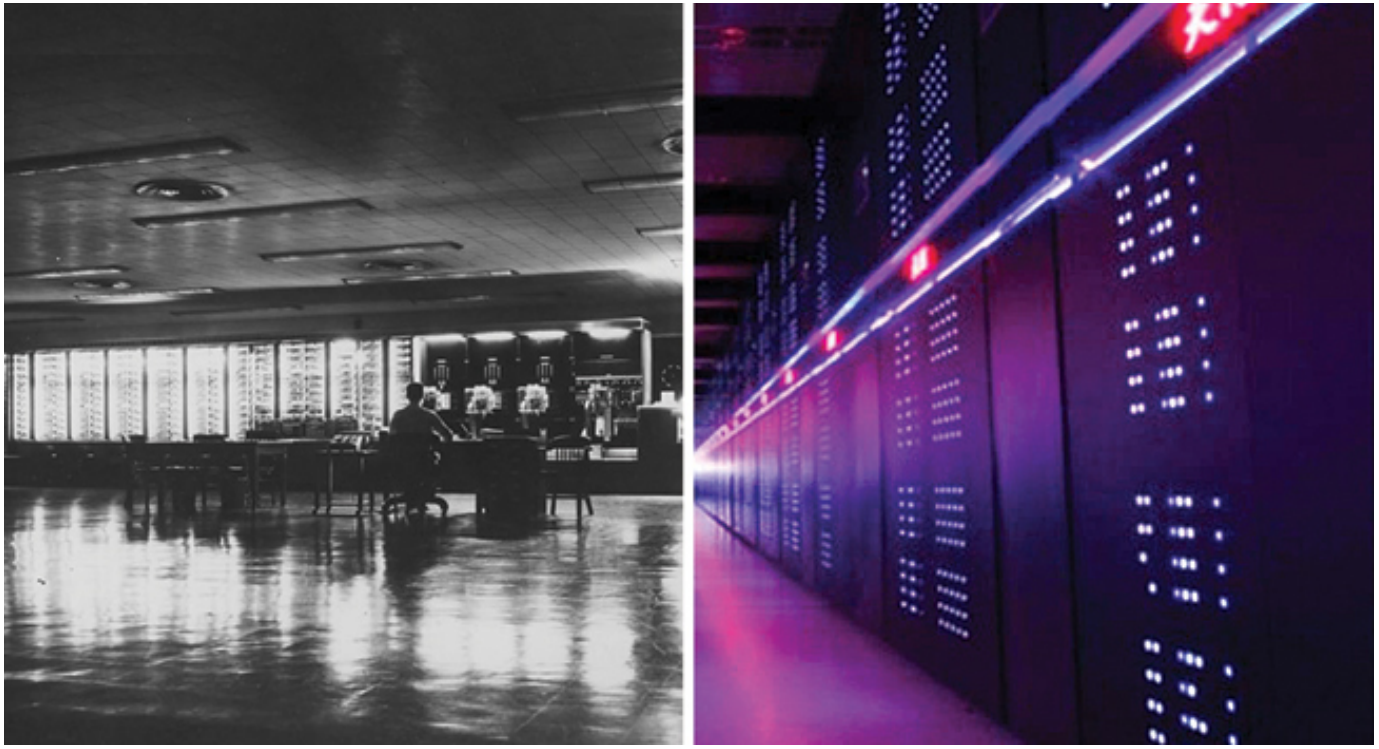

"Big" computers have come a long way since the completion of the Mark I (left) at Harvard in 1944. Capable of approximately **three calculations per second**, the Mark I was "the first operating machine that could execute long computations automatically," according to IBM. Today, the fastest supercomputer is the Tianhe-2 (right), developed by China's National University of Defense Technology. The Tianhe-2 has achieved almost **34 petaflops**, or 34,000,000,000,000,000 floating-point operations per second.

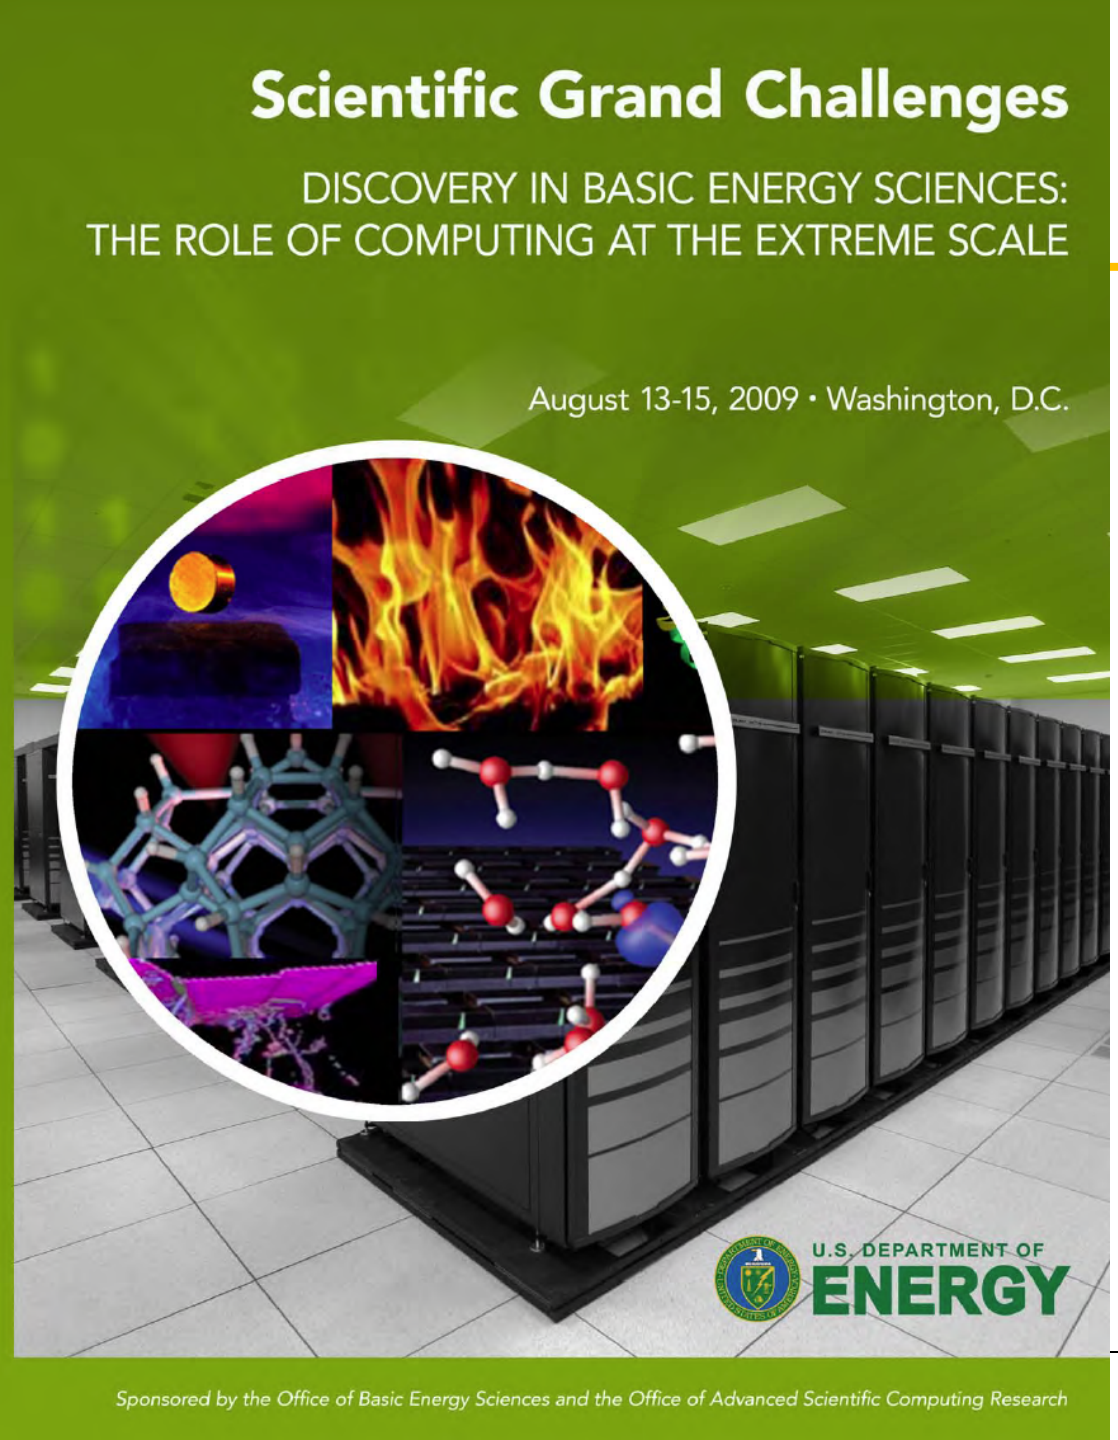

- Adaptation to regional **climate changes** such as sea level rise, drought and flooding, and severe weather patterns;
- **Reduction of the carbon footprint** of the transportation sector;
- Efficiency and safety of the **nuclear energy** sector;
- Innovative designs for cost-effective **renewable energy** resources such as batteries, catalysts, and biofuels;
- Certification of the **U.S. nuclear stockpile**, life extension programs, and directed stockpile work;
- **Design of advanced experimental facilities**, such as accelerators, and magnetic and inertial confinement fusion;
- First-principles understanding of the properties of **fission and fusion reactions**;
- Reverse engineering of the **human brain**;
- Design, control and manufacture of **advanced materials**.

NNSA

The mission and science opportunities in going to exascale are compelling.

For the growing number of problems where experiments are impossible, dangerous, or inordinately costly, extreme-scale computing will enable the solution of vastly more accurate predictive models and the analysis of massive quantities of data, producing quantum advances in areas of science and technology that are essential to DOE and Office of Science missions. For example, exascale computing will push the frontiers of

- Adaptation to regional climate changes such as sea level rise, drought and flooding, and severe weather patterns;
- Reduction of the carbon footprint of the transportation sector;
- Efficiency and safety of the nuclear energy sector;
- Innovative designs for cost-effective renewable energy resources such as batteries, catalysts, and biofuels;
- Certification of the U.S. nuclear stockpile, life extension programs, and directed stockpile work;
- Design of advanced experimental facilities, such as accelerators, and magnetic and inertial confinement fusion;
- First-principles understanding of the properties of fission and fusion reactions;
- Reverse engineering of the human brain;
- Design, control and manufacture of advanced materials.

U.S. economic competitiveness will also be significantly enhanced as companies utilize accelerate development of superior new products and spur creativity arising from modeling and simulation at unprecedented speed and fidelity. An important ingredient is the boost to continued international leadership of the U.S.-based information technology industry at all levels, from laptop to exaflop.

Making the transition to exascale poses numerous unavoidable scientific and technological challenges.

Every major change in computer architecture has led to changes, usually dramatic and often unanticipated, and the move to exascale will be no exception. At the hardware level, feature size in silicon will almost certainly continue to decrease at Moore's Law pace at least over the next decade. To remain effective in high-end computing systems and in consumer electronics, computer chips must change in radical ways, such as containing 100 times more processing elements than today. Three challenges to be resolved are:

- **Reducing power requirements.** Based on current technology, scaling today's systems to an exaflop level would consume more than a gigawatt of power, roughly the output of Hoover Dam. Reducing the power requirement by a factor of at least 100 is a challenge for future hardware and software technologies.
- **Coping with run-time errors.** Today's systems have approximately 500,000 processing elements. By 2020, due to design and power constraints, the clock frequency is unlikely to change, which means that an exascale system will have approximately one billion processing elements. An immediate consequence is that the frequency of errors will increase (possibly by a factor of 1000) while timely identification and correction of errors become much more difficult.
- **Exploiting massive parallelism.** Mathematical models, numerical methods, and software implementations will all need new conceptual and programming paradigms to make effective use of unprecedented levels of concurrency.

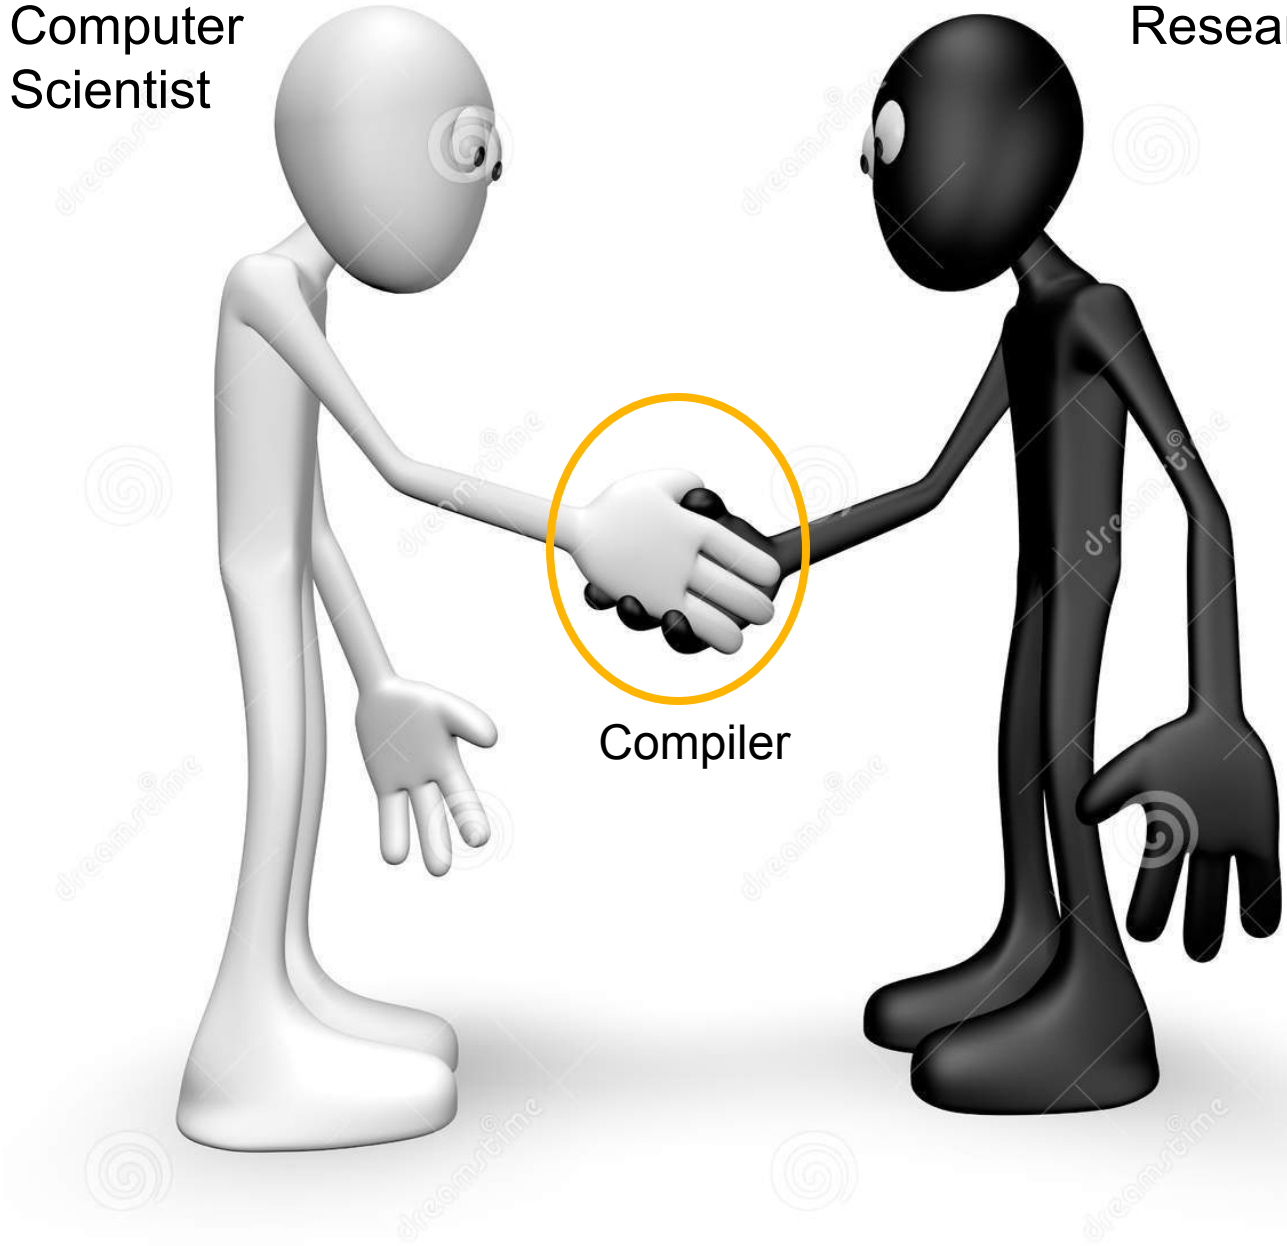
Computer Scientist
Researcher
Compiler

Fully atomistic models

- **Classical Molecular Dynamics**
  - Solves Newton equations of motion

$$m\ddot{r} = -\nabla U + f$$

  - All the physics are included in the empirical potential
  - Extremely parallelizable in space for large systems
  - For small systems the parallel efficiency drops
- **Parallel Replica Dynamics**
  - Parallelizes time evolution
  - Assumptions:
    - Infrequent events
    - Exponential distribution of first-escape times
- **Parareal**
  - Parallelizes time evolution

Traditional Parallel Replica Dynamics

The properties of the exponential allow us to parallelize time, by having many processors seek the first escape event. The procedure:

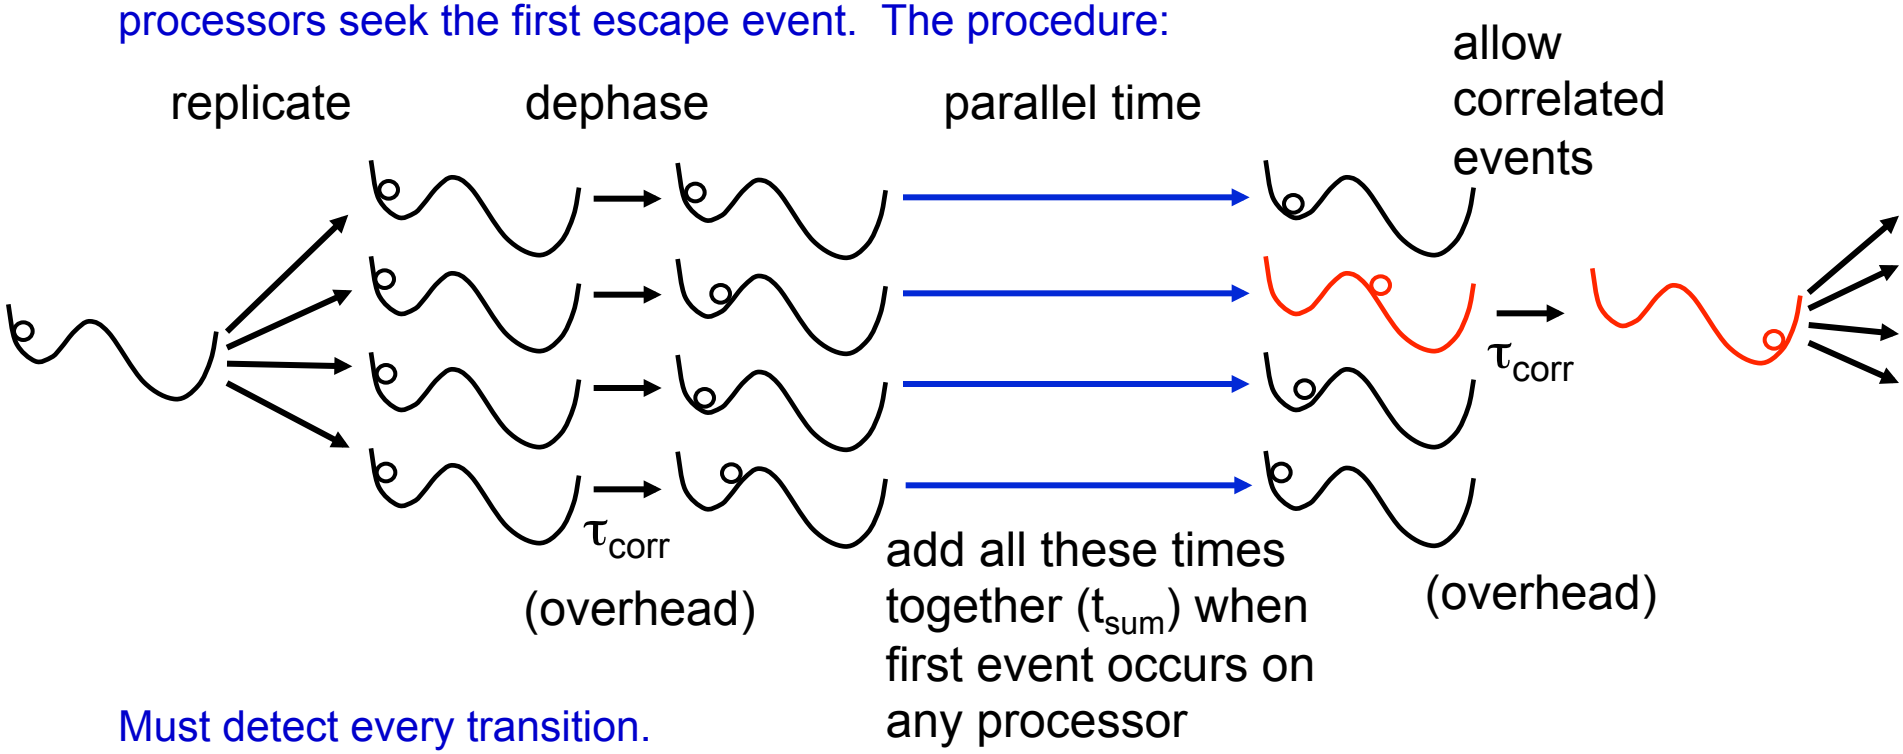

Must detect every transition.

AFV, Phys. Rev. B 57, 13985 (1998).

The Exascale Challenge

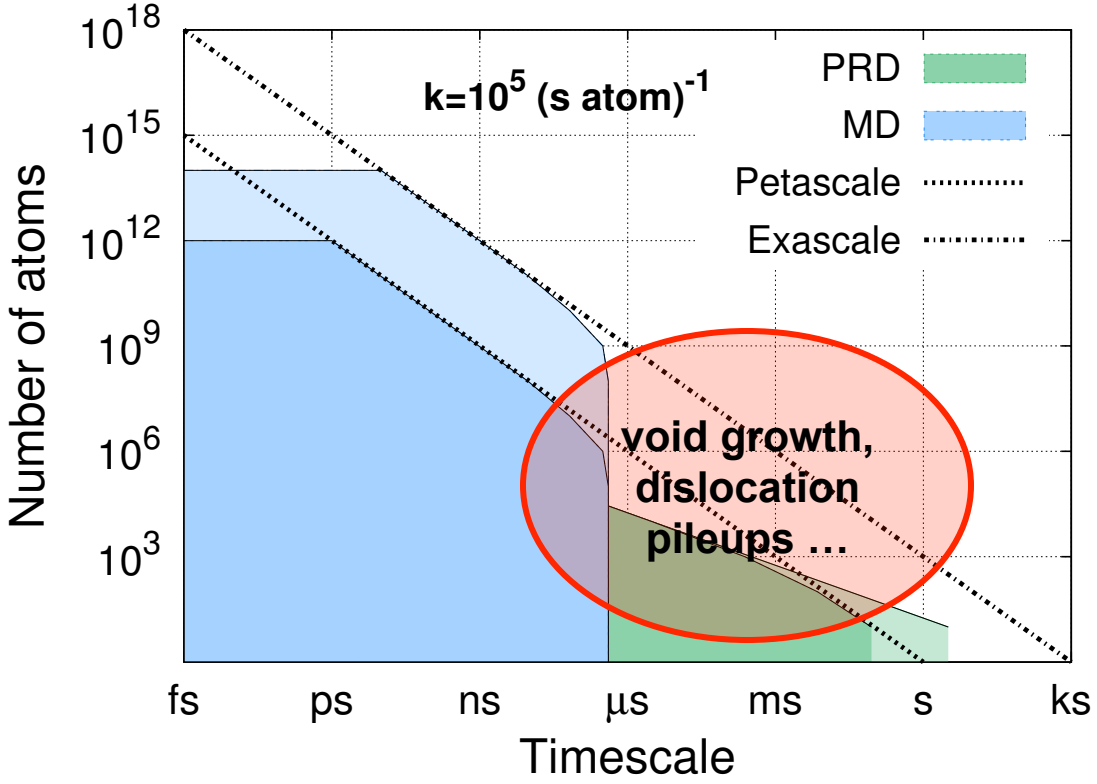

PRD curves assume an average rate of 10^5 events per atom.second.

Petascale = 10^4 processors

Exascale = 10^7 processors

PRD uses 1 proc per replica.

Even with PRD, we have a challenge at the exascale (dephasing overhead dominates for large nproc)

The Exascale Challenge

Exascale Challenge

Speedup:

$$\eta_{ParPRD} = N_{PP}\delta(A_R)\frac{R_S \cdot x \cdot t_S}{x \cdot t_S(1+\frac{t_M}{t_B}) + \overline{n}_S \cdot (t_D + t_C)}$$

Spatially Parallelized Efficiency (MD)

Time Parallelized Efficiency (PRD)

As the number of processors increases the Spatially Parallelized Efficiency decreases

As the number of particles or the number of replicas goes up, the average number of events increases, PRD Efficiency decreases

Unreachable part of the exascale triangle

Pushing PRD a bit Further

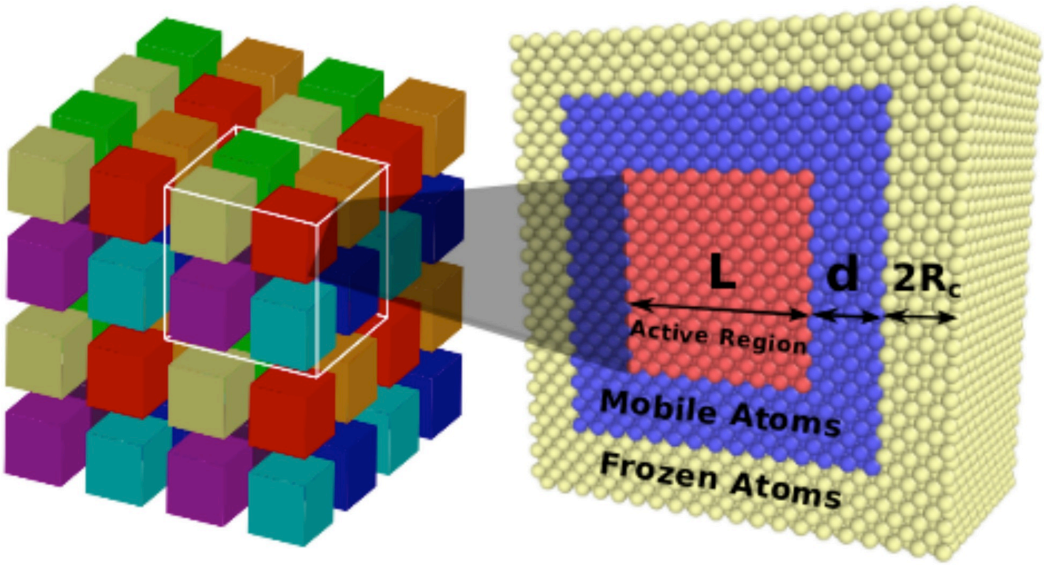

Can we arrange the processors in a different way so to improve the efficiency?

E. Martinez et. al, Journal of Computational Physics 230 (2011) 1359–1369
Y. Shim, J.G. Amar, Physical Review B 71 (2005) 115436
Y. Shim et al. Physical Review B **76**, 205439 2007

Sublattice PRD (SLPRD)

Speedup:

$$\eta_{SLPRD} = N_{PP}\delta(A_R)D\frac{R \cdot x \cdot t_S}{x \cdot t_S(1+\frac{t_M}{t_B}) + \overline{n}_{\max} \cdot (t_D + t_C)}$$

Spatially Parallelized Efficiency (MD)

Time Parallelized Efficiency (PRD)

If the number of events is considered a stochastic variable Poisson distributed, the average of the maximum number of events can be calculated analytically

$$\overline{n} = \sum_n n P_n^{exact} \qquad P_n^{exact} = (A + P_n)^D - A^D \qquad A = P_0 + P_1 + \cdots + P_{n-1}$$

$$P_i = \frac{\left(N_{at} k t_{adv}\right)^i \cdot \exp\left\{-N_{at} k t_{adv}\right\}}{i!}$$

Validation

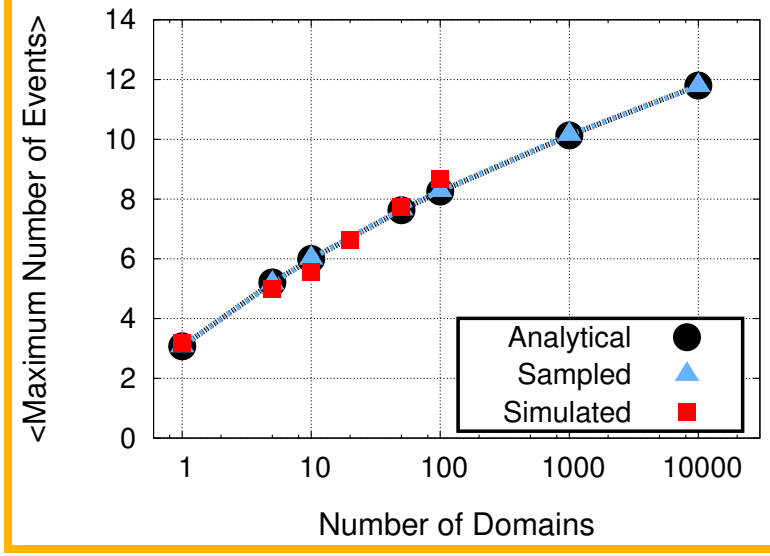

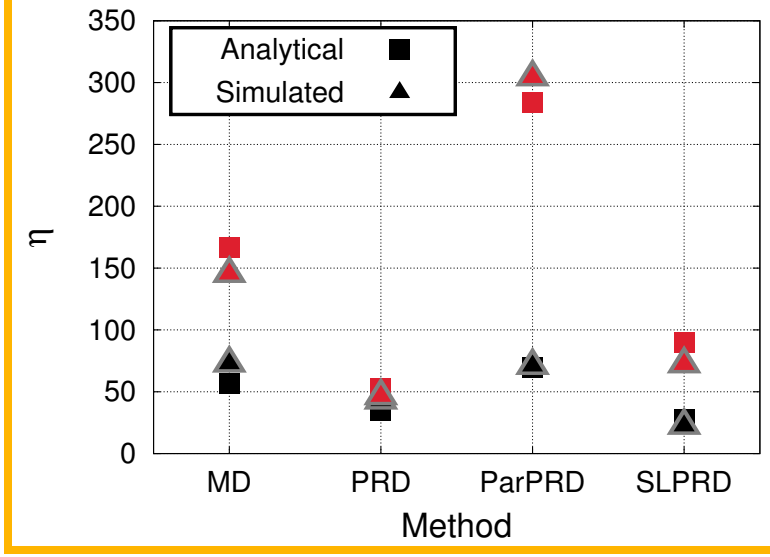

Checking expression for $\overline{n}_{\max}$

- 1 Vacancy in 107 Ni atoms at T=1022 K gives an average event rate of 1.07 10^7 $(s\ atom)^{-1}$.
- Run dynamics for 2.688 ns.

Checking η

- 1 Vacancy in 107 Ni atoms at T=723 K gives an average rate of $k=10^5$ $(s\ atom)^{-1}$.
- Run dynamics for 192 ps
- Good agreement between theoretical expressions and simulation results
- Black points for 96 procs and red points for 512 procs

Speedup with the number of processors

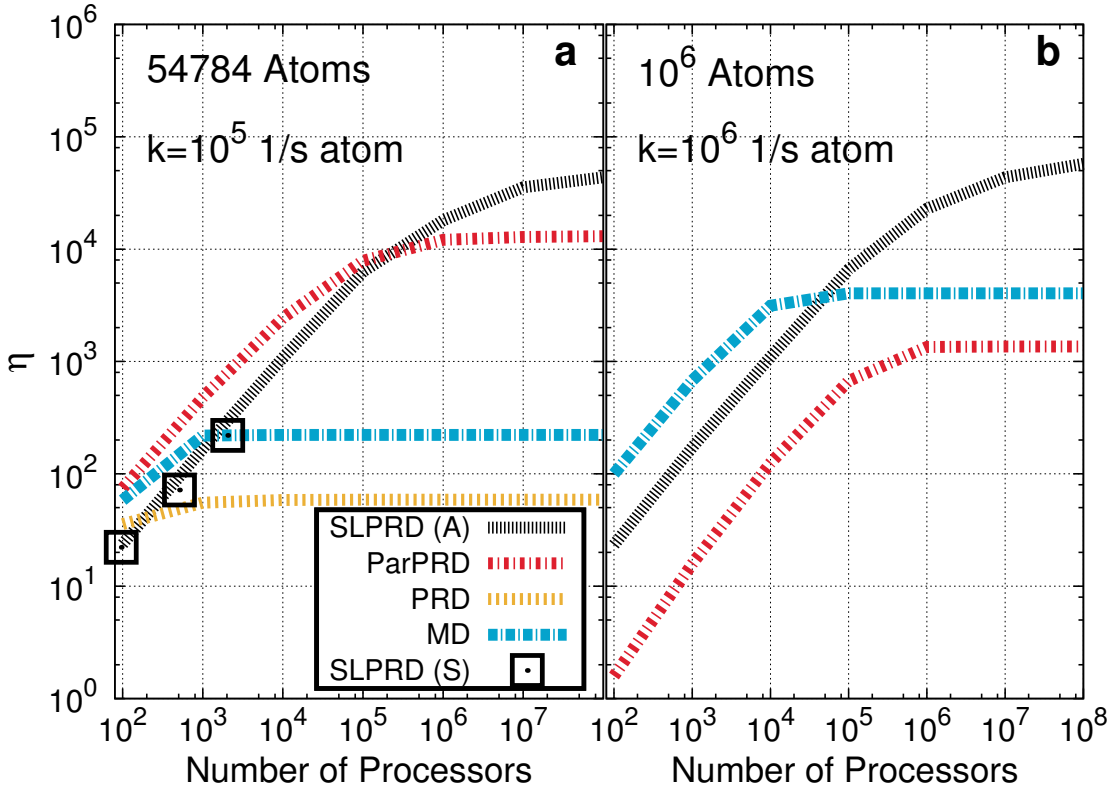

For a large numbers of processors SLPRD becomes more efficient than existing methodologies

Speedup with the number of atoms

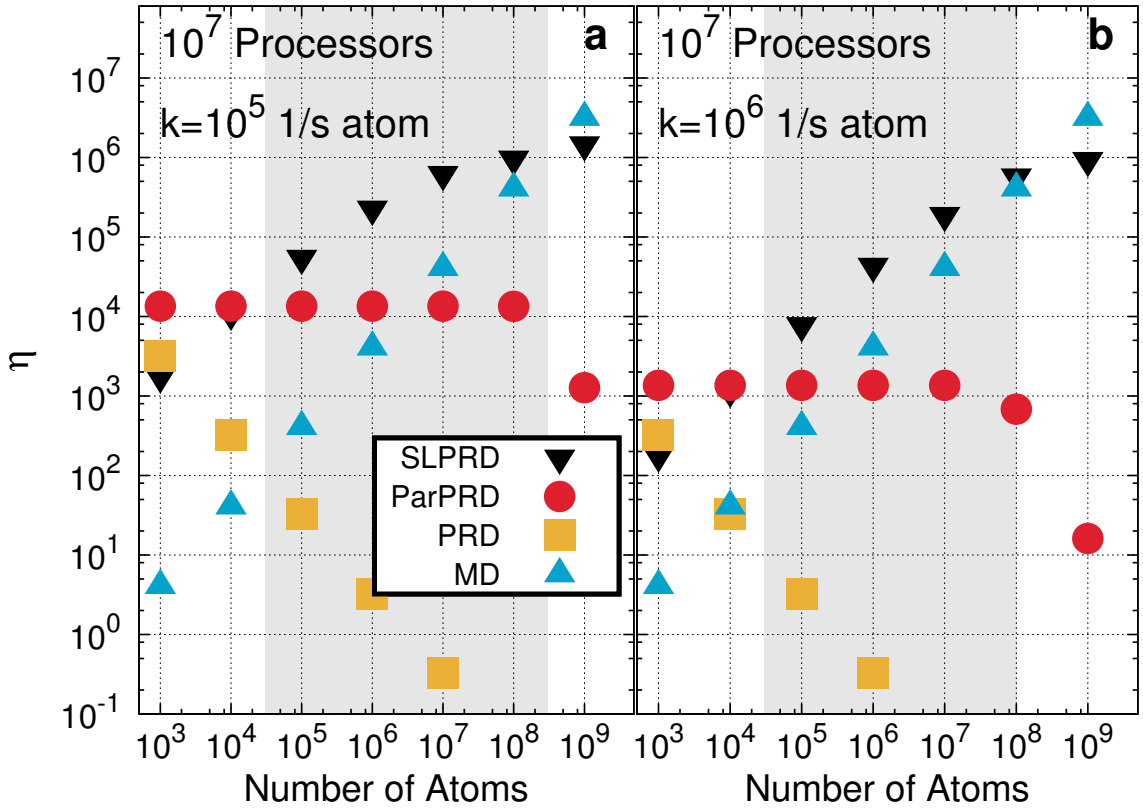

For an intermediate numbers of atoms SLPRD becomes more efficient than existing methodologies

Crossover in efficiency

Pool of processors needed for the SLPRD algorithm to become the most efficient

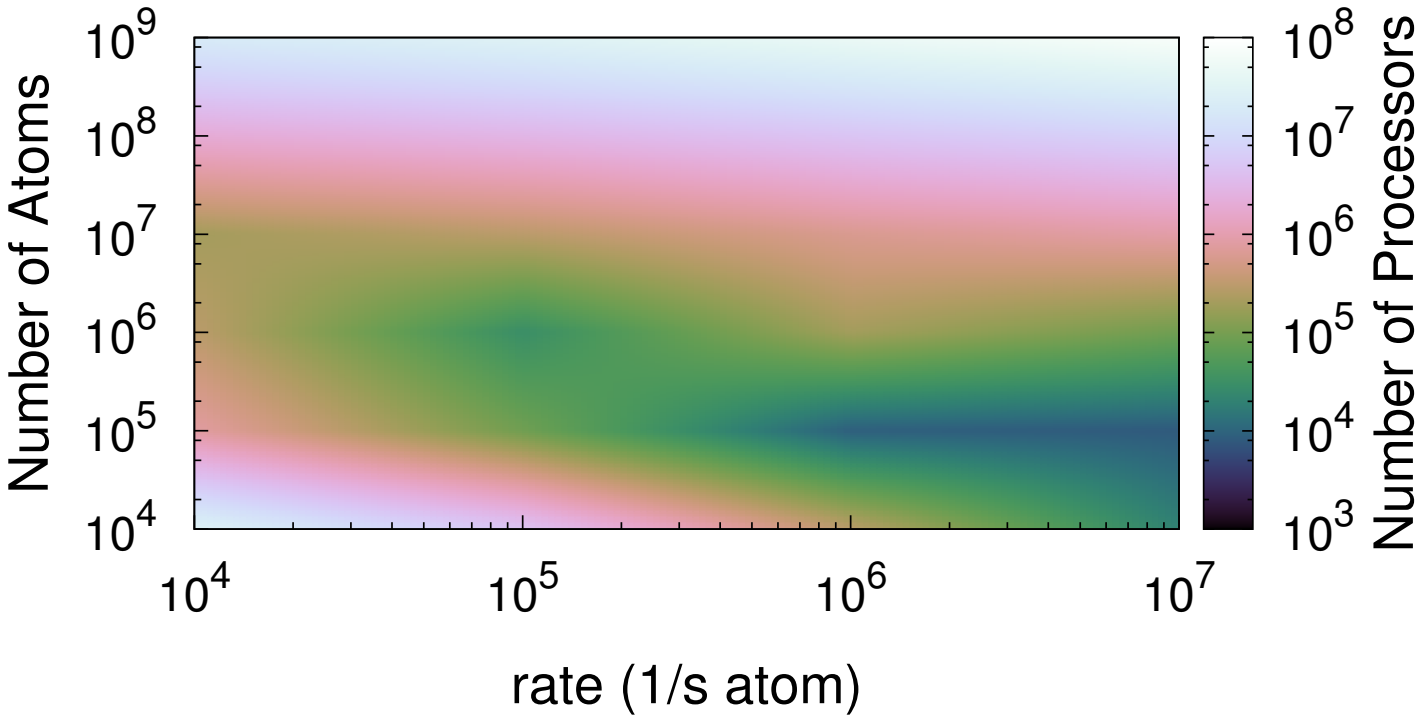

There is a minimum in the number of processors for a given number of atoms and an average rate

SLPRD optimum parameters

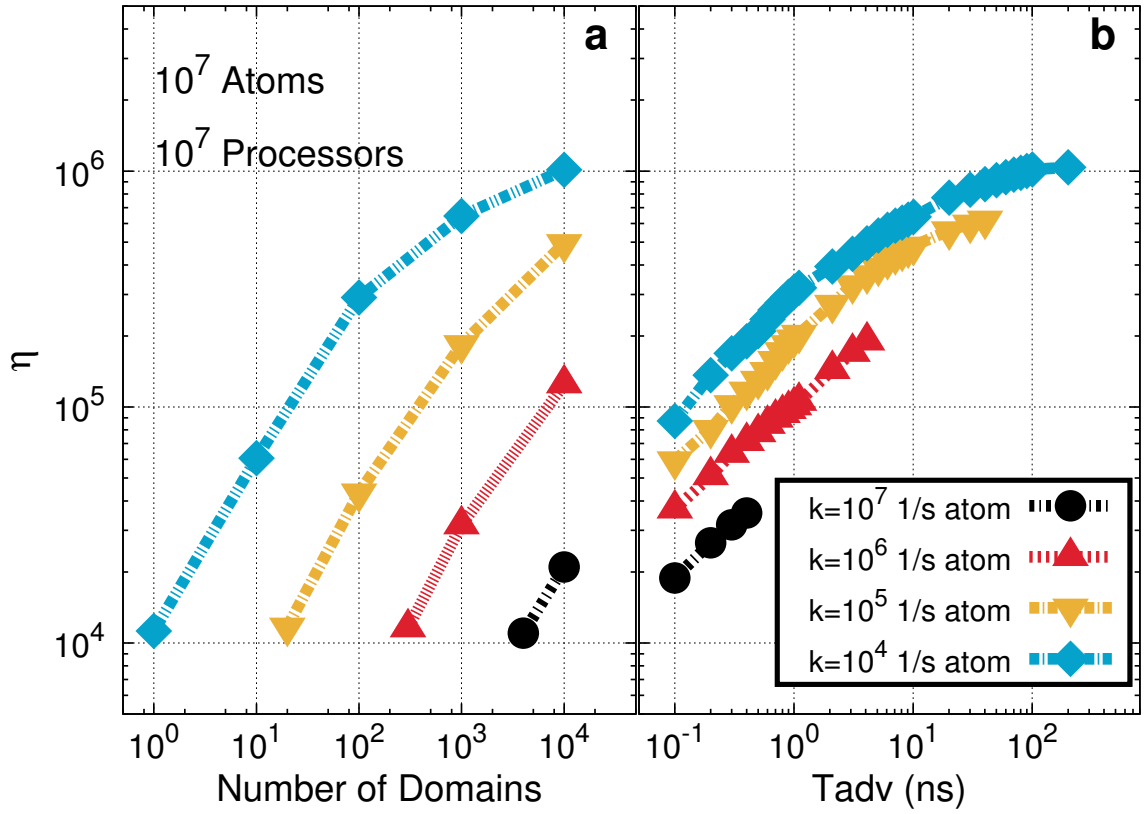

We can predict the simulation conditions to have optimum scalability

The Exascale Challenge Revisited

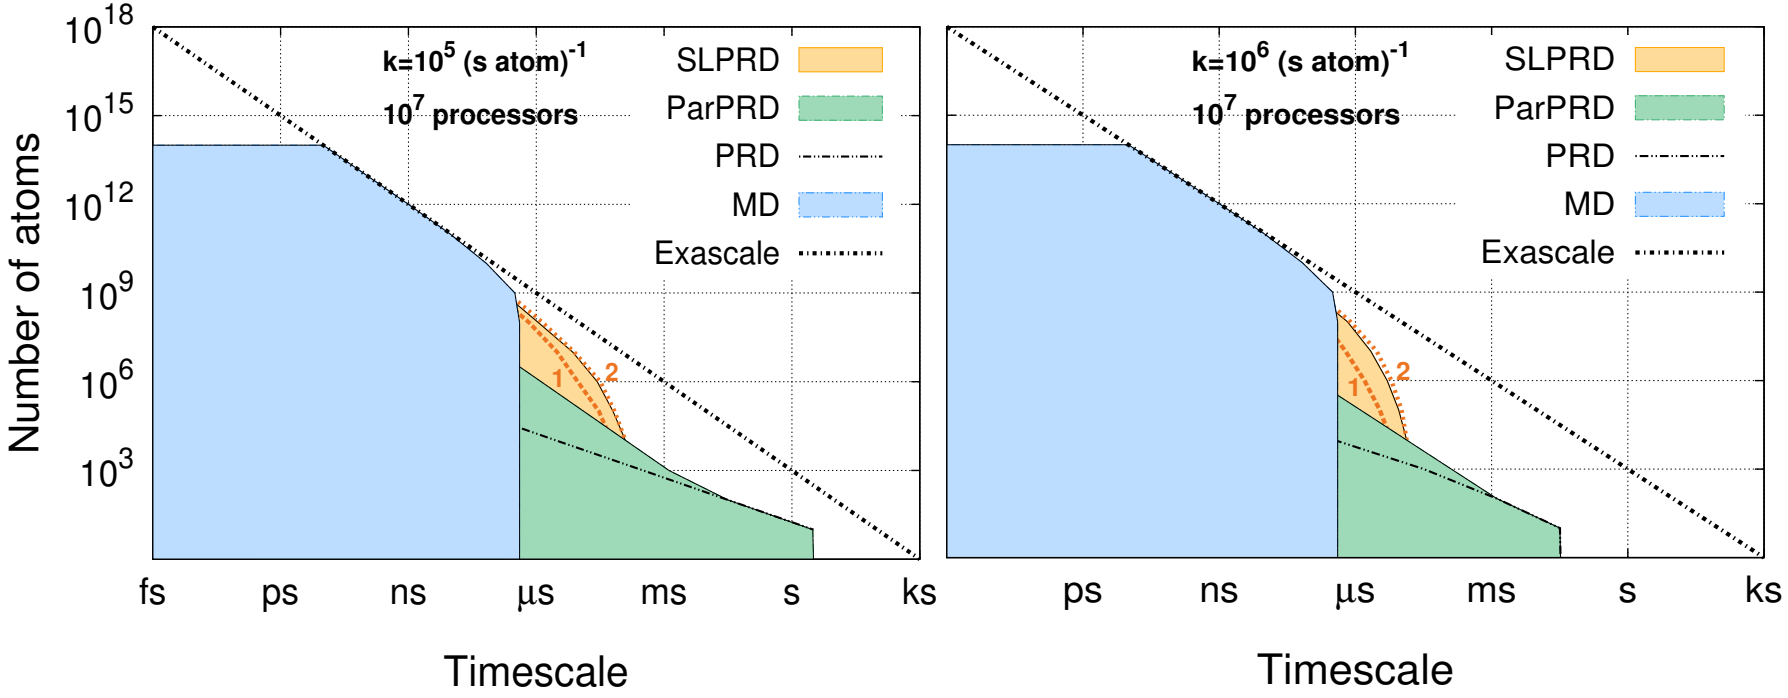

SLPRD fills part of the exascale hole

Conclusions

- **An increase in computational power will have to come along with new algorithms.**
- **A strong interaction between researchers and computational scientists will have to be developed in order to take advantage of the new architectures and programming paradigms.**