ERLANGEN REGIONAL
COMPUTING CENTER

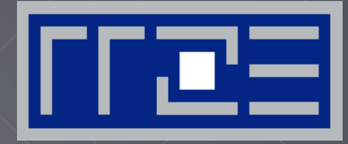

http://tiny.cc/NLPE-MdlS

Node-Level Performance Engineering

Georg Hager, Jan Eitzinger
Erlangen Regional Computing Center (RRZE)
University of Erlangen-Nuremberg

Three-day short course
Maison de la Simulation, Saclay, France
2018-06-11/12/13

Course Schedule

http://tiny.cc/NLPE-MdlS

Day 1	
Welcome & intro	GH
Computer architecture	GH
Tools: topology, affinity, clock speed	GH
Microbenchmarking for architectural exploration (and more)	GH
The Roofline performance model: basics	GH
Optimal use of parallel resources: ccNUMA	GH
Day 2	
Tools: hardware performance counters	JE
Optimal use of parallel resources: SIMD	JE
Performance Engineering with patterns	JE
Roofline case study: Jacobi smoother	GH
Roofline case study: sparse matrix-vector multiplication	GH
Hands-On exercises	
Day 3	
Hands-On code work	

Cycle gymnastics

http://tiny.cc/NLPE-MdlS

- 1 cycle = smallest unit of time on a CPU ("heartbeat")
  - Clock speed of typical CPU: 3.0 Gcy/s (or GHz)
- Basic unit of work: Floating-point operation (Flop)
  - Typical peak performance of 8-core CPU: P_{peak} = 192 Gflop/s
  - How many Flops per cycle per core is that? $\frac{192 \cdot 10^9 \frac{Flops}{s}}{8\ cores \cdot 3.0 \cdot 10^9 \frac{cy}{s}} = 8 \frac{Flops}{cy \cdot core}$
  - Typical duration of a double precision multiply: 5 cycles
    - How much time is that? $\frac{5\ cy}{3.0 \cdot 10^9 \frac{cy}{s}} = 1.67 \cdot 10^{-9} s = 1.67\ ns$
- Basic unit of traffic: Byte
- Unit of bandwidth: Bytes/s
  - Typical memory bandwidth: 48 Gbytes/s = $4.8 \cdot 10^{10}$ Bytes/s
  - How many bytes per cycle is that? $\frac{48 \cdot 10^9 \frac{Bytes}{s}}{3.0 \cdot 10^9 \frac{cy}{s}} = 16 \frac{Bytes}{cy}$

PRELUDE: SCALABILITY 4 THE WIN!

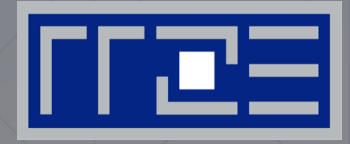

How to ask the right questions

A conversation

From a student seminar on “Efficient programming of modern multi- and manycore processors”

Student: I have implemented this algorithm on the GPGPU, and it solves a system with 26546 unknowns in 0.12 seconds, so it is really fast.

Me: What makes you think that 0.12 seconds is fast?

Student: It is fast because my baseline C++ code on the CPU is about 20 times slower.

Scalability Myth: Code scalability is the key issue

```
!$OMP PARALLEL DO
do k = 1 , Nk
  do j = 1 , Nj; do i = 1 , Ni
     y(i,j,k)= b*(  x(i-1,j,k)+ x(i+1,j,k)+ x(i,j-1,k)+
                    x(i,j+1,k)+ x(i,j,k-1)+ x(i,j,k+1))
    enddo; enddo
enddo
!$OMP END PARALLEL DO
```

Changing only a the compile options makes this code scalable on an 8-core chip

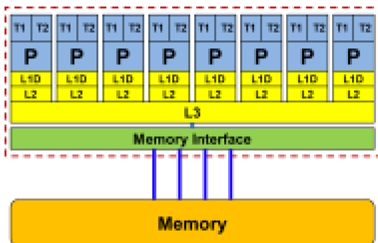

3D Stencil Update ("Jacobi")

Version 1
Version 2

-O0

-O3 -xAVX

Prepared for the highly parallel era!

Speed-Up

#cores

1 2 3 4 5 6 7 8

Scalability Myth: Code scalability is the key issue

```
!$OMP PARALLEL DO
do k = 1 , Nk
  do j = 1 , Nj; do i = 1 , Ni
     y(i,j,k)= b*(  x(i-1,j,k)+ x(i+1,j,k)+ x(i,j-1,k)+
                    x(i,j+1,k)+ x(i,j,k-1)+ x(i,j,k+1))
    enddo; enddo
enddo
                     L DO
```

Upper limit from simple performance model:
35 GB/s & 24 Byte/update

Performance [MLUP/s]

1500
1000
500
0

Version 1
Version 2

3D Stencil Update
("Jacobi")

Single core/socket efficiency
is key issue!

1 2 3 4 5 6 7 8
#cores

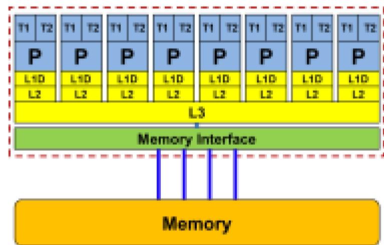

How model-building works: Physics

Newtonian mechanics

$$\vec{F} = m\vec{a}$$

Fails @ small scales!

Nonrelativistic quantum mechanics

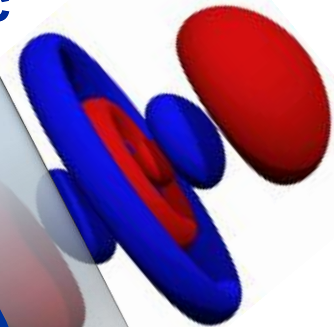

$$i\hbar\frac{\partial}{\partial t}\psi(\vec{r},t) = H\psi(\vec{r},t)$$

Fails @ even smaller scales!

If a model fails, we learn something!

Relativistic quantum field theory

$$U(1)_Y \otimes SU(2)_L \otimes SU(3)_c$$

Performance Engineering – no black boxes!

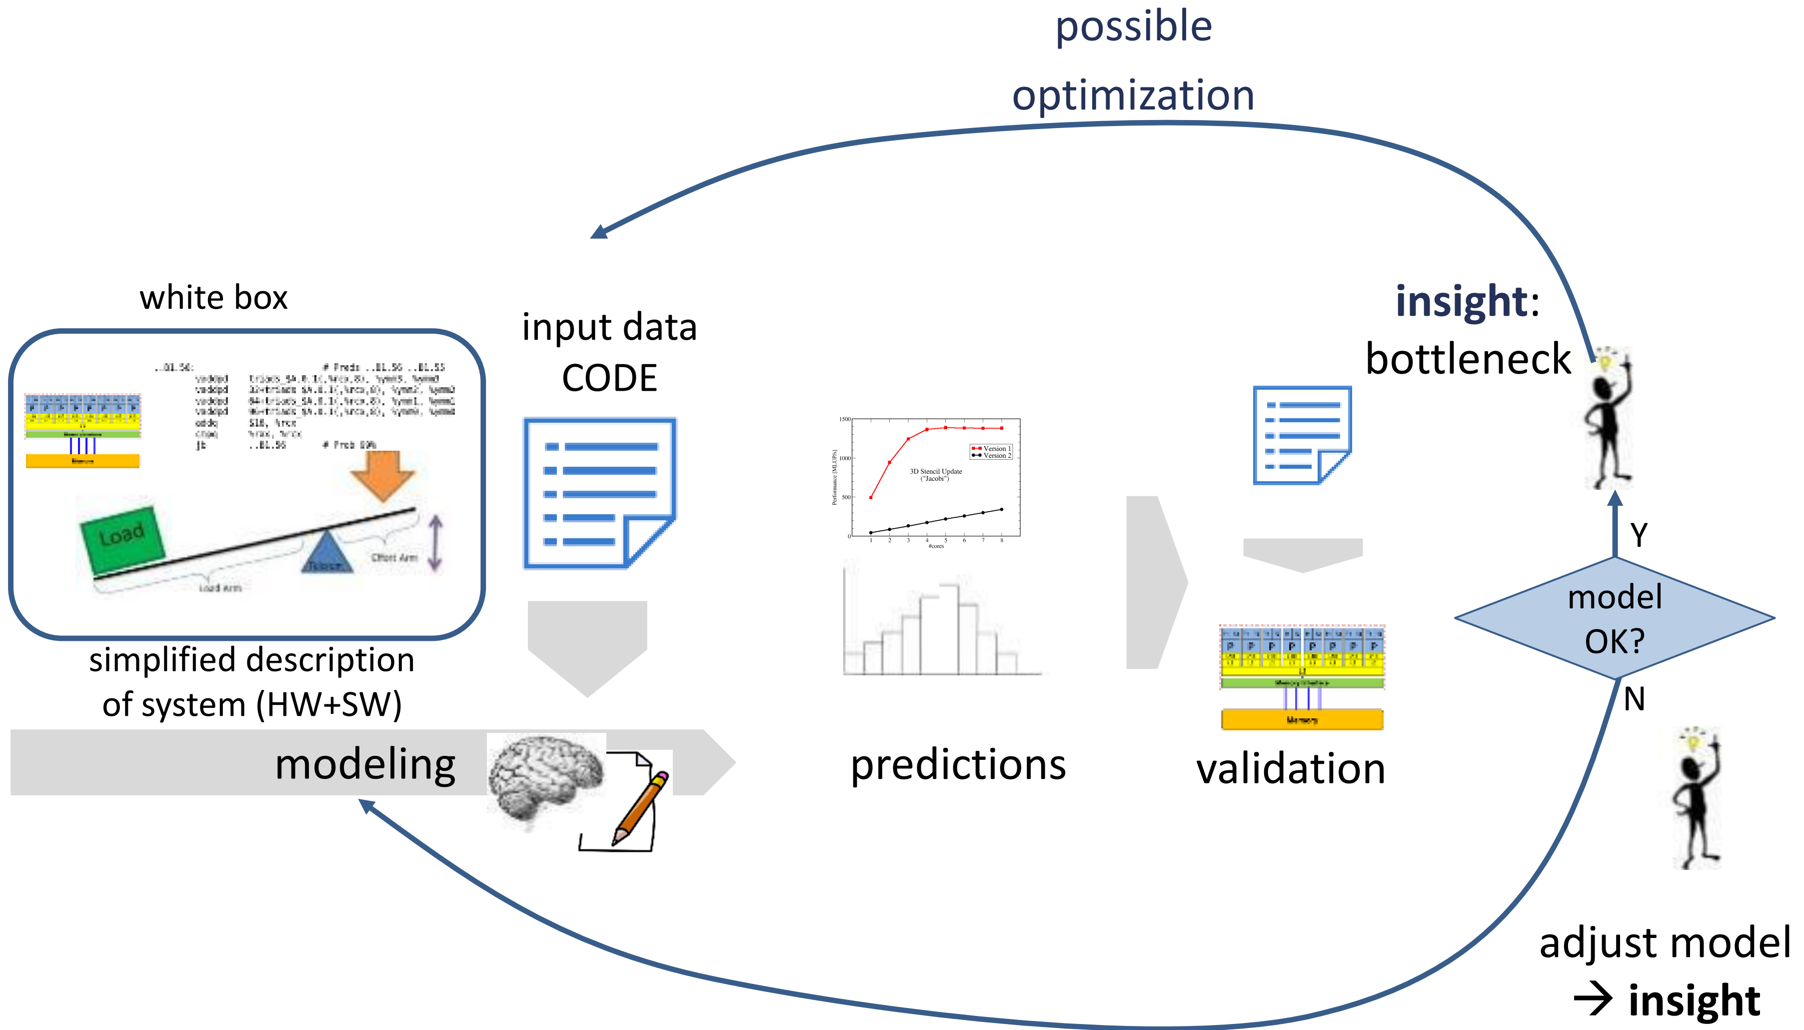

Questions to ask in high performance computing

- **Do I understand the performance behavior of my code?**
  - Does the performance match a model I have made?
- **What is the optimal performance for my code on a given machine?**
  - High Performance Computing == Computing at the bottleneck
- **Can I change my code so that the "optimal performance" gets higher?**
  - Circumventing/ameliorating the impact of the bottleneck
- **My model does not work – what's wrong?**
  - This is the good case, because you learn something
  - Performance monitoring / microbenchmarking may help clear up the situation
- **Use your brain! Tools may help, but you do the thinking.**