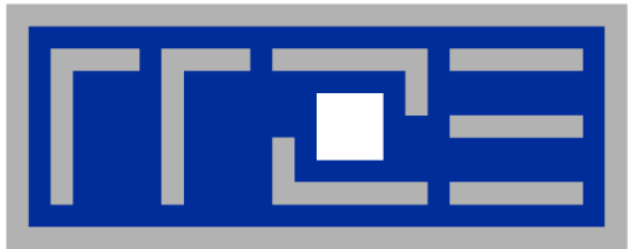

Energy efficiency: A down-to-earth perspective

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

Cool Supercomputing BoF @ SC12, Nov 14, 2012

Assumptions & naive conclusions

1. **The lifetime of a typical machine is constant (4-6 years)**
2. **Energy costs account for a significant fraction of TCO (especially in Europe)**
3. **Machines are almost 100% utilized**
4. **Domain scientists have no idea about**
   - Performance optimization
   - Connection between performance and power bill for *their* jobs

Straightforward conclusions:

- **Install automatic mechanisms to automagically clock down CPUs in apps not sensitive to clock speed**
- **Use "application slack" to clock down/power down individual cores**

My point of view: Get the low-hanging fruits first!

Application optimization is
the first and easiest way to save energy

Example:

A medical image reconstruction code on Sandy Bridge

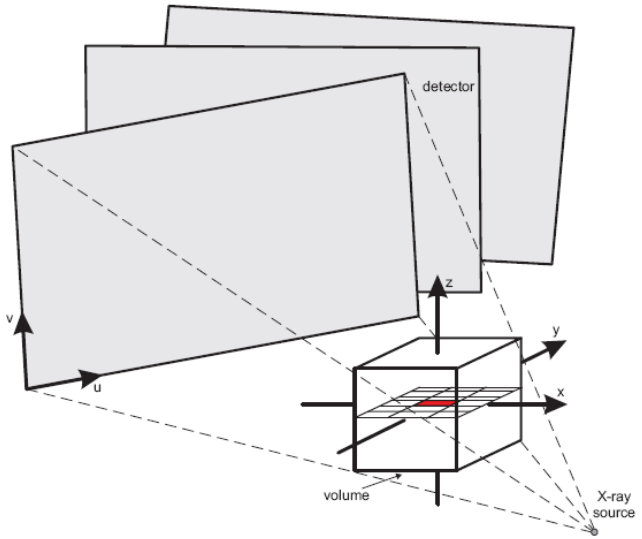

Sandy Bridge EP (8 cores, 2.7 GHz base freq.)

Test case	Runtime [s]	Power [W]	Energy [J]
8 cores, plain C	**90.43**	90	8110
8 cores, SSE	29.63	93	2750
8 cores (SMT), SSE	22.61	102	2300
8 cores (SMT), AVX	**18.42**	111	2040

Faster code ➘ less energy

Get the low-hanging fruits first!

Load imbalance is better removed from the start

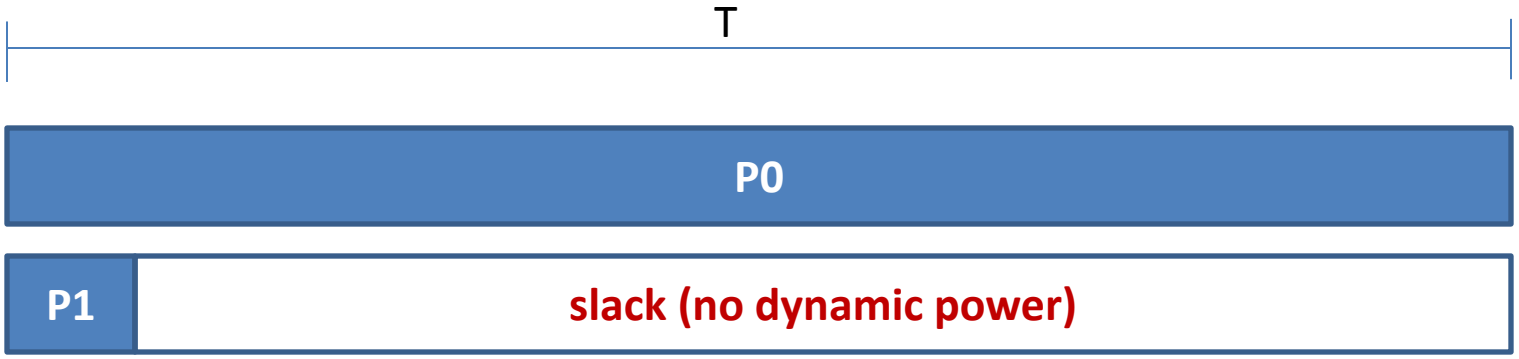

Energy to solution: $E_{slack} \approx T \cdot \left(W_{static} + W_{dyn}\right) = T \cdot W_{static} + T \cdot W_{dyn}$

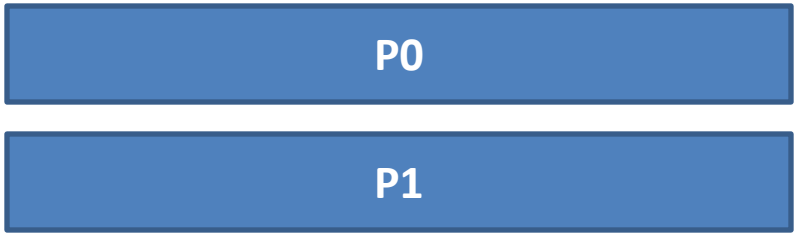

Energy to solution: $E_{opt} \approx \frac{T}{2} \cdot \left(W_{static} + 2 \cdot W_{dyn}\right) = \frac{T}{2} \cdot W_{static} + T \cdot W_{dyn}$

Performance and energy to solution of LBM solver

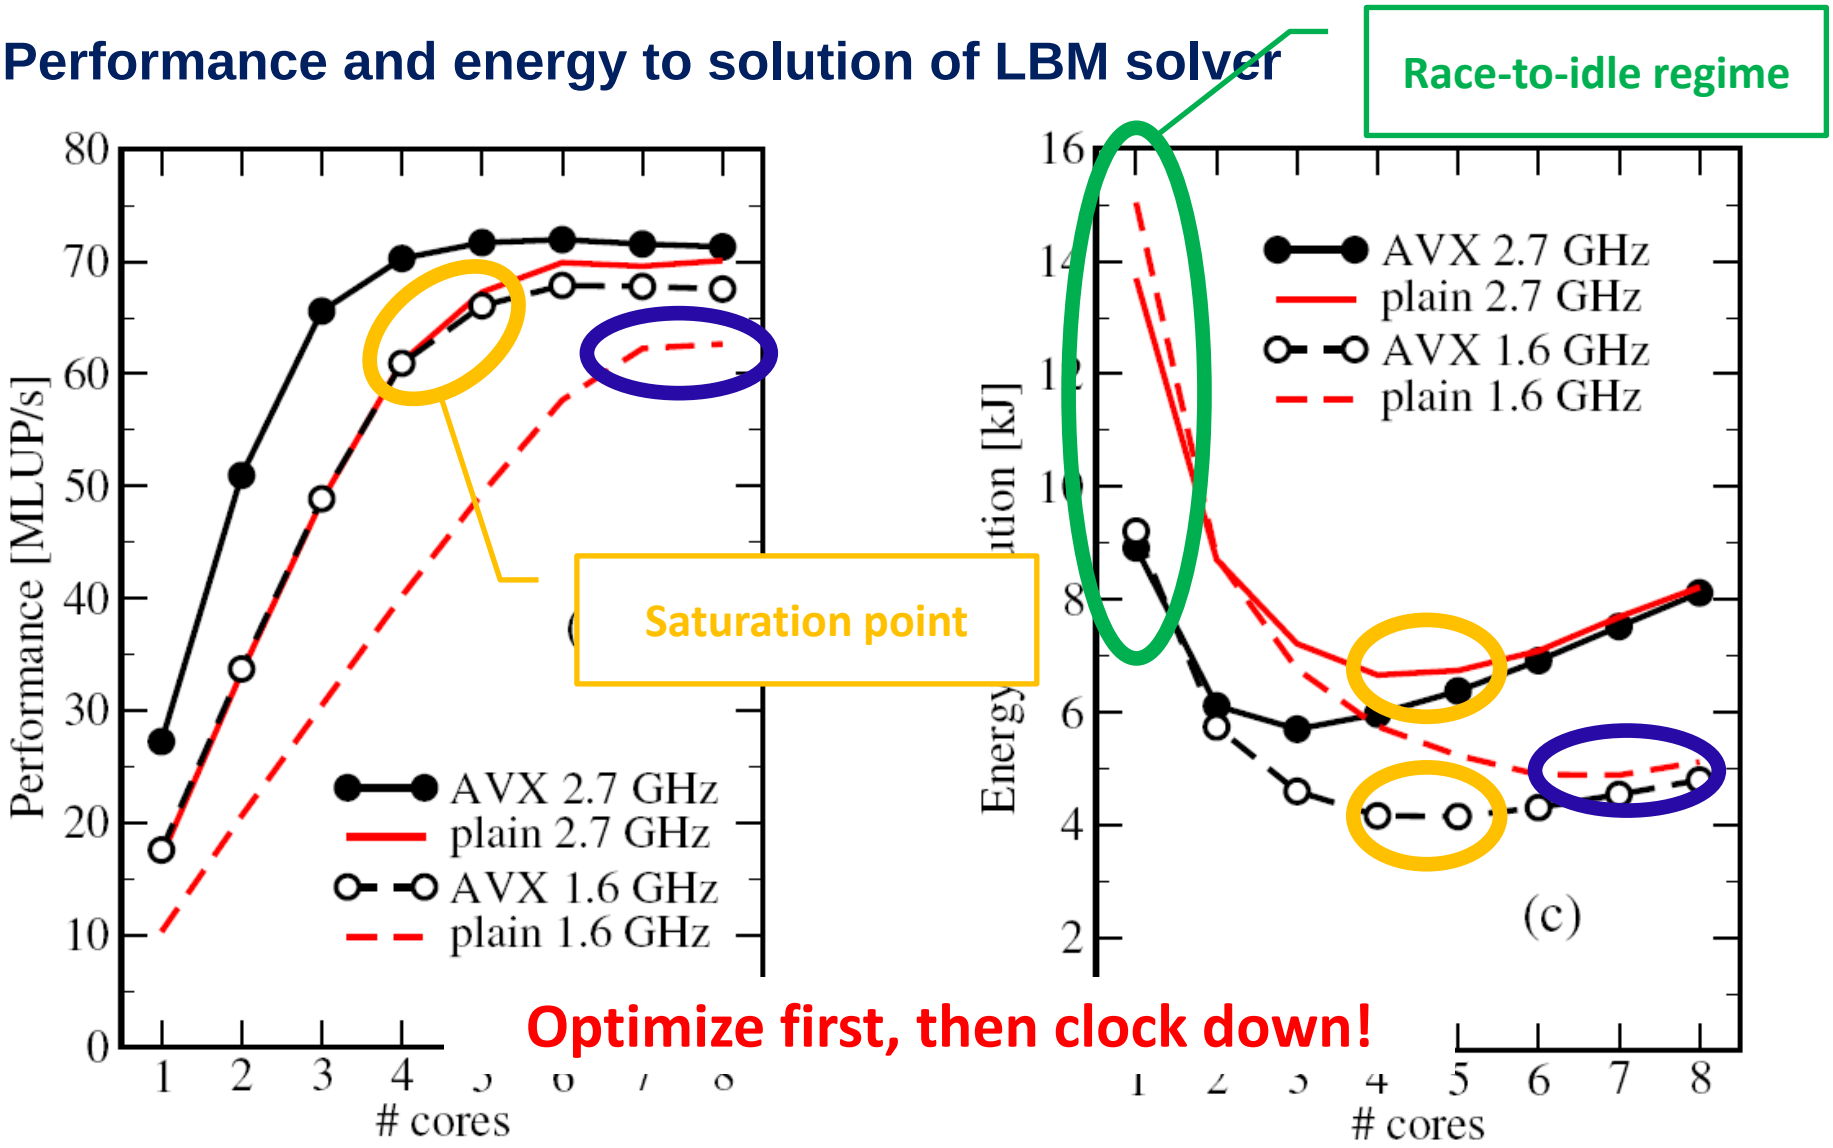

Optimize first, then clock down!

Conclusions

Low-hanging fruits for power efficiency

- **Think about a "Science per Joule" metric**
- **Remove load imbalance for better resource utilization**
- **Single-core (and then parallel) optimization for**
  - Shorter time to solution
  - Earlier in-socket saturation
- **Train application programmers to get the fallen fruits themselves (zeroth order)!**

And then, if there's time, think about the third order:

- **Power capping**
- **Efficient power distribution**
- **Automatic, profile-guided DVFS**
- **…**