Course Review

Review of "Kernel Lectures"

- **Dense direct solvers**
- **Sparse direct solvers**
- **Spectral methods**
- **N-body methods**
- **Structured grids / iterative solvers**
- **Unstructured grids / iterative solvers**
- **Monte Carlo ("MapReduce")**
- **Combinatorial logic**
- **Graph traversal**
- **Graphical models**
- **Finite state machines**
- **Dynamic programming**
- **Backtrack and branch-and-bound**

* *The Landscape of Parallel Computing Research: The View from Berkeley*, UCB/EECS-2006-183

“Kernel Lecture” Topics

- Strong and weak scaling
- Speedup
- Efficiency
- The lack of unbounded strong scalability of almost all methods, due to Amdahl’s Law
- The lack of weak scalability of explicit methods, due to the CFL stability restriction
- Optimal complexity

Kernel topics, cont.

- Different kernels stress different parts of modern architecture
  - Computational capability (flops/second)
  - Memory bandwidth (bytes/second)
  - Memory latency (seconds)
  - Network bandwidth (bytes/second)
  - Network latency (seconds)

Colloquial definitions of scalability

- "Strong scaling"
  - execution time (*T*) decreases in inverse proportion to the number of processor-memory units (*P*)
  - *fixed size problem* (*N*) *overall*
  - often instead graphed as reciprocal, "speedup"

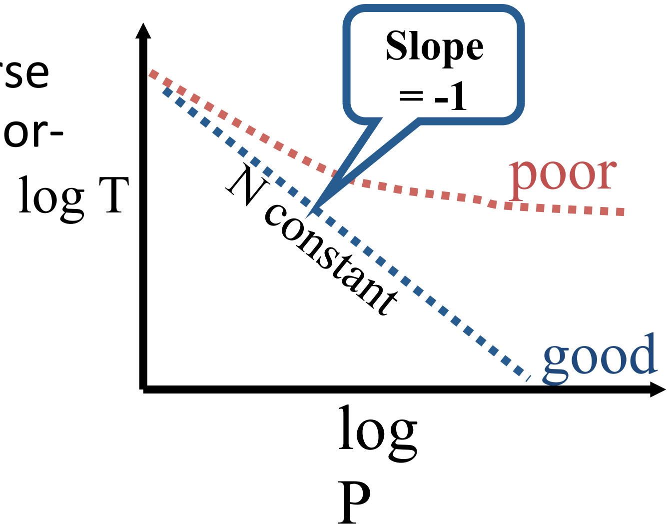

- "Weak scaling"
  - execution time remains constant, as problem size and processor number are increased in proportion
  - *fixed size problem per processor*
  - also known as "Gustafson scaling"

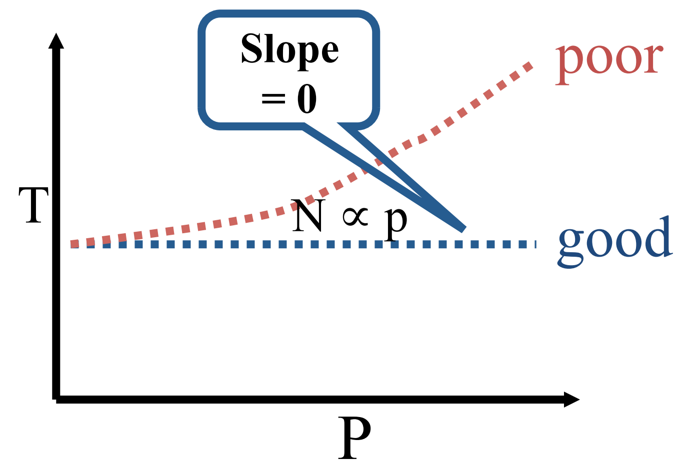

Four potential limiters on scalability

- Insufficient localized concurrency
- Load imbalance at synchronization points
- Interprocessor message latency
- Interprocessor message bandwidth

"horizontal aspects"

Four potential limiters on arithmetic performance

- Memory latency
  - failure to predict which data items are needed
- Memory bandwidth
  - failure to deliver data at consumption rate of processor
- Load/store instruction issue rate
  - failure of processor to issue enough loads/stores per cycle
- Floating point instruction issue rate
  - low percentage of floating point operations among all operations

"vertical aspects"

CS&E topics

- Well-posedness of a simulation
  - A *well-conditioned* problem, solved with a *stable* algorithm, can yield an *accurate* result
- Implications of the triple-finiteness of the computer
  - wordlength, number of words, number of flops
- Modeling hierarchy and triangle inequality

$$|e_{tot}| \le |e_{round}| + |e_{trunc}| + |e_{data}|$$

- Problem transformation and the quest for the reduced basis
  - *"Think. Then discretize."* (V. Rokhlin)

CS&E topics

- The ***condition number*** is a scalar -- the multiplier by which, in the worst case, an error in the input can be multiplied to produce an error in the output
  - The input and output errors are measured in norms, which are also scalars
- For a linear system, $Ax=b$, the condition number can be expressed as a product of two operator norms: the norm of A and of its inverse
- For a symmetric linear system, specializing further, the norm of A is its largest eigenvalue of A and the norm of A^{-1} is the largest eigenvalue of A^{-1}, which is the smallest eigenvalue of A
- For the Laplacian on a uniform mesh, these eigenvalues can be calculated exactly

Nonlinear solver lectures

- Review of Newton's method for scalar nonlinear algebraic problems (classical "rootfinding")
  - Examples of quadratic convergence
  - Intuition about non-robustness
  - Globalization through linesearch and trust region
- Introduction of Newton's method for multivariate problems
  - Jacobian matrix inversion replaces the divided by f'(x)
- Nonlinear eq. solution reduces, in principle, to
  - Residual evaluation, Jacobian evaluation, linear equation solution, state-vector update
  - In practice, all aspects are modified