Dagstuhl Seminar on High-level parallel programming models April 1999

The Barnes-Hut Algorithm as a case study for parallel programming models

Jan F. Prins, UNC Chapel Hill

Martin Simons, DaimlerChrysler R & D

Motivation

Why a case study?

- An opportunity to evaluate programming models on
  - Expressiveness
  - Performance
  - Practicality

Why consider the Barnes-Hut algorithm?

- presents non-trivial expression and performance challenges
- has been used as a case study by others
- relatively small and manageable

The *n*-body simulation problem

Simulate the evolution of a system of n bodies over time.

- Equations of motion
  - pairwise interaction potential $f(i,j)$ between body *i* and *j*
  - total force $f(i)$ on body *i*
- Numerical integration of equations of motion

Ex: Gravitation

$$f(i,j) = G \cdot \frac{m_i \cdot m_j}{r_{ij}^2}$$

$$f(i) = \sum_{j \neq i} f(i,j)$$

the basic simulation algorithm:

```
while (t < t_Final) do
  forall 1 ≤ i ≤ n do
    ⟨ compute  force f(i) on body i ⟩
    ⟨ numerical integration step  Δt ⟩
    ⟨ update velocity and position of body i ⟩
  t = t + Δt
```

Applications

- astrophysics (gravity)
- molecular dynamics (electrostatics)
- computer graphics (radiosity)

$O(n^2)$ interactions per time-step

Reducing the number of interactions

exploit combined effect of “distant” bodies:

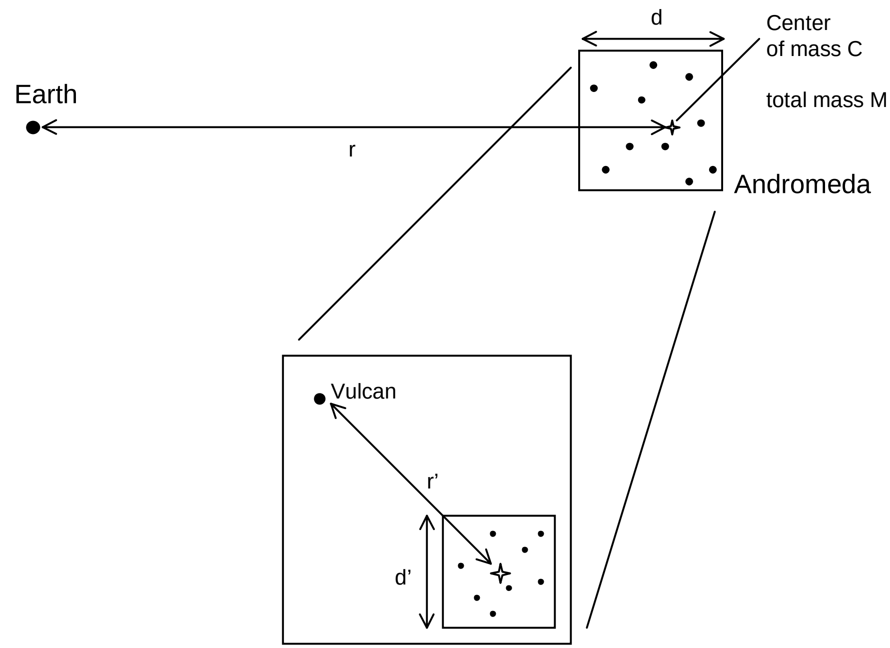

apply this idea *recursively*:

- determines control-structure
- determines hierarchical decomposition of space

Formally

Multipole expansion of the interaction of a body with the entire *Andromeda* galaxy

$$f(i) = G\frac{Mm_i}{r_{iC}^2} + \cdots$$

Multipole expansion saves work if it can be reused with multiple bodies

Accuracy of truncated multipole expansion improves with

- increasing *r*
- decreasing *d*
- retained terms in expansion
- uniformity of particle distribution

Hierarchical decomposition of space: quad- and octtrees

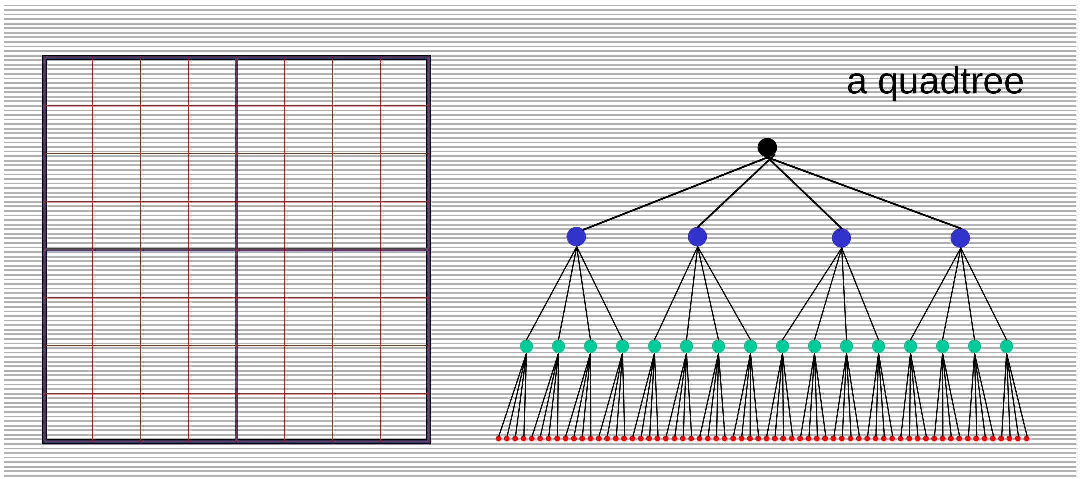

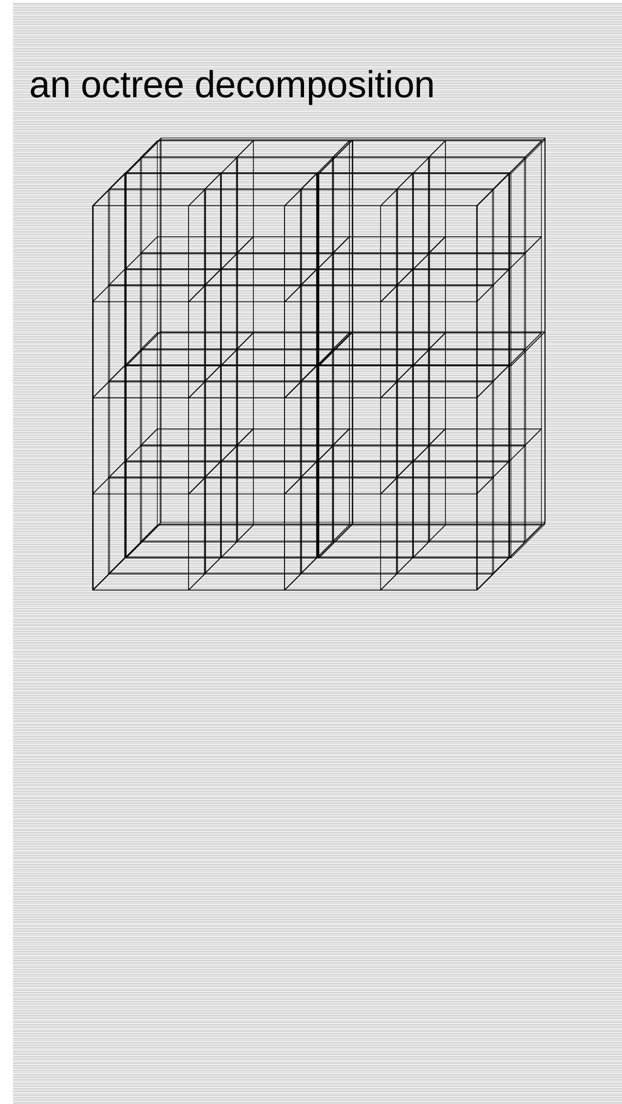

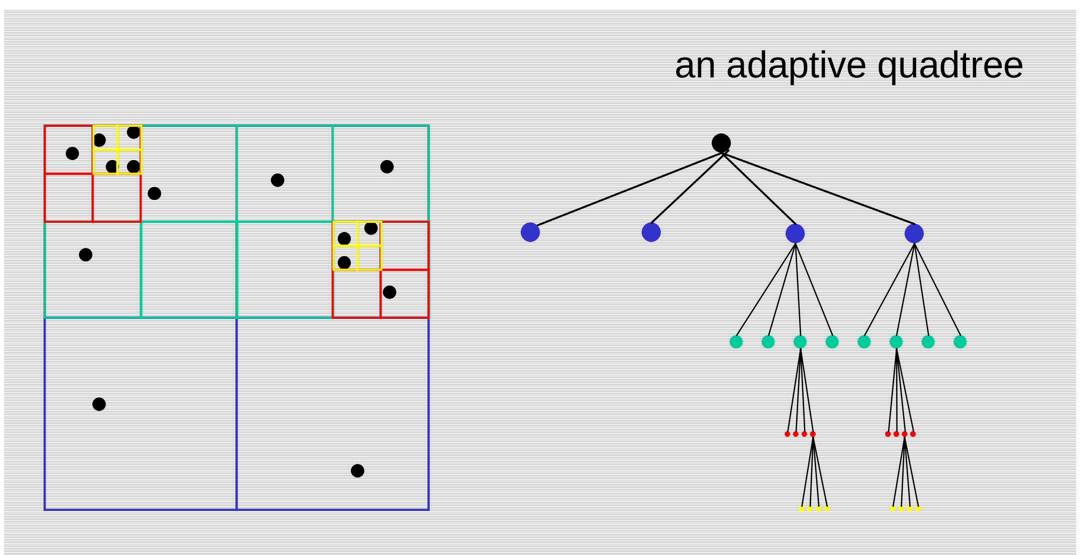

The Barnes-Hut algorithm

```
stepSystem():

  T := makeTree(P(1:n))
  forall 1 ≤ i ≤ n do
     f(i) := gravCalc(P(i),T)
     ⟨update velocities and positions⟩
```

```
function gravCalc(p,q)
  if (“q is a leaf”) then
     ⟨return body-body interaction⟩
  else
    if (“p is distant enough from q”) then
      ⟨return body-cell interaction⟩
    else
      forall q’∈ nonemptyChildren(q) do
         accumulate gravCalc(p,q’)
      ⟨return accumulated interaction⟩
    end if
 end if
```

interaction in the case of gravitation:

$$F = G \cdot \frac{m_p \cdot m_q}{r_{pq}^2} \cdot \left[\frac{x_p - x_q}{r_{pq}}, \quad \frac{y_p - y_q}{r_{pq}}, \quad \frac{z_p - z_q}{r_{pq}} \right]$$

$$r = \sqrt{(x_p - x_q)^2 + (y_p - y_q)^2 + (z_p - z_q)^2}$$

***body-body interaction**: use masses of particles and distance between them.*

***body-cell interaction**: use mass of particle and cell and distance between particle and center of mass of cell.*

force is additive; individual contributions can be accumulated.

The Barnes-Hut algorithm - Performance issues

```
stepSystem(P(1:n))

  T := makeTree(P(1:n))
  forall 1 ≤ i ≤ n do
     f(i) := gravCalc(P(i),T)
     ⟨update velocities and positions⟩
```

```
function gravCalc(p,q)
  if (“q is a leaf”) then
    ⟨return body-body interaction⟩
  else
    if (“p is distant enough from q”) then
      ⟨return body-cell interaction⟩
    else
      forall q’∈ nonemptyChildren(q) do
         accumulate gravCalc(p,q’)
      ⟨return accumulated interaction⟩
    end if
 end if
```

Parallelism

nested parallelism

- over bodies
- over recursively divided cells

load balance

different number of interactions for different bodies

Locality

nearby bodies interact with similar set of nodes in *T*

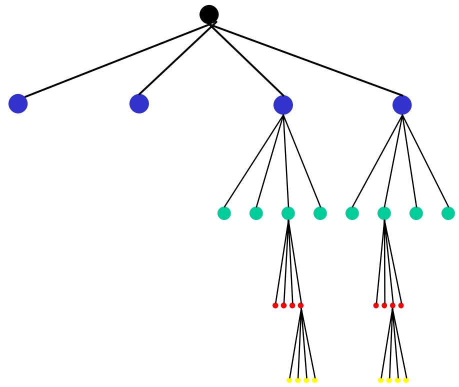

The acceptance criterion

when is a cell “distant enough”?

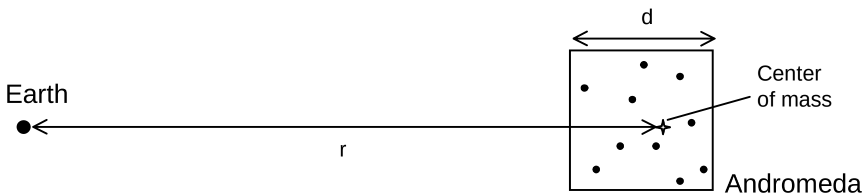

original criterion used by Barnes-Hut:

$$\frac{d}{r} < \theta \;\equiv\; r > \frac{d}{\theta}$$

where usually $0.7 \le \theta \le 1.0$

problem: detonating galaxy anomaly

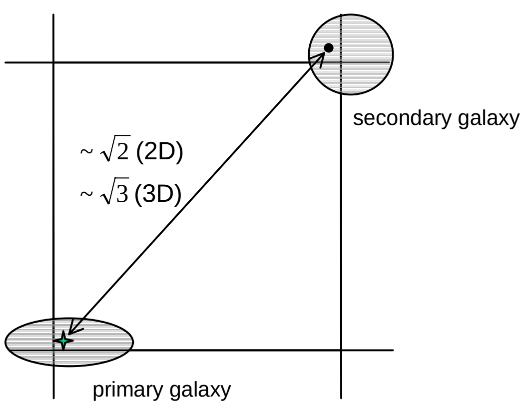

(one) solution: *add distance between center of mass (cm) and geometric center of cell (c).*

$$r > \frac{d}{\theta} + |cm - c|$$

Effects of acceptance criterion … on runtime

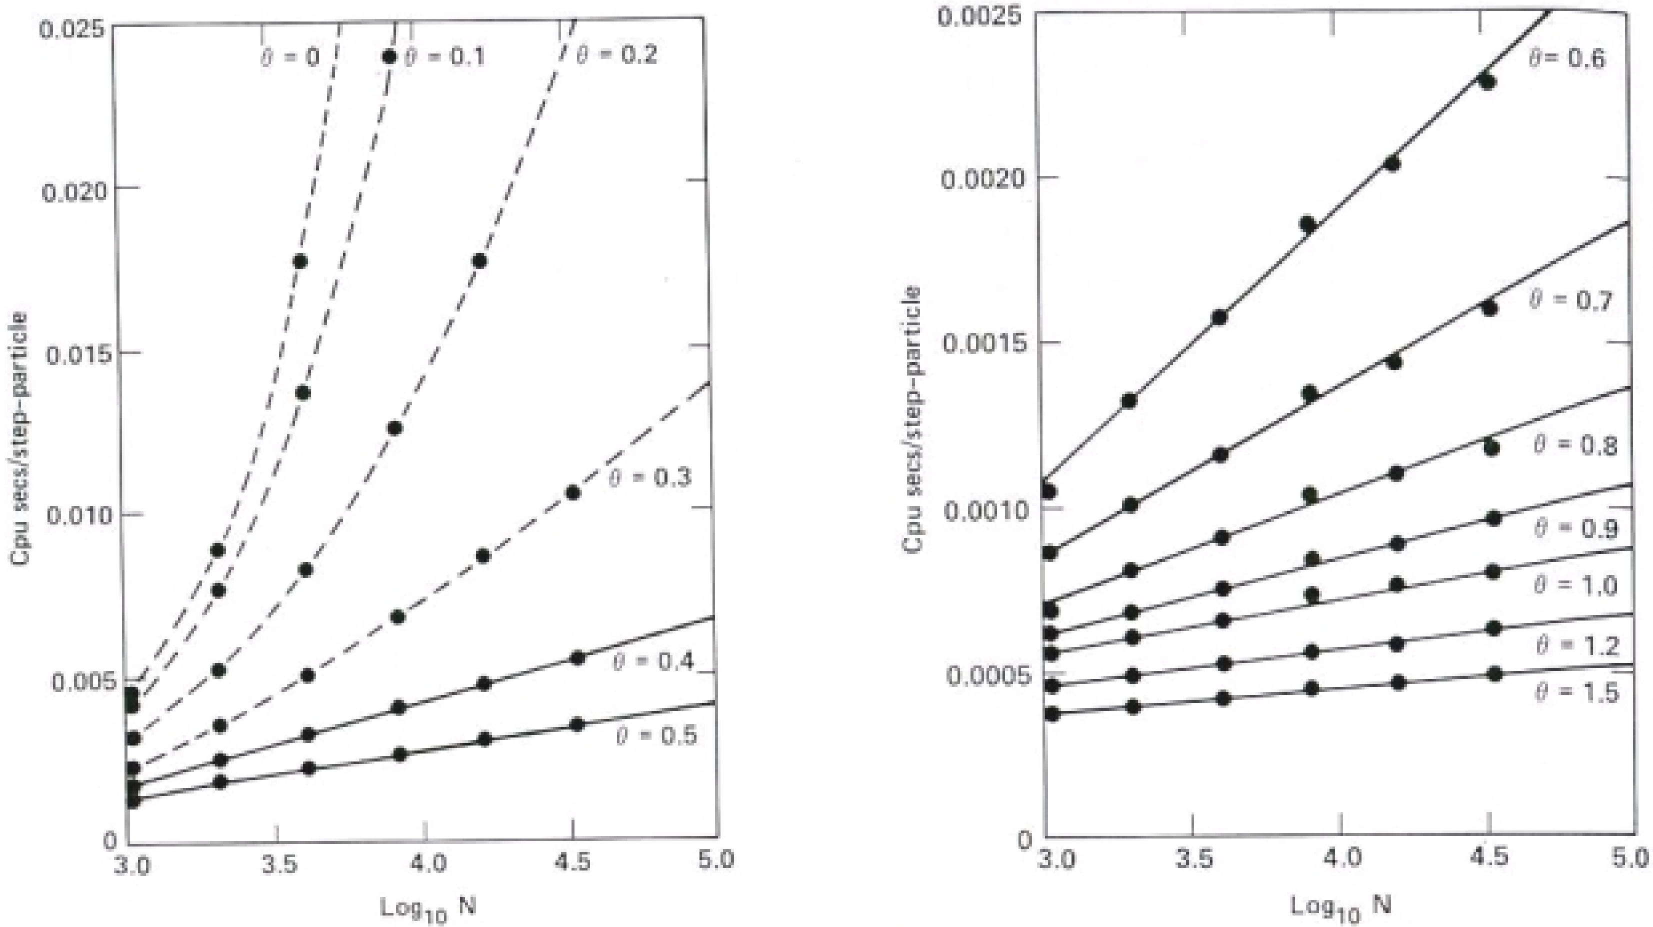

FIG. 3.—Scaling of CRAY X-MP CPU time (CPU seconds per step per particle) for spherical, isotropic Plummer models, as a function of the number of particles, for values of the clumping parameter θ in the range $0 \leq \theta \leq 1.5$. Only monopole terms have been included in the force computation.

Source: L. Hernquist. *Performance characteristics of tree codes*. Astrophysical Journal Supplement Series, Vol. 64, Pages 715-734, 1987.

Effects of acceptance criterion … on accuracy

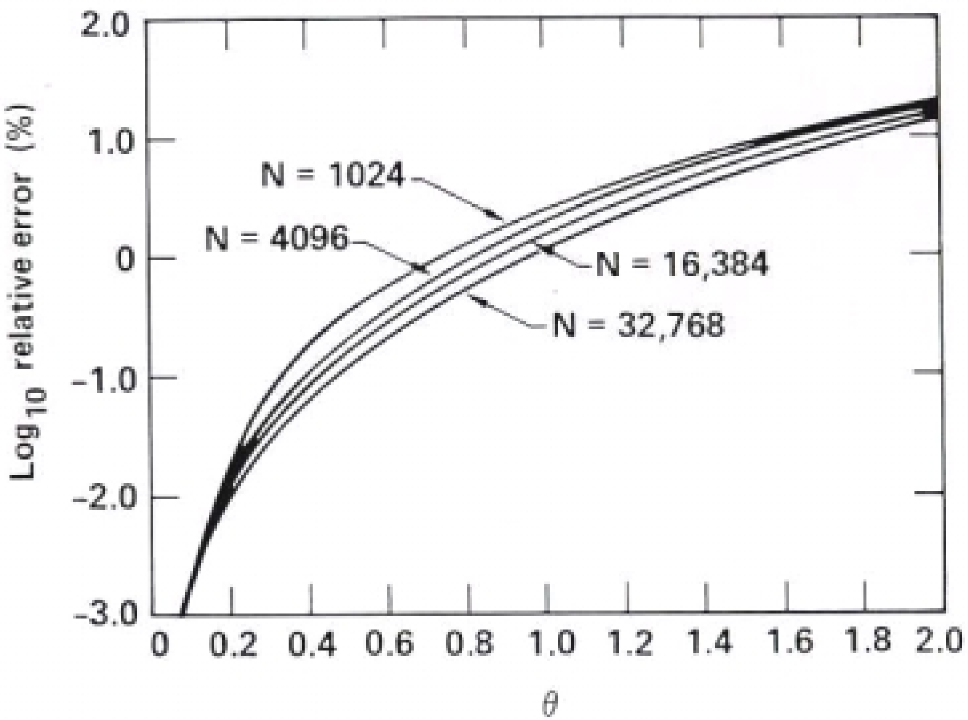

FIG. 6.—Magnitude of the typical error (in percent) in the tree force computation, relative to a direct sum, as a function of θ, for selected values of the particle number N. The calculations have assumed spherical, isotropic Plummer models with softening parameter $\epsilon = 0$, and only monopole terms have been included in the force computations.

Source: L. Hernquist. *Performance characteristics of tree codes*. Astrophysical Journal Supplement Series, Vol. 64, Pages 715-734, 1987.

1% accuracy sufficient for most astrophysical simulations. Different techniques with better error control necessary for other systems (*fast multipole methods*).

Effect of body distribution … on tree depth

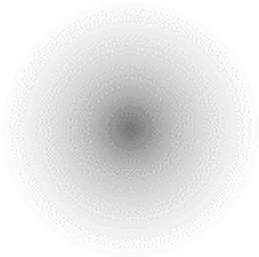

Plummer model

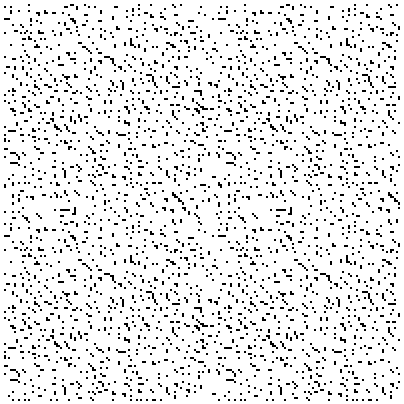

Uniform distribution

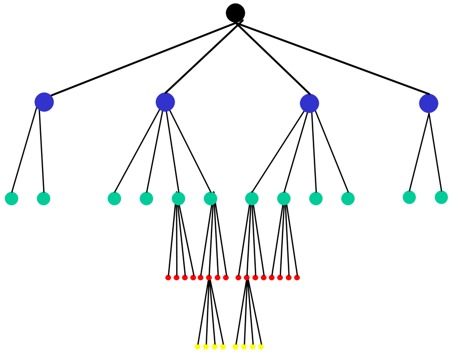

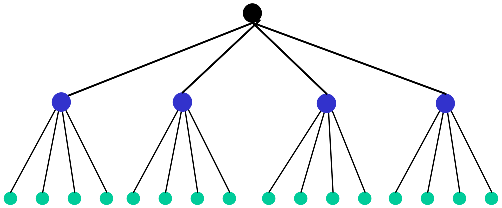

Parameters of the Barnes-Hut algorithm

- Initial configuration of bodies
  - number of bodies *n*
  - dimensionality of the simulation
  - distribution of bodies
    - uniform density (random placement)
    - plummer model
- Acceptance criterion
  - when can a body-cell interaction be used?
- Tree properties
  - number of bodies per leaf
  - update strategy
- Retained terms in mutipole expansion
  - monopole, quadrupole, ...
- Integration method and time step
- Form of potential
  - gravity, electrostatics, ...

<u>Suggested parameter settings</u>

- Initial configuration
  - parameter n
  - 3D
  - Plummer 1 (Barnes)
- Acceptance criterion
  - (Barnes)
- Tree properties
  - single body per leaf
  - strict update
- Retained terms
  - monopole, quadrupole
- Integration method
  - (Barnes)
- Form of potential
  - Gravity (Barnes)

Summary of issues for BH

How can it be expressed?

- sophisticated algorithm with
  - complex data and control structure
  - sources of parallelism and locality

What performance can be obtained and how?

- cost model
  - accuracy and tractability
- accessibility of performance factors
  - memory hierarchy
  - communication
  - load-balance
- performance portability
  - architectural requirements and commitments

How practical is the approach?

Workshop presenters who will include some discussion of BH

- Manuel Chakravarty
- Ralf Ebner
- Kevin Hammond
- Gabi Keller
- Paul Kelly
- Christoph Kessler
- Jan Prins

Website for details

http://www.score.is.tsukuba.ac.jp/~chak/dagstuhl/bh