



Autotuning (2.5/2): TCE & Empirical compilers

Prof. Richard Vuduc

Georgia Institute of Technology

CSE/CS 8803 PNA: Parallel Numerical Algorithms

[L.19] Tuesday, March 11, 2008




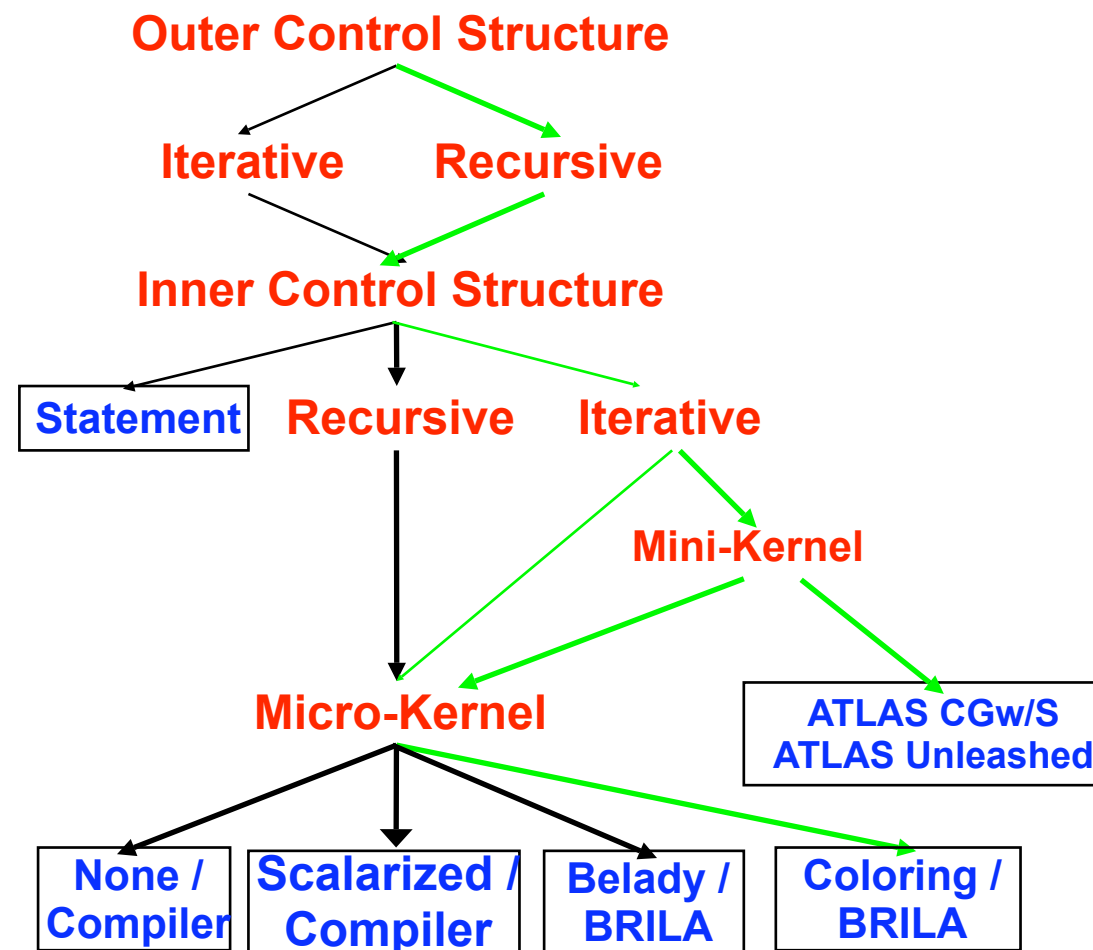
Today's sources

- CS 267 at UCB (Demmel & Yelick)
- Papers from various autotuning projects
 - PHiPAC, ATLAS, FFTW, SPIRAL, TCE
- See: Proc. IEEE 2005 special issue on Program Generation, Optimization, and Platform Adaptation
- Me (for once!)



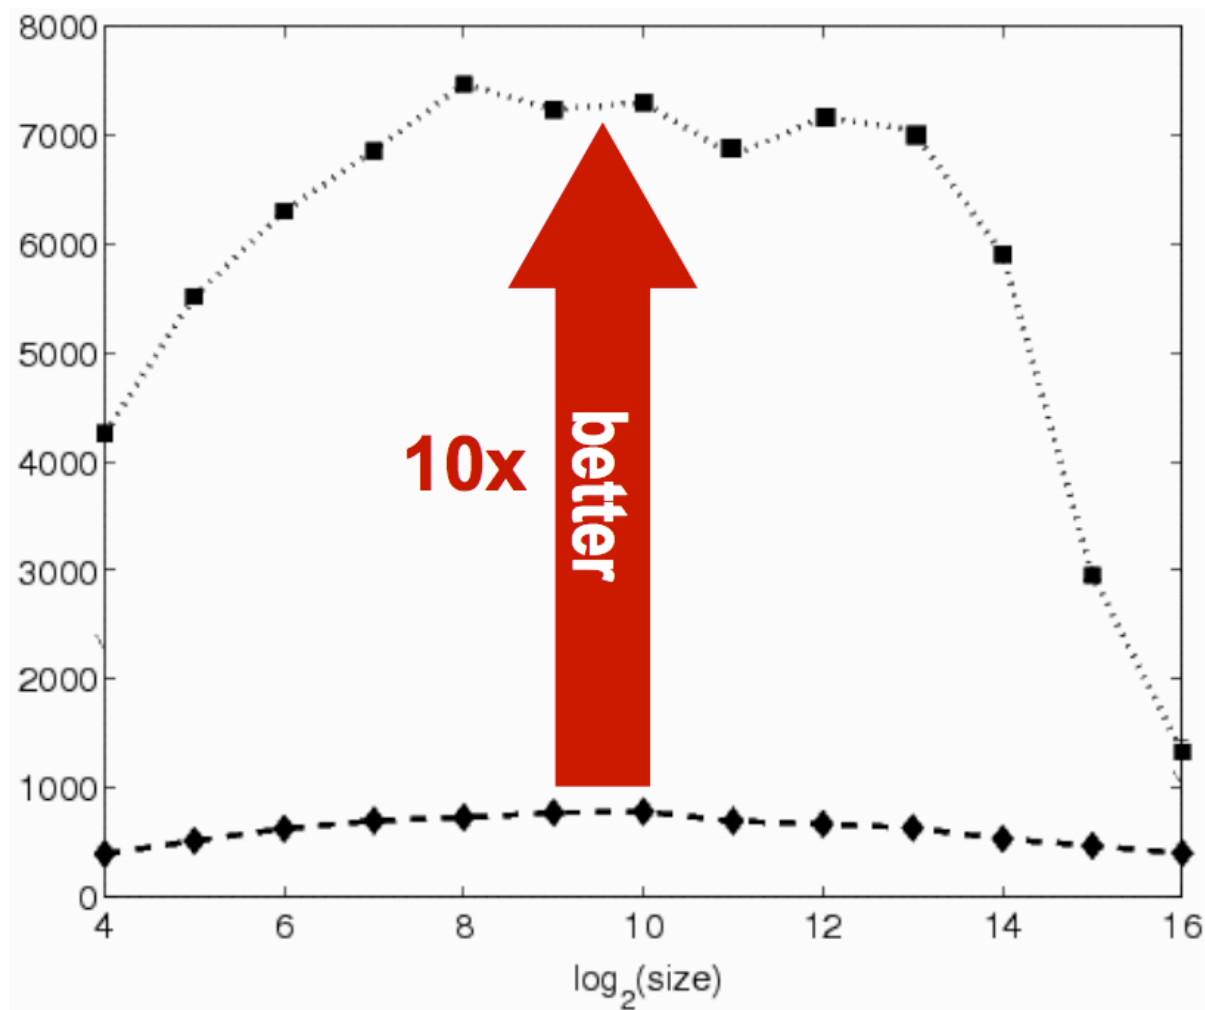
Review: Autotuners

Performance-engineering challenges



Motivation for performance tuning

pseudo
Mflop/s



**best available
implementation
(FFTW, Intel IPP, Spiral)**

**roughly the same
operations count**

**reasonable
implementation
(Numerical recipes,
GNU scientific library)**


Source: J. Johnson (2007), CScADS autotuning workshop



Context for autotuning

- Problem: HPC needs detailed low-level machine knowledge
- Autotuning methodology
 - Identify and generate a space of implementations
 - Search (modeling, experiments) to choose the best one
- Early idea seedlings
 - Polyalgorithms
 - Profile and feedback-directed compilation
 - Domain- and architecture-specific code generators

Example: What a search space looks like




Platform: Sun Ultra Ili

16 double regs

667 Mflop/s peak

Unrolled, pipelined inner-kernel

Sun cc v5.0 compiler



Source: PHiPAC Project at UC Berkeley (1997)

Cooley-Tukey FFT algorithm: Encoding in FFTW's codelet generator

$$y[k] \leftarrow \text{DFT}_N(x, k) \equiv \sum_{j=0}^{N-1} x[j] \cdot \omega_N^{-kj} \quad x, y \in \mathbb{C}^N$$

$$y[k_1 + k_2 \cdot N_1] \leftarrow \sum_{n_2=0}^{N_2-1} \left[\underbrace{\left(\sum_{n_1=0}^{N_1-1} x[n_1 \cdot N_2 + n_2] \cdot \omega_{N_1}^{-k_1 n_1} \right)}_{N_1\text{-point DFT}} \cdot \underbrace{\omega_N^{-k_1 n_2}}_{\text{Twiddle}} \right] \cdot \omega_{N_2}^{-k_2 n_2}$$

N_2 -point DFT

let dftgen(N, x) \equiv fun $k \rightarrow \dots$ # $\text{DFT}_N(x, k)$

let cooley_tukey(N_1, N_2, x) \equiv

let $\hat{x} \equiv$ fun $n_2, n_1 \rightarrow x(n_2 + n_1 \cdot N_2)$ in

let $\mathbf{G}_1 \equiv$ fun $n_2 \rightarrow \text{dftgen}(N_1, \hat{x}(n_2, --))$ in

let $\mathbf{W} \equiv$ fun $k_1, n_2 \rightarrow \mathbf{G}_1(n_2, k_1) \cdot \omega_N^{-k_1 n_2}$ in


let $\mathbf{G}_2 \equiv$ fun $k_1 \rightarrow \text{dftgen}(N_2, \mathbf{W}(k_1, --))$

in

fun $k \rightarrow \mathbf{G}_2(k \bmod N_1, k \text{ div } N_1)$

**(Functional
pseudo-code)**


Matrix 02-raefsky3



1792 ideal nz + 0 explicit zeros = 1792 nz



900 MHz Itanium 2, Intel C v8: ref=275 Mflop/s





Tensor Contraction Engine (TCE) for quantum chemistry



Tensor Contraction Engine (TCE)

- Application domain: Quantum chemistry
 - Electronic structure calculations
 - Dominant computation expressible as a “tensor contraction”
- TCE generates a complete parallel program from a high-level spec
 - Automates time-space trade-offs
 - Output
- S. Hirata (2002), and many others
- Following presentation taken from Proc. IEEE 2005 special issue

Motivation: Simplify program development

$$\begin{aligned} \text{hbar}[a,b,i,j] = & \text{sum}[f[b,c] * t[i,j,a,c], c] - \text{sum}[f[k,c] * t[k,b] * t[i,j,a,c], k,c] + \text{sum}[f[a,c] * t[i,j,c,b], c] - \text{sum}[f[k,c] * t[k,a] * t[i,j,c,b], k,c] - \text{sum}[f[k,j] * t[i,k,a,b], k] - \text{sum}[f[k,c] * \\ & t[j,c] * t[i,k,a,b], k,c] - \text{sum}[f[k,i] * t[j,k,b,a], k] - \text{sum}[f[k,c] * t[i,c] * t[j,k,b,a], k,c] + \text{sum}[t[i,c] * t[j,d] * v[a,b,c,d], c,d] + \text{sum}[t[i,j,c,d] * v[a,b,c,d], c,d] + \text{sum}[t[j,c] * v[a,b,i,c], \\ & c] - \text{sum}[t[k,b] * v[a,k,i,j], k] + \text{sum}[t[i,c] * v[b,a,j,c], c] - \text{sum}[t[k,a] * v[b,k,j,i], k] - \text{sum}[t[k,d] * t[i,j,c,b] * v[k,a,c,d], k,c,d] - \text{sum}[t[i,c] * t[j,k,b,d] * v[k,a,c,d], k,c,d] - \text{sum}[t[j,c] \\ & * t[k,b] * v[k,a,c,i], k,c] + 2 * \text{sum}[t[j,k,b,c] * v[k,a,c,i], k,c] - \text{sum}[t[j,k,c,b] * v[k,a,c,i], k,c] - \text{sum}[t[i,c] * t[j,d] * t[k,b] * v[k,a,d,c], k,c,d] + 2 * \text{sum}[t[k,d] * t[i,j,c,b] * v[k,a,d,c], \\ & k,c,d] - \text{sum}[t[k,b] * t[i,j,c,d] * v[k,a,d,c], k,c,d] - \text{sum}[t[j,d] * t[i,k,c,b] * v[k,a,d,c], k,c,d] + 2 * \text{sum}[t[i,c] * t[j,k,b,d] * v[k,a,d,c], k,c,d] - \text{sum}[t[i,c] * t[j,k,d,b] * v[k,a,d,c], k,c,d] - \\ & \text{sum}[t[j,k,b,c] * v[k,a,i,c], k,c] - \text{sum}[t[i,c] * t[k,b] * v[k,a,j,c], k,c] - \text{sum}[t[i,k,c,b] * v[k,a,j,c], k,c] - \text{sum}[t[i,c] * t[j,d] * t[k,a] * v[k,b,c,d], k,c,d] - \text{sum}[t[k,d] * t[i,j,a,c] * v[k,b,c,d], \\ & k,c,d] - \text{sum}[t[k,a] * t[i,j,c,d] * v[k,b,c,d], k,c,d] + 2 * \text{sum}[t[j,d] * t[i,k,a,c] * v[k,b,c,d], k,c,d] - \text{sum}[t[j,d] * t[i,k,c,a] * v[k,b,c,d], k,c,d] - \text{sum}[t[i,c] * t[j,k,d,a] * v[k,b,c,d], k,c,d] \\ & - \text{sum}[t[i,c] * t[k,a] * v[k,b,c,j], k,c] + 2 * \text{sum}[t[i,k,a,c] * v[k,b,c,j], k,c] - \text{sum}[t[i,k,c,a] * v[k,b,c,j], k,c] + 2 * \text{sum}[t[k,d] * t[i,j,a,c] * v[k,b,d,c], k,c,d] - \text{sum}[t[j,d] * t[i,k,a,c] * \\ & v[k,b,d,c], k,c,d] - \text{sum}[t[j,c] * t[k,a] * v[k,b,i,c], k,c] - \text{sum}[t[j,k,c,a] * v[k,b,i,c], k,c] - \text{sum}[t[i,k,a,c] * v[k,b,j,c], k,c] + \text{sum}[t[i,c] * t[j,d] * t[k,a] * t[l,b] * v[k,l,c,d], k,l,c,d] - 2 * \\ & \text{sum}[t[k,b] * t[l,d] * t[i,j,a,c] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[k,a] * t[l,d] * t[i,j,c,b] * v[k,l,c,d], k,l,c,d] + \text{sum}[t[k,a] * t[l,b] * t[i,j,c,d] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[j,c] * t[l,d] * \\ & t[i,k,a,b] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[j,d] * t[l,b] * t[i,k,a,c] * v[k,l,c,d], k,l,c,d] + \text{sum}[t[j,d] * t[l,b] * t[i,k,c,a] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[i,c] * t[l,d] * t[j,k,b,a] * v[k,l,c,d], \\ & k,l,c,d] + \text{sum}[t[i,c] * t[l,a] * t[j,k,b,d] * v[k,l,c,d], k,l,c,d] + \text{sum}[t[i,c] * t[l,b] * t[j,k,d,a] * v[k,l,c,d], k,l,c,d] + \text{sum}[t[i,k,c,d] * t[j,l,b,a] * v[k,l,c,d], k,l,c,d] + 4 * \text{sum}[t[i,k,a,c] * \\ & t[j,l,b,d] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[i,k,c,a] * t[j,l,b,d] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[i,k,a,b] * t[j,l,c,d] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[i,k,a,c] * t[j,l,d,b] * v[k,l,c,d], k,l,c,d] \\ & + \text{sum}[t[i,k,c,a] * t[j,l,d,b] * v[k,l,c,d], k,l,c,d] + \text{sum}[t[i,c] * t[j,d] * t[k,l,a,b] * v[k,l,c,d], k,l,c,d] + \text{sum}[t[i,j,c,d] * t[k,l,a,b] * v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[i,j,c,b] * t[k,l,a,d] * \\ & v[k,l,c,d], k,l,c,d] - 2 * \text{sum}[t[i,j,a,c] * t[k,l,b,d] * v[k,l,c,d], k,l,c,d] + \text{sum}[t[j,c] * t[k,b] * t[l,a] * v[k,l,c,i], k,l,c] + \text{sum}[t[l,c] * t[j,k,b,a] * v[k,l,c,i], k,l,c] - 2 * \text{sum}[t[l,a] * t[j,k,b,c] \\ & * v[k,l,c,i], k,l,c] + \text{sum}[t[l,a] * t[j,k,c,b] * v[k,l,c,i], k,l,c] - 2 * \text{sum}[t[k,c] * t[j,l,b,a] * v[k,l,c,i], k,l,c] + \text{sum}[t[k,a] * t[j,l,b,c] * v[k,l,c,i], k,l,c] + \text{sum}[t[k,b] * t[j,l,c,a] * v[k,l,c,i], \\ & k,l,c] + \text{sum}[t[j,c] * t[l,k,a,b] * v[k,l,c,i], k,l,c] + \text{sum}[t[i,c] * t[k,a] * t[l,b] * v[k,l,c,j], k,l,c] + \text{sum}[t[l,c] * t[i,k,a,b] * v[k,l,c,j], k,l,c] - 2 * \text{sum}[t[l,b] * t[i,k,a,c] * v[k,l,c,j], k,l,c] \\ & + \text{sum}[t[l,b] * t[i,k,c,a] * v[k,l,c,j], k,l,c] + \text{sum}[t[i,c] * t[k,l,a,b] * v[k,l,c,j], k,l,c] + \text{sum}[t[j,c] * t[l,d] * t[i,k,a,b] * v[k,l,d,c], k,l,c,d] + \text{sum}[t[j,d] * t[l,b] * t[i,k,a,c] * v[k,l,d,c], \\ & k,l,c,d] + \text{sum}[t[j,d] * t[l,a] * t[i,k,c,b] * v[k,l,d,c], k,l,c,d] - 2 * \text{sum}[t[i,k,c,d] * t[j,l,b,a] * v[k,l,d,c], k,l,c,d] - 2 * \text{sum}[t[i,k,a,c] * t[j,l,b,d] * v[k,l,d,c], k,l,c,d] + \text{sum}[t[i,k,c,a] * \\ & t[j,l,b,d] * v[k,l,d,c], k,l,c,d] + \text{sum}[t[i,k,a,b] * t[j,l,c,d] * v[k,l,d,c], k,l,c,d] + \text{sum}[t[i,k,c,b] * t[j,l,d,a] * v[k,l,d,c], k,l,c,d] + \text{sum}[t[i,k,a,c] * t[j,l,d,b] * v[k,l,d,c], k,l,c,d] + \text{sum}[t[k,a] \\ & * t[l,b] * v[k,l,i,j], k,l] + \text{sum}[t[k,l,a,b] * v[k,l,i,j], k,l] + \text{sum}[t[k,b] * t[l,d] * t[i,j,a,c] * v[l,k,c,d], k,l,c,d] + \text{sum}[t[k,a] * t[l,d] * t[i,j,c,b] * v[l,k,c,d], k,l,c,d] + \text{sum}[t[i,c] * t[l,d] * \\ & t[j,k,b,a] * v[l,k,c,d], k,l,c,d] - 2 * \text{sum}[t[i,c] * t[l,a] * t[j,k,b,d] * v[l,k,c,d], k,l,c,d] + \text{sum}[t[i,c] * t[l,a] * t[j,k,d,b] * v[l,k,c,d], k,l,c,d] + \text{sum}[t[i,j,c,b] * t[k,l,a,d] * v[l,k,c,d], k,l,c,d] \\ & + \text{sum}[t[i,j,a,c] * t[k,l,b,d] * v[l,k,c,d], k,l,c,d] - 2 * \text{sum}[t[l,c] * t[i,k,a,b] * v[l,k,c,j], k,l,c] + \text{sum}[t[l,b] * t[i,k,a,c] * v[l,k,c,j], k,l,c] + \text{sum}[t[l,a] * t[i,k,c,b] * v[l,k,c,j], k,l,c] + \\ & v[a,b,i,j] \end{aligned}$$

Source: Baumgartner, et al. (2005)

Rewriting to reduce operation counts

Naïvely, $\approx 4 \times N^{10}$ flops

$$S_{abij} = \sum_{c,d,e,f,k,l} A_{acik} \times B_{befl} \times C_{dfjk} \times D_{cdel}$$

\Downarrow

$$S_{abij} = \sum_{c,k} \left(\sum_{d,f} \left(\sum_{e,l} B_{befl} \times D_{cdel} \right) \times C_{dfjk} \right) \times A_{acik}$$

Assuming associativity and distributivity, $\approx 6 \times N^6$ flops,
but also **requires temporary storage.**

Source: Baumgartner, et al. (2005)

Operation and storage minimization *via* loop fusion

$$T_{bcdf}^{(1)} = \sum_{e,l} B_{befl} \times D_{cdel}$$

$$T_{bcjk}^{(2)} = \sum_{d,f} T_{bcdf}^{(1)} \times C_{dfjk}$$

$$S_{abij} = \sum_{c,k} T_{bcjk}^{(2)} \times A_{acik}$$

$T1 = T2 = S = 0$

for b, c, d, e, f, l **do**

$T1[b, c, d, f] += B[b, e, f, l] \cdot D[c, d, e, l]$

for b, c, d, f, j, k **do**

$T2[b, c, j, k] += T1[b, c, d, f] \cdot C[d, f, j, k]$

for a, b, c, i, j, k **do**

$S[a, b, i, j] += T2[b, c, j, k] \cdot A[a, c, i, k]$



Operation and storage minimization *via* loop fusion

$$T_{bcdf}^{(1)} = \sum_{e,l} B_{befl} \times D_{cdel}$$

$$T_{bcjk}^{(2)} = \sum_{d,f} T_{bcdf}^{(1)} \times C_{dfjk}$$

$$S_{abij} = \sum_{c,k} T_{bcjk}^{(2)} \times A_{acik}$$

$T1 = T2 = S = 0$

for b, c, d, e, f, l **do**

$T1[b, c, d, f] += B[b, e, f, l] \cdot D[c, d, e, l]$

for b, c, d, f, j, k **do**

$T2[b, c, j, k] += T1[b, c, d, f] \cdot C[d, f, j, k]$

for a, b, c, i, j, k **do**

$S[a, b, i, j] += T2[b, c, j, k] \cdot A[a, c, i, k]$

$S = 0$

for b, c **do**

$T1f \leftarrow 0, T2f \leftarrow 0$

for d, f **do**

for e, l **do**

$T1f += B[b, e, f, l] \cdot D[c, d, e, l]$

for j, k **do**

$T2f[j, k] += T1f \cdot C[d, f, j, k]$

for a, i, j, k **do**

$S[a, b, i, j] += T2f[j, k] \cdot A[a, c, i, k]$

Time-space trade-offs

Max index of
 $a-f$: $O(1000)$
 $i-k$: $O(100)$

for a, e, c, f do

for i, j do

$$X_{aecf} += T_{ijae} \cdot T_{ijcf} \quad \leftarrow \text{“Contraction” of } T \text{ over } i, j$$

for c, e, b, k do

$$T_{cebk}^{(1)} \leftarrow f_1(c, e, b, k)$$

for a, f, b, k do

$$T_{afb}^{(2)} \leftarrow f_2(a, f, b, k)$$

Integrals, $O(1000)$ flops

for c, e, a, f do

for b, k do

$$Y_{ceaf} += T_{cebk}^{(1)} \cdot T_{afb}^{(2)} \quad \leftarrow \text{“Contraction” over } T^{(1)} \text{ and } T^{(2)}$$

for c, e, a, f do

$$E += X_{aecf} \cdot Y_{ceaf}$$

Time-space trade-offs

for a, e, c, f do ←

for i, j do

$$X_{aecf} += T_{ijae} \cdot T_{ijcf}$$

for c, e, b, k do

$$T_{cebk}^{(1)} \leftarrow f_1(c, e, b, k)$$

for a, f, b, k do

$$T_{afb k}^{(2)} \leftarrow f_2(a, f, b, k)$$

for c, e, a, f do ←

for b, k do

$$Y_{ceaf} += T_{cebk}^{(1)} \cdot T_{afb k}^{(2)}$$

for c, e, a, f do ←

$$E += X_{aecf} \cdot Y_{ceaf}$$

Max index of
 $a-f$: $O(1000)$
 $i-k$: $O(100)$

Same indices
⇒ Loop fusion candidates

Time-space trade-offs

<p>for a, e, c, f do</p> <p style="padding-left: 20px;">for i, j do</p> <p style="padding-left: 40px;">$X_{aecf} += T_{ijae} \cdot T_{ijcf}$</p> <p>for c, e, b, k do</p> <p style="padding-left: 20px;">$T_{cebk}^{(1)} \leftarrow f_1(c, e, b, k)$</p> <p>for a, f, b, k do</p> <p style="padding-left: 20px;">$T_{afbk}^{(2)} \leftarrow f_2(a, f, b, k)$</p> <p>for c, e, a, f do</p> <p style="padding-left: 20px;">for b, k do</p> <p style="padding-left: 40px;">$Y_{ceaf} += T_{cebk}^{(1)} \cdot T_{afbk}^{(2)}$</p> <p>for c, e, a, f do</p> <p style="padding-left: 20px;">$E += X_{aecf} \cdot Y_{ceaf}$</p>	<p>→</p>	<p>for a, e, c, f do</p> <p style="padding-left: 20px;">for i, j do</p> <p style="padding-left: 40px;">$X_{aecf} += T_{ijae} \cdot T_{ijcf}$</p> <p>for a, c, e, f, b, k do</p> <p style="padding-left: 20px;">$T_{cebk}^{(1)} \leftarrow f_1(c, e, b, k)$</p> <p>for a, e, c, f, b, k do</p> <p style="padding-left: 20px;">$T_{afbk}^{(2)} \leftarrow f_2(a, f, b, k)$</p> <p>for c, e, a, f do</p> <p style="padding-left: 20px;">for b, k do</p> <p style="padding-left: 40px;">$Y_{ceaf} += T_{cebk}^{(1)} \cdot T_{afbk}^{(2)}$</p> <p>for c, e, a, f do</p> <p style="padding-left: 20px;">$E += X_{aecf} \cdot Y_{ceaf}$</p>
--	----------	--

**Add
extra
flops**

Time-space trade-offs

for a, e, c, f do \longrightarrow

for i, j do

$$X_{aecf} += T_{ijae} \cdot T_{ijcf}$$

for c, e, b, k do

$$T_{cebk}^{(1)} \leftarrow f_1(c, e, b, k)$$

for a, f, b, k do

$$T_{afbk}^{(2)} \leftarrow f_2(a, f, b, k)$$

for c, e, a, f do \longrightarrow

for b, k do

$$Y_{cea f} += T_{cebk}^{(1)} \cdot T_{afbk}^{(2)}$$

for c, e, a, f do \longrightarrow

$$E += X_{aecf} \cdot Y_{cea f}$$

for a, e, c, f do \Leftarrow Fused

for i, j do

$$x += T_{ijae} \cdot T_{ijcf}$$

for b, k do

$$T_{cebk}^{(1)} \leftarrow f_1(c, e, b, k)$$

$$T_{afbk}^{(2)} \leftarrow f_2(a, f, b, k)$$

$$y += T_{cebk}^{(1)} \cdot T_{afbk}^{(2)}$$

$$E += x \cdot y$$



Tiled & partially fused

for a, e, c, f **do** \longrightarrow

for i, j **do**

$$X_{aecf} += T_{ijae} \cdot T_{ijcf}$$

for c, e, b, k **do**

$$T_{cebk}^{(1)} \leftarrow f_1(c, e, b, k)$$

for a, f, b, k **do**

$$T_{afbk}^{(2)} \leftarrow f_2(a, f, b, k)$$

for c, e, a, f **do** \longrightarrow

for b, k **do**

$$Y_{ceaf} += T_{cebk}^{(1)} \cdot T_{afbk}^{(2)}$$

for c, e, a, f **do** \longrightarrow

$$E += X_{aecf} \cdot Y_{ceaf}$$

for a^B, e^B, c^B, f^B **do**

for a, e, c, f **do**

for i, j **do**

$$\hat{X}_{aecf} += T_{ijae} \cdot T_{ijcf}$$

for b, k **do**

for c, e **do**

$$\hat{T}_{ce}^{(1)} \leftarrow f_1(c, e, b, k)$$

for a, f **do**

$$\hat{T}_{af}^{(2)} \leftarrow f_2(a, f, b, k)$$

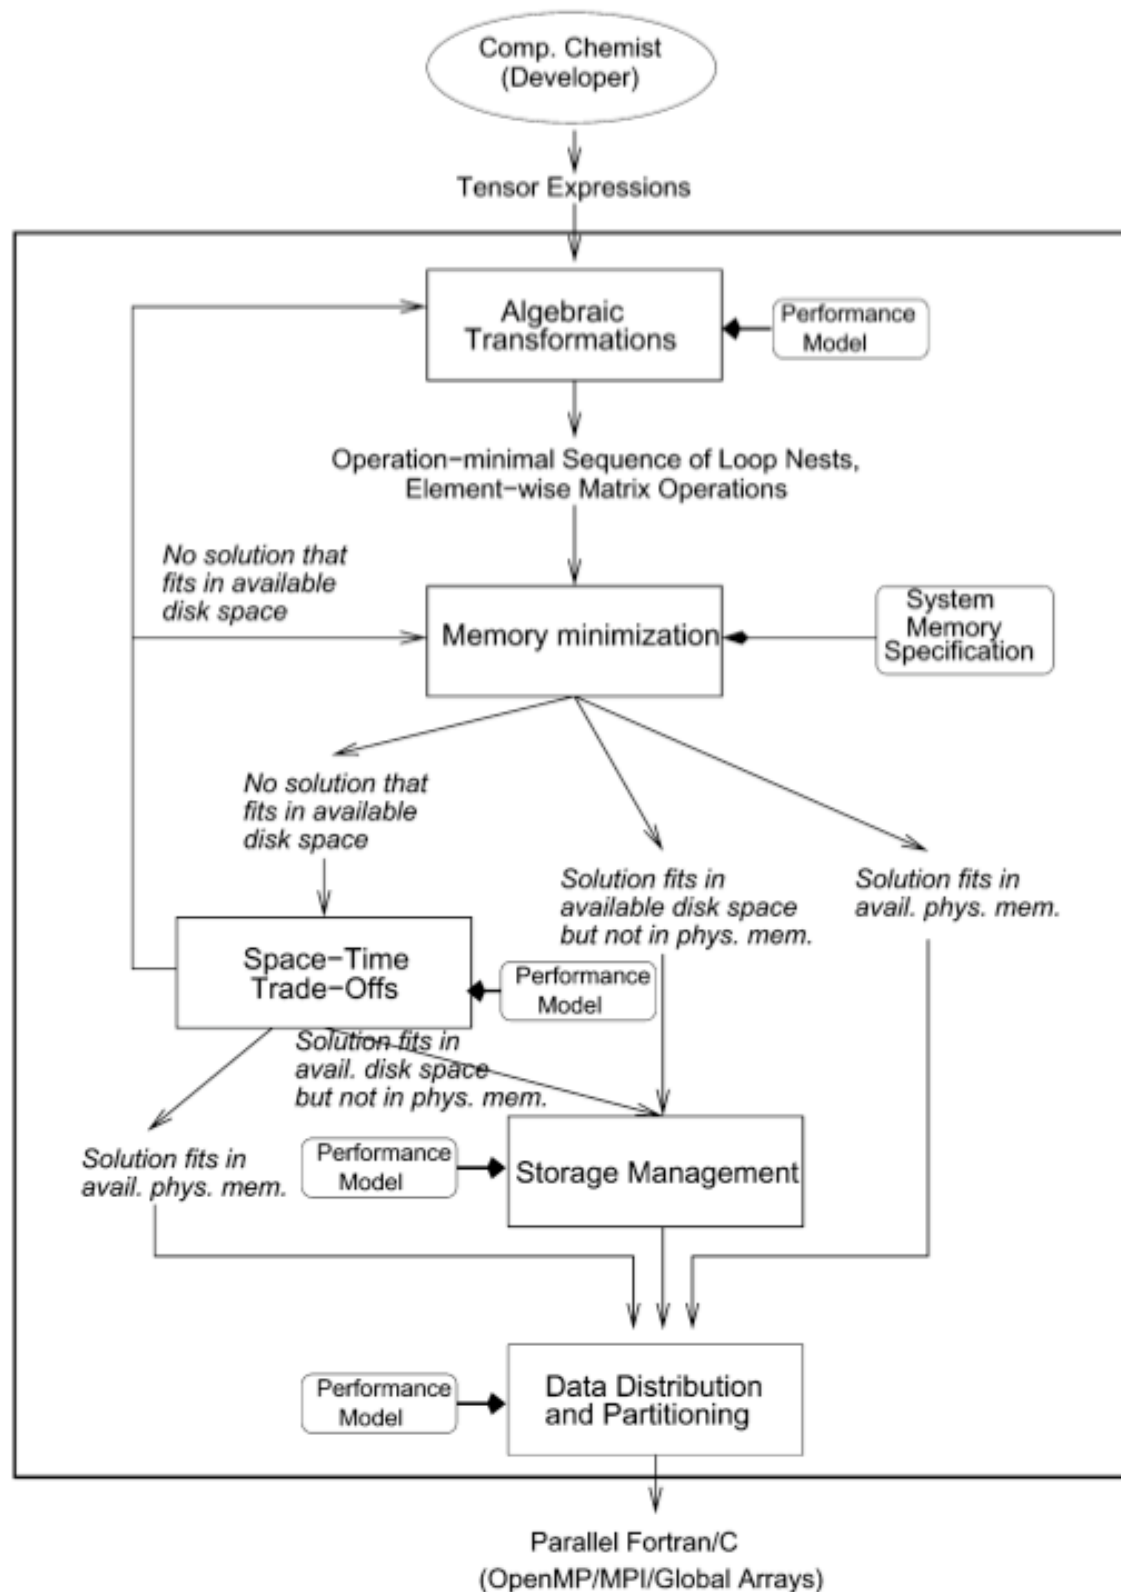
for c, e, a, f **do**

$$\hat{Y}_{ceaf} += \hat{T}_{ce}^{(1)} \cdot \hat{T}_{af}^{(2)}$$

for c, e, a, f **do**

$$E += \hat{X}_{aecf} \cdot \hat{Y}_{ceaf}$$





- **Transform algebraically**, to minimize flops
- **Minimize temporary storage**
- **Distribute and partition data** for a parallel system
- **Search wrt space-time trade-off** (feedback)
- For **out-of-core** problems, apply **optimize data locality**
- Generate final program (C/ Fortran + MPI/Global-arrays)



Tensor loop nest \Rightarrow Expression tree

for a, e, c, f do

 for i, j do

$$X_{aecf} += T_{ijae} \cdot T_{ijcf}$$

 for c, e, b, k do

$$T_{ceb}^{(1)} \leftarrow f_1(c, e, b, k)$$

 for a, f, b, k do

$$T_{afb}^{(2)} \leftarrow f_2(a, f, b, k)$$


 for c, e, a, f do

 for b, k do


$$Y_{ceaf} += T_{ceb}^{(1)} \cdot T_{afb}^{(2)}$$

 for c, e, a, f do


$$E += X_{aecf} \cdot Y_{ceaf}$$



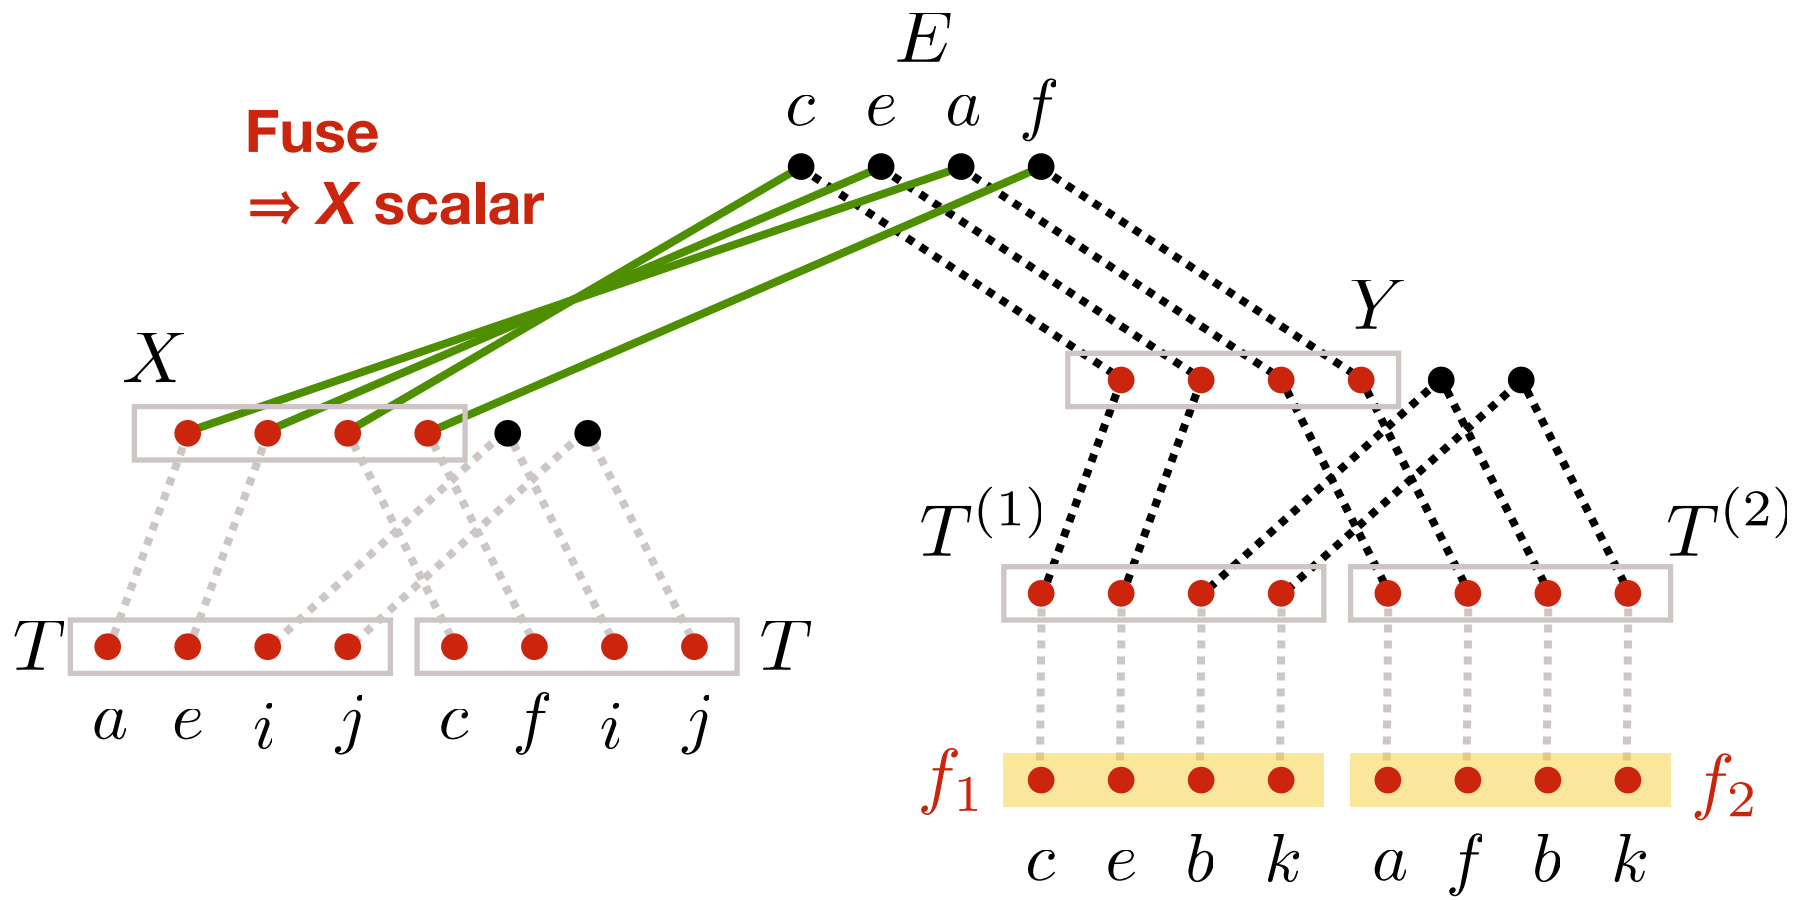
Expression tree \Rightarrow fusion graph



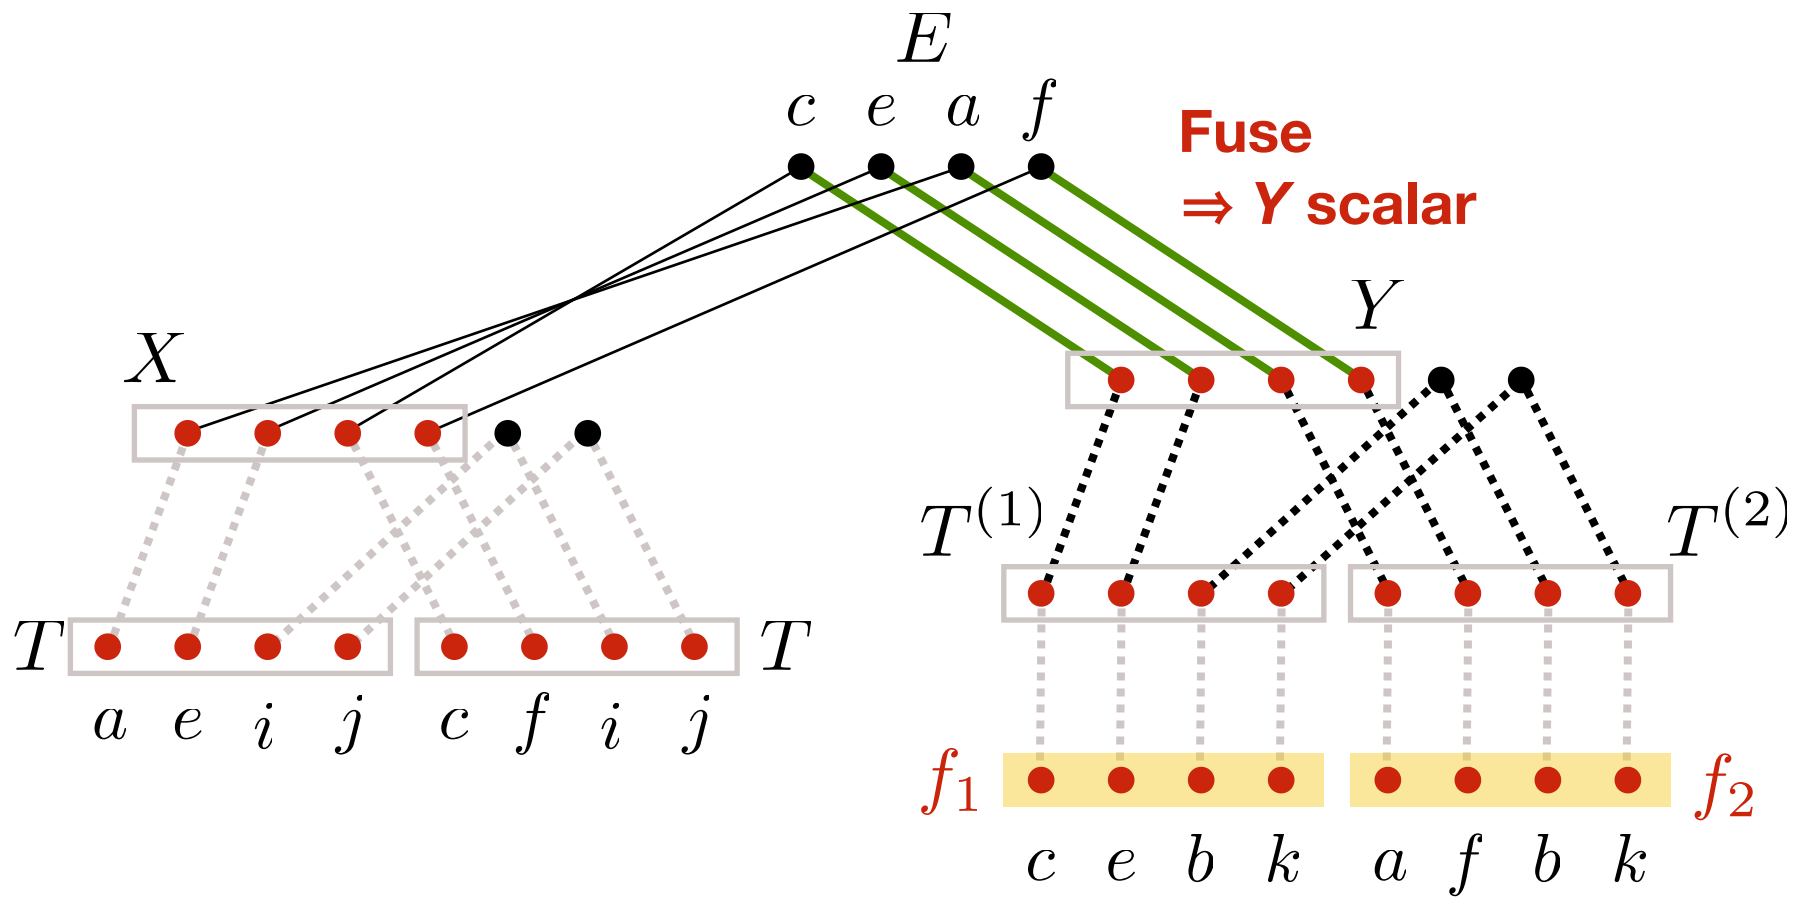
Fusion graph




Fusion graph




Fusion graph




Fusion graph




Fusion graph



Fusion graph





(b) Partially fused computation from Fig. 5.



Empirical compilers and tools



Code generation tools for autotuning

- Code generation tools
 - GNU Superoptimizer -- Exhaustive search over schedules of straight-line code
 - Denali -- Theorem proving based scheduling
 - iFKO (Whaley @ UTSA) -- Iterative floating-point kernel optimizer
 - POET (Yi @ UTSA) -- Parameterized Optimizations for Empirical Tuning



Iterative/empirical compilers

- Compile-time
 - “Iterative compilation” -- Kisuki, Knijnenberg, O’Boyle, *et al.*
 - Hybrid model/search-based compiler -- Hall, et al. (USC)
 - Eigenmann @ Purdue (Polaris)
 - Quinlan, et al. (LLNL / PERI)
 - Qasem (TSU), Kennedy, Mellor-Crummey (Rice) -- Whole program tuning
 - Compilers that learn -- Cavazos (UDel); Stephenson/Amarsinghe (MIT)
- Run-time: Voss, et al.: ADAPT



Administrivia



Upcoming schedule changes

- Some adjustment of topics (TBD)
- Today — Project proposals due
- Th 3/13 — SIAM Parallel Processing (attendance encouraged)
- Tu 4/1 — No class
- Th 4/3 — Attend talk by Doug Post from DoD HPC Modernization Program




Homework 1:


Parallel conjugate gradients

- Put name on write-up!
- Grading: **100** pts max
 - Correct implementation — **50** pts
 - **Evaluation** — **45** pts
 - Tested on two samples matrices — 5
 - Implemented and tested on stencil — 10
 - “Explained” performance (e.g., per proc, load balance, comp. vs. comm) — 15
 - Performance model — 15 pts
 - Write-up “quality” — 5 pts

Spy plot of matrix 'msdoor--UF1644'




Non-zeros per row 'msdoor--UF1644'



Non-zeros per row 'msdoor--UF1644'

“active”
elements




Spy plot of matrix 'audikw_1--UF1252'



Non-zeros per row 'audikw_1--UF1252'

nnz

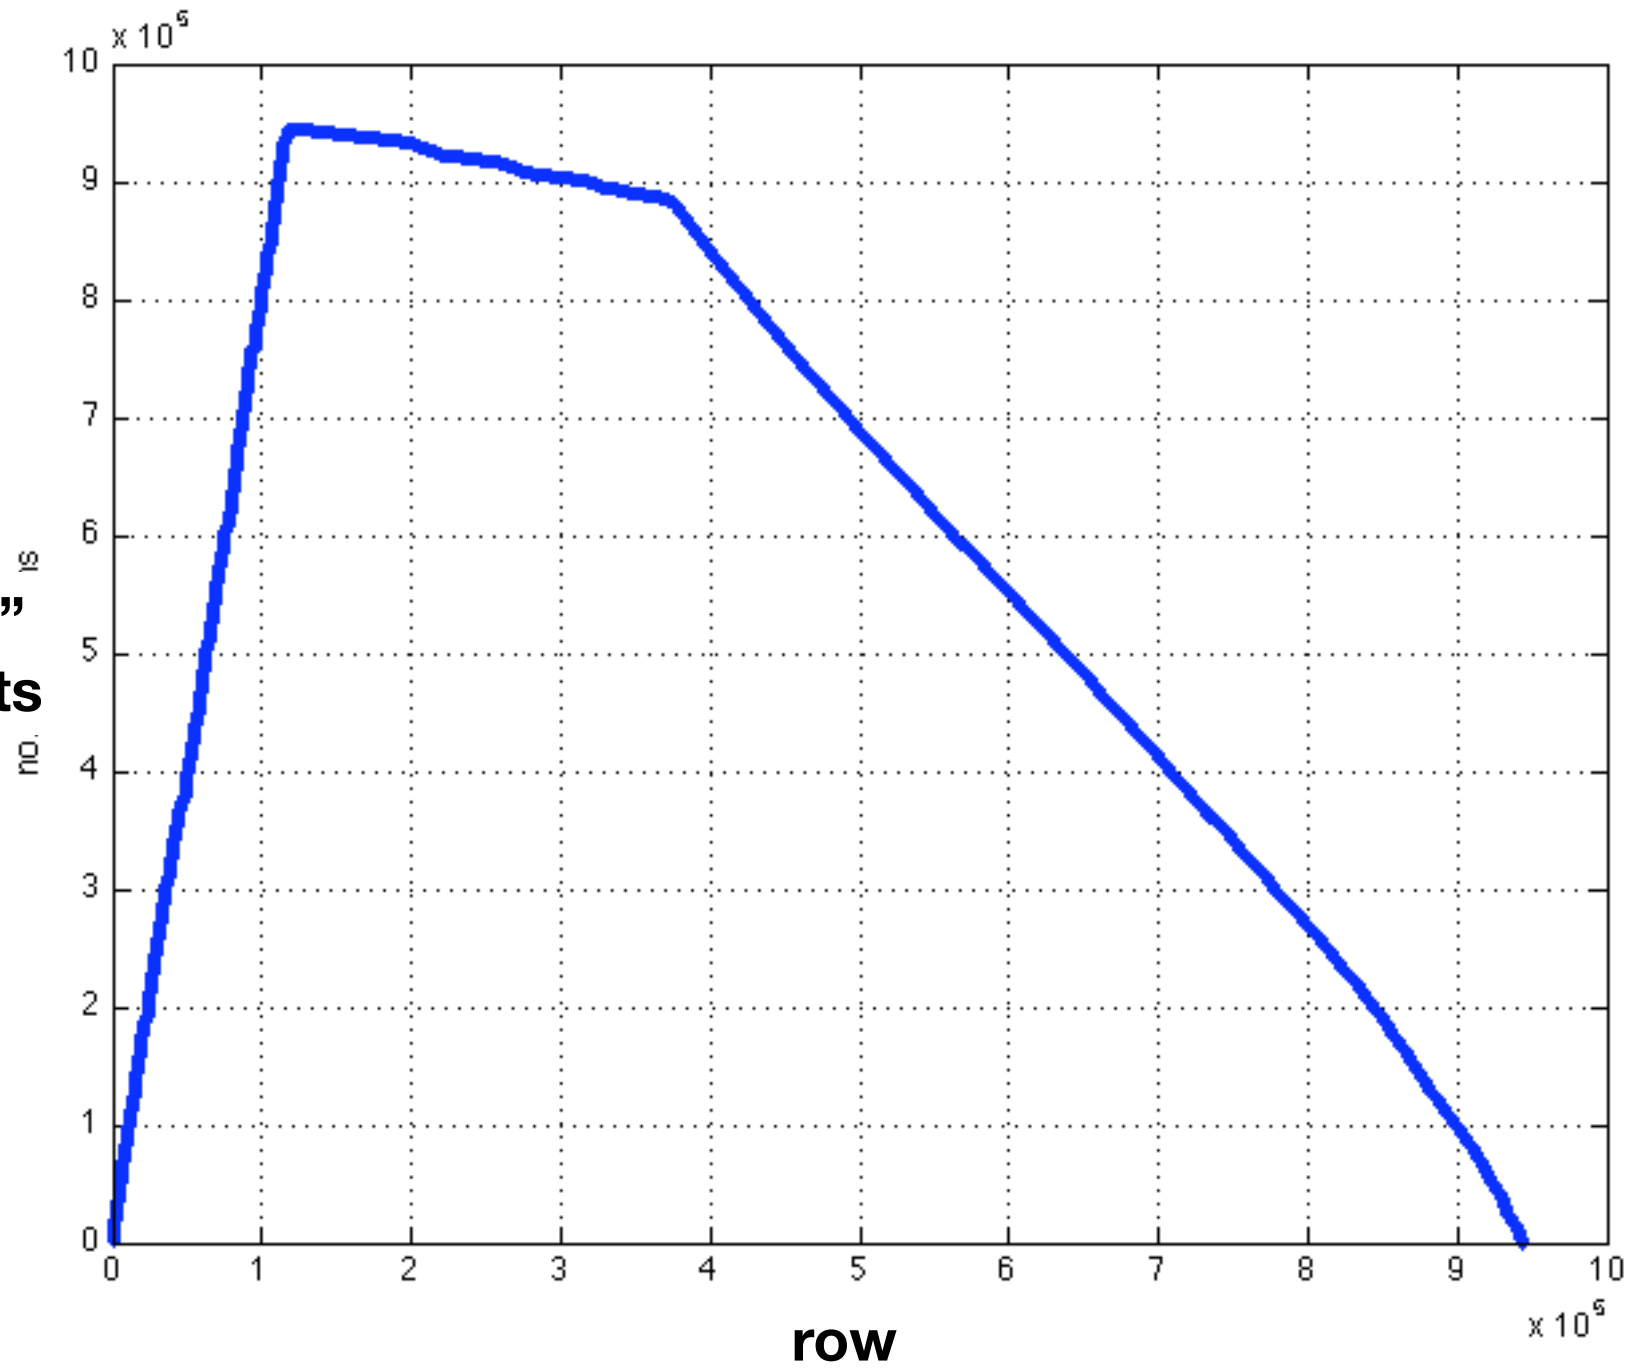


row



Active elements for 'audikw_1--UF1252'

“active”
elements



Homework 2: Parallel n-body using the particle-mesh method

- Acceleration of particle i , due to forces from all other particles:

$$\ddot{\mathbf{r}}_i = \sum_{j \neq i} \mathbf{F}_{ij} = - \sum_{j \neq i} \frac{Gm_j(\mathbf{r}_i - \mathbf{r}_j)}{\|\mathbf{r}_i - \mathbf{r}_j\|^3}$$

- Not yet decided what exactly I will ask to do (implementation? pencil-and-paper? thoughts?)



Projects


- ■ Your goal should be to do something useful, interesting, and/or publishable!
 - ■ Something you're already working on, suitably adapted for this course
 - ■ Faculty-sponsored/mentored
 - ■ Collaborations encouraged



“In conclusion...”



Backup slides



My criteria for “approving” your project

- “Relevant to this course:” Many themes, so think (and “do”) broadly
 - Parallelism and architectures
 - Numerical algorithms
 - Programming models
 - Performance modeling/analysis



General styles of projects

- Theoretical: Prove something hard (high risk)
- Experimental:
 - Parallelize something
 - Take existing parallel program, and improve it using models & experiments
 - Evaluate algorithm, architecture, or programming model

Examples

- *Anything of interest to a faculty member/project outside CoC*
- Parallel sparse triple product ($R^*A^*R^T$, used in multigrid)
- Future FFT
- Out-of-core or I/O-intensive data analysis and algorithms
- Block iterative solvers (convergence & performance trade-offs)
- Sparse LU
- Data structures and algorithms (trees, graphs)
- Look at mixed-precision
- Discrete-event approaches to continuous systems simulation
- Automated performance analysis and modeling, tuning
- “Unconventional,” but related
 - Distributed deadlock detection for MPI
 - UPC language extensions (dynamic block sizes)
 - Exact linear algebra

