

Table of Contents

- Motivation & Trends in HPC
- Mathematical Modeling
- **Numerical Methods used in HPSC**
 - Systems of Differential Equations: ODEs & PDEs
 - Automatic Differentiation
 - Solving Optimization Problems
 - Solving Nonlinear Equations
 - Basic Linear Algebra, Eigenvalues and Eigenvectors
 - **Chaotic systems**
- HPSC Program Development/Enhancement: from Prototype to Production
- Visualization, Debugging, Profiling, Performance Analysis & Optimization



Terminology

- **Chaos**: a well-defined mathematical property of solutions of some nonlinear differential equations
 - Comparable with randomness
- **Attractor**: A set towards which the solutions eventually drift. Various types available:
 - A fixed point (equilibrium)
 - A limit cycle
 - Torus (doughnut)
 - Strange = 'none of the above three'
- **Bifurcation**: if the system splits into two branches
 - Usually happens when a control parameter reaches a critical value



How to recognize chaos?

- The difference between **randomness** and **chaos** is not always clear (sometimes a philosophical matter)
- Chaos in the context of dynamical systems means the superficially random behavior of a deterministic system
- If the laws governing the system are known
 - Future states of the system can be predicted for arbitrarily long time periods
 - If the present state is known arbitrarily accurately
- Purely random systems cannot be predicted



How to recognize chaos? (2)

- Signatures of chaos:
 - Patterns in power spectral analysis of time
 - Series
 - Structures in phase space (attractors, projections of attractors)
 - Fractal dimensions of phase space structures
 - Sensitivity to initial conditions (clear sign of chaos)
 - Controllability of time series
- None of these methods is absolutely foolproof



When to expect chaos?

- Some conditions under which chaotic behavior can occur:
 - Nonlinear systems – e.g. with strong feedback
 - Time delays in continuous systems
 - Finite speed of information
 - Different age groups in population models
 - Spatial discretization
 - External forcing functions



Table of Contents

- Motivation & Trends in HPC
- Mathematical Modeling
- Numerical Methods used in HPSC
 - Systems of Differential Equations: ODEs & PDEs
 - Automatic Differentiation
 - Solving Optimization Problems
 - Solving Nonlinear Equations
 - Basic Linear Algebra, Eigenvalues and Eigenvectors
 - Chaotic systems
- **HPSC Program Development/Enhancement: from Prototype to Production**
- Visualization, Debugging, Profiling, Performance Analysis & Optimization



Program Development – Prototyping

- How to solve a computational problem?
 - Identify the problem
 - Formulate the problem as a mathematical model
 - Choose a suitable computational method for solving the model
 - Implement the method with the right tools
 - Check the results



Program Development - Prototyping (2)

- To understand the problem and the model – start with a simplified version
 - Easier to implement
 - Easier to test
 - Computations take less time
- Move gradually towards the complete model
- Balance the implementation and total computation time



Program Development - Prototyping (3)

- If the problem must be solved only once
 - Performance is probably not critical unless the problem is really big
 - Choose the easiest possible tool to minimize the implementation time
- If the problem needs to be solved repeatedly – various parameter values
 - More time can be spent optimizing the implementation & parallelizing it
 - Attention should be paid to performance



Tools and implementation

- One may use a tool for the prototype problem that is
 - Easy to use
 - Interactive
 - Less efficient
- Good candidates for prototyping
 - Matlab (numerical)
 - Mathematica or Maple (symbolic)
- The final implementation should be as efficient as possible



Tools and implementation (2)

- For final implementation and production runs the best choices are:
 - Compiled programming languages:
 - Fortran 77 & 90
 - C/C++
 - Subroutine libraries
 - NAG
 - IMSL
 - BLAS
 - LAPACK
 - Matlab



Case study: The Problem

- We consider chemical reactions in silicon carbide (SiC) production
- The reacting species may be gases, liquids, and solids
- The amounts of species may vary greatly
- We are looking for the equilibrium state
- The equilibrium state is the minimum of the Gibbs free energy under mass conservation constraints



Case study: Mathematical Model

- The Gibbs free energy G is given by $G = \sum_{i=1}^N n_i \mu_i$
- With
 - N is the number of species present
 - n_i is the molar amount of species i
 - μ_i is the corresponding chemical potential
- The Gibbs free energy is minimized under the requirement that the amounts of basic elements must be conserved:
- Where
$$\sum_i^N a_{ki} n_i = b_k; \quad k = 1, 2, \dots, M$$
 - a_{ki} is the subscript of the k -th element in the molecular formula of species i
 - b_k is the initial amount of element k
 - M is the number of elements
- This is a constrained optimization problem



Case study: Simplified Version

- We consider ideal gases at constant pressure conditions
 - Then the chemical potentials μ_i can be expressed as:
 - With
$$\mu_i = \mu_i^0 + RT \log p_i$$
 - μ_i^0 is the standard chemical potential of species i
 - p_i is the partial pressure
 - T is the temperature
 - R is the universal gas constant
 - For ideal gases the partial pressures are given by:
 - Where
$$p_i = \frac{n_i}{n_t} p_t$$
 - p_t is the total pressure and
 - $n_t = \sum n_i$ is the total molar amount of the species



Case study: Simplified Version (2)

- We thus obtain $\mu_i = \mu_i^* + RT \log \frac{n_i}{n_i^*}$
- Where $\mu_i^* = \mu_i^0 + RT \log p_i$
- We consider only the most important reactants and ignore the ones with low concentrations
 - Decrease the number of unknowns
 - Solve the problem faster
 - Check the results easier
 - Neglecting minor species to avoid numerical problems
- Tools for solving the simplified version
 - Matlab Optimization Toolbox
 - GAMS – General Algebraic Modeling System



Case study: Complete Problem

- Include liquids and solids
 - Their chemical potentials require new equations
- Include all reactants
 - Increases the number of unknowns
- The low concentrations of minor species may cause numerical problems
 - Hand tuned solver may be needed: Fortran 90 + NAG/MKL
- Choose the suitable numerical methods
 - Method of projected gradient
 - Method of augmented Lagrangian
 - Sequential Quadratic Programming (SQP)



Computational Patterns

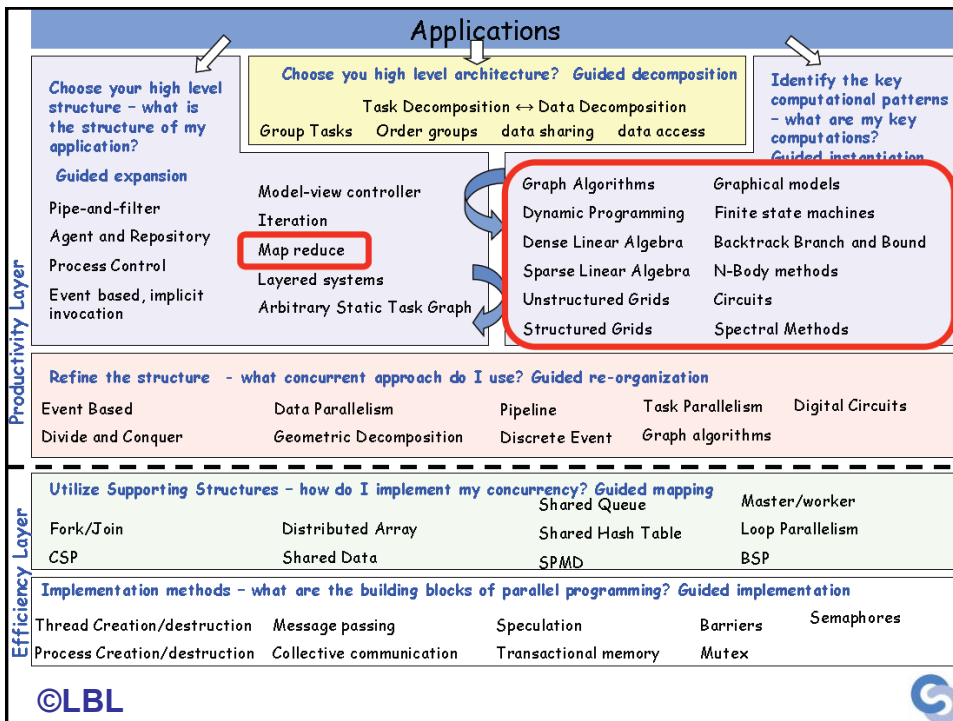
- Productive parallel computing depends on recognizing and exploiting known patterns
 - Design, computational, and mathematical
- Optimizing (some of) the following 7 Motifs
 - To minimize time, minimize communication
- Between levels of the memory hierarchy
- Between processors over a network
 - Autotuning to explore large design spaces
- Too hard (tedious) to write by hand, let machine do it



“7 Motifs” of High Performance Computing

- Phil Colella (LBL) identified 7 kernels of which most simulation and data-analysis programs are composed:
 1. Dense Linear Algebra
 - Ex: Solve $Ax=b$ or $Ax = \lambda x$ where A is a dense matrix
 2. Sparse Linear Algebra
 - Ex: Solve $Ax=b$ or $Ax = \lambda x$ where A is a sparse matrix (mostly zero)
 3. Operations on Structured Grids
 - Ex: $A_{new}(i,j) = 4*A(i,j) - A(i-1,j) - A(i+1,j) - A(i,j-1) - A(i,j+1)$
 4. Operations on Unstructured Grids
 - Ex: Similar, but list of neighbors varies from entry to entry
 5. Spectral Methods
 - Ex: Fast Fourier Transform (FFT)
 6. Particle Methods
 - Ex: Compute electrostatic forces on n particles
 7. Monte Carlo, Embarrassing Parallelism, Map Reduce
 - Ex: Many independent simulations using different inputs





What you (might) want to know about a motif

216

- How to use it
 - What problems does it solve?
 - How to choose solution approach, if more than one?
 - How to find the best software available now
 - Best: fastest? most accurate? fewest keystrokes?
 - How are the best implementations built?
 - What is the “design space” (w.r.t. math and CS)?
 - How do we search for best (autotuning)?
-

The Dense Linear Algebra Motif

- In the beginning was the do-loop...
- Libraries like EISPACK (for eigenvalue problems)
- Then the BLAS (1) were invented (1973-1977)
 - Standard library of 15 operations (mostly) on vectors
 - “AXPY” ($y = \alpha \cdot x + y$), dot product, scale ($x = \alpha \cdot x$), etc
 - Up to 4 versions of each (S/D/C/Z), 46 routines, 3300 LOC
 - Goals
 - Common “pattern” to ease programming, readability
 - Robustness, via careful coding (avoiding over/underflow)
 - Portability + Efficiency via machine specific implementations
 - Why BLAS 1 ? They do $O(n^1)$ ops on $O(n^1)$ data
 - Used in libraries like LINPACK (for linear systems)
 - Source of the name “LINPACK Benchmark” (not the code!)



The Dense Linear Algebra Motif (2)

- But BLAS-1 weren't enough
 - Consider AXPY ($y = \alpha \cdot x + y$): $2n$ flops on $3n$ read/writes
 - Computational intensity = $(2n)/(3n) = 2/3$
 - Too low to run near peak speed (read/write dominates)
 - Hard to vectorize on supercomputers of the day (1980s)
- So the BLAS-2 were invented (1984-1986)
 - Standard library of 25 operations on matrix/vector pairs
 - “GEMV”: $y = \alpha \cdot A \cdot x + \beta \cdot x$, “GER”: $A = A + \alpha \cdot x \cdot y^T$, $x = T^{-1} \cdot x$
 - Up to 4 versions of each (S/D/C/Z), 66 routines, 18K LOC
 - Why BLAS 2 ? They do $O(n^2)$ ops on $O(n^2)$ data
 - So computational intensity still just $\sim (2n^2)/(n^2) = 2$
 - OK for vector machines, but not for machine with caches



The Dense Linear Algebra Motif (3)

- The next step: BLAS-3 (1987-1988)
 - Standard library of 9 operations on matrix/matrix pairs
 - “GEMM”: $C = \alpha \cdot A \cdot B + \beta \cdot C$, $C = \alpha \cdot A \cdot A^T + \beta \cdot C$, $B = T^{-1} \cdot B$
 - Up to 4 versions of each (S/D/C/Z), 30 routines, 10K LOC
 - Why BLAS 3? They do $O(n^3)$ ops on $O(n^2)$ data
 - So computational intensity $(2n^3)/(4n^2) = n/2$ – big at last!
 - Good for machines with caches, other mem. hierarchy levels
- How much BLAS1/2/3 code so far
 - Source: 142 routines, 31K LOC, Testing: 28K LOC
 - Reference (unoptimized) implementation only
 - Ex: 3 nested loops for GEMM
 - Lots more optimized code
 - Motivates “automatic tuning” of the BLAS



The Dense Linear Algebra Motif (4)

- LAPACK – “Linear Algebra PACKage”: BLAS-3 (1989 – now)
 - Ex: Obvious way to express Gaussian Elimination (GE) is adding multiples of one row to other rows – BLAS-1
 - How do we reorganize GE to use BLAS-3? (details later)
 - Contents of LAPACK (summary)
 - Algorithms we can turn into (nearly) 100% BLAS 3: Linear Systems & Least squares
 - Algorithms that are only 50% BLAS 3 (so far): Eigenproblems & Singular Value Decomposition (SVD)
 - Error bounds for everything
 - Lots of variants depending on A’s structure
 - How much code? (Nov 2008) (www.netlib.org/lapack)
 - Source: 1582 routines, 490K LOC, Testing: 352K LOC




The Dense Linear Algebra Motif (5)

- Is LAPACK parallel?
 - Only if the BLAS are parallel (shared memory)
- ScaLAPACK – “Scalable LAPACK” (1995 – now)
 - For distributed memory – uses MPI
 - More complex data structures, algorithms than LAPACK
 - Only subset of LAPACK’s functionality available
 - All at www.netlib.org/scalapack



Success Stories for Sca/LAPACK

- Widely used
 - Adopted by Mathworks, Cray, Fujitsu, HP, IBM, IMSL, Intel, NAG, NEC, SGI
 - 5.5M webhits/year @ Netlib (incl. CLAPACK, LAPACK95)
- New Science discovered through the solution of dense matrix systems
 - Nature article on the flat universe used ScaLAPACK
 - Other articles in Physics Review B that also use it
 - 1998 Gordon Bell Prize
 - www.nersc.gov/news/reports/newNERSResults050703.pdf

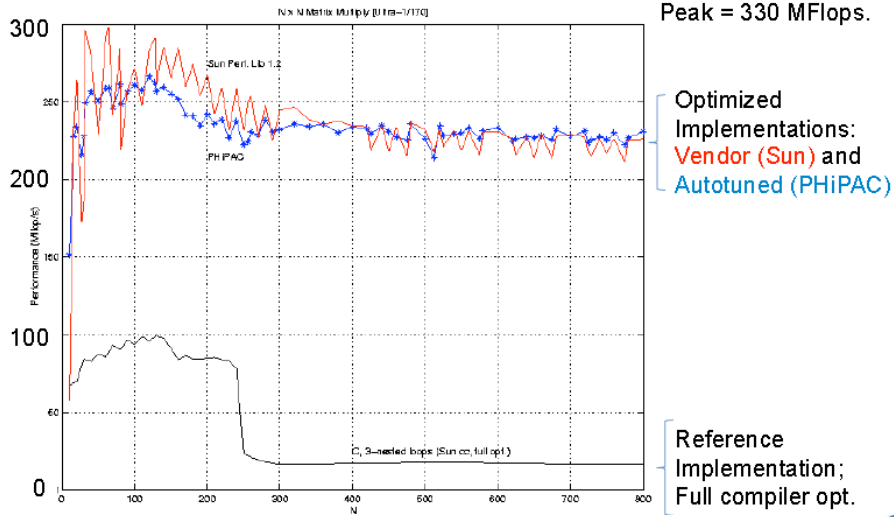


Cosmic Microwave Background Analysis, BOOMERanG collaboration, MADCAP code (Apr. 27, 2000).

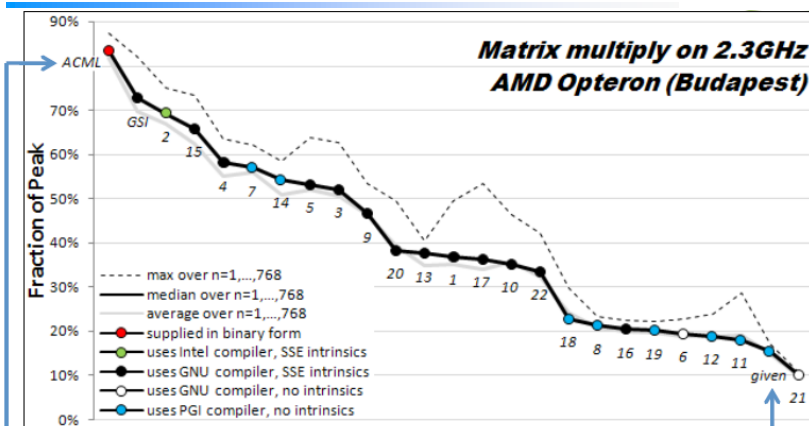
ScaLAPACK



Optimized Matrix-Multiply



How hard is hand-tuning?



- Results of 22 student teams trying to tune matrix-multiply, in CS267 Spr09
- Students given “blocked” code to start with
 - Still hard to get close to vendor tuned performance (ACML)
- For more discussion, see www.cs.berkeley.edu/~volkov/cs267.sp09/hw1/results/
- Naïve matmul: just 2% of peak

©LBL



The Future of Dense Linear Algebra

- Communication-Avoiding for everything
- Extensions for multicore systems
 - PLASMA – Parallel Linear Algebra for Scalable Multicore Architectures
 - Dynamically schedule tasks into which the algorithm is decomposed: minimize synchronization & keep all processors busy
- Extensions for GPUs
 - “Benchmarking GPUs to tune Dense Linear Algebra”
 - Best Student Paper Prize at SC08 (Vasily Volkov)
 - MAGMA – Matrix Algebra on GPU and Multicore Architectures
- How much code generation can we automate?
 - MAGMA , and FLAME (www.cs.utexas.edu/users/flame/)



Sparse Linear Algebra Motif

- Similar problems to dense matrices
 - $Ax=b$, Least squares, $Ax = \lambda x$, SVD, ...
- But different algorithms!
 - Exploit structure: only store, work on non-zeros
 - Direct methods
 - LU, Cholesky for $Ax=b$, QR for Least squares
 - See crd.lbl.gov/~xiaoye/SuperLU/index.html for LU codes
 - See crd.lbl.gov/~xiaoye/SuperLU/SparseDirectSurvey.pdf for a survey of available serial and parallel sparse solvers
 - Iterative methods – for $Ax=b$, least squares, eig, SVD
 - Use simplest operation: Sparse-Matrix-Vector-Multiply (SpMV)
 - Krylov Subspace Methods: find “best” solution in space spanned by vectors generated by SpMVs



Sparse Linear Algebra Motif (2)

- Fast code must **minimize communication**
 - Especially for sparse matrix computations because communication dominates
- Generating fast code for a single SpMV
 - Design space of possible algorithms must be searched at run-time, when sparse matrix available
 - Design space should be searched automatically
- Biggest speedups from minimizing communication in an entire sparse solver
 - Many more opportunities to minimize communication in multiple SpMVs than in one
 - Requires transforming the entire algorithm
- More information can be found: bebop.cs.berkeley.edu



ODEs and Sparse Matrices

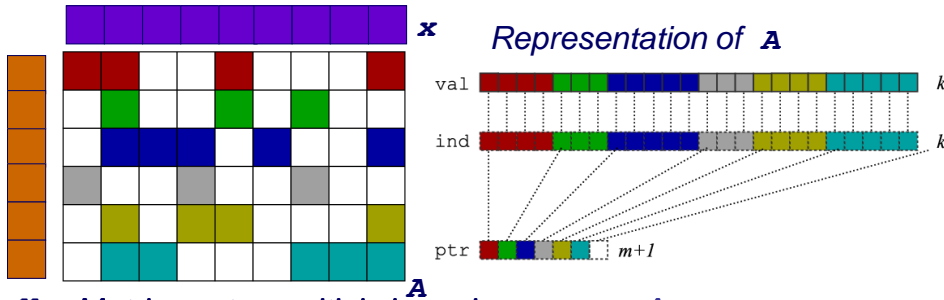
- ODEs/PDEs problems reduce to sparse matrix problems
 - Explicit: sparse matrix-vector multiplication (SpMV).
 - Implicit: solve a sparse linear system
 - Direct solvers (Gaussian elimination)
 - Iterative solvers (use sparse matrix-vector multiplication)
 - Eigenvalue/vector algorithms may also be explicit or implicit
- Conclusion: SpMV is key to many ODE problems
 - Relatively simple algorithm to study in detail
 - Two key problems: locality and load balance



SpMV in Compressed Sparse Row (CSR) Format

231

SpMV: $y = y + A \cdot x$, only store, do arithmetic, on nonzero entries
CSR format is simplest one of many possible data structures for A



Y Matrix-vector multiply kernel: $y(i) \leftarrow y(i) + A(i,j) \cdot x(j)$
 for each row i
 for $k = \text{ptr}[i]$ to $\text{ptr}[i+1] - 1$ do
 $y[i] = y[i] + \text{val}[k] * x[\text{ind}[k]]$




Parallel Sparse Matrix-vector multiplication

232

- $y = A \cdot x$, where A is a sparse $n \times n$ matrix

Questions

- which processors store
 - $y[i]$, $x[i]$, and $A[i,j]$
- which processors compute
 - $y[i] = \text{sum}(\text{from } 1 \text{ to } n) A[i,j] * x[j]$
 $= (\text{row } i \text{ of } A) * x$... a sparse dot product



Partitioning

- Partition index set $\{1, \dots, n\} = N1 \cup N2 \cup \dots \cup Np$.
- For all i in Nk , Processor k stores $y[i]$, $x[i]$, and row i of A
- For all i in Nk , Processor k computes $y[i] = (\text{row } i \text{ of } A) * x$
 - “owner computes” rule: Processor k computes the $y[i]$ s it owns.

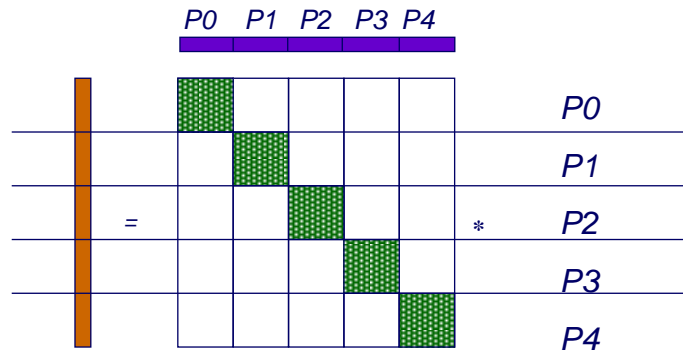
May require communication



Matrix Reordering via Graph Partitioning

233

- “Ideal” matrix structure for parallelism: block diagonal
 - p (number of processors) blocks, can all be computed locally
 - If no non-zeros outside these blocks, no communication needed
- Can we reorder the rows/columns to get close to this?
 - Most non-zeros in diagonal blocks, few outside



Goals of Reordering

234

- Performance goals
 - Balance load – how is load measured?
 - Approx equal number of non-zeros (not necessarily rows)
 - Balance storage – how much does each processor store?
 - Approx equal number of non-zeros
 - Minimize communication – how much is communicated?
 - Minimize non-zeros outside diagonal blocks
 - Related optimization criterion is to move non-zeros near diagonal
 - Improve register and cache re-use
 - Group non-zeros in small vertical blocks so source (x) elements loaded into cache or registers may be reused (temporal locality)
 - Group non-zeros in small horizontal blocks so nearby source (x) elements in the cache may be used (spatial locality)
- Other algorithms reorder for other reasons
 - Reduce # non-zeros in matrix after Gaussian elimination
 - Improve numerical stability




Graph Partitioning and Sparse Matrices

235

- Relationship between matrix and graph

	1	2	3	4	5	6
1	1			1	1	
2		1	1	1	1	
3			1			1
4	1	1		1		1
5	1	1			1	1
6			1	1	1	1



- Edges in the graph are nonzero in the matrix: here the matrix is symmetric (edges are unordered) and weights are equal (1)
- If divided over 3 procs, there are 14 nonzeros outside the diagonal blocks, which represent the 7 (bidirectional) edges




Graph Partitioning and Sparse Matrices (2)

236

- Relationship between matrix and graph

	1	2	3	4	5	6
1	1	1			1	
2	1	1		1		1
3			1	1		1
4		1	1	1	1	
5	1			1	1	1
6		1	1		1	1



- A "good" partition of the graph has
 - equal (weighted) number of nodes in each part (load and storage balance).
 - minimum number of edges crossing between (minimize communication).
- Reorder the rows/columns by putting all nodes in one partition together.



Summary: Common Problems

- Load Balancing
 - Dynamically – if load changes significantly during job
 - Statically - Graph partitioning
 - Discrete systems
 - Sparse matrix vector multiplication
- Linear algebra
 - Solving linear systems (sparse and dense)
 - Eigenvalue problems will use similar techniques
- Fast Particle Methods
 - $O(n \log n)$ instead of $O(n^2)$

