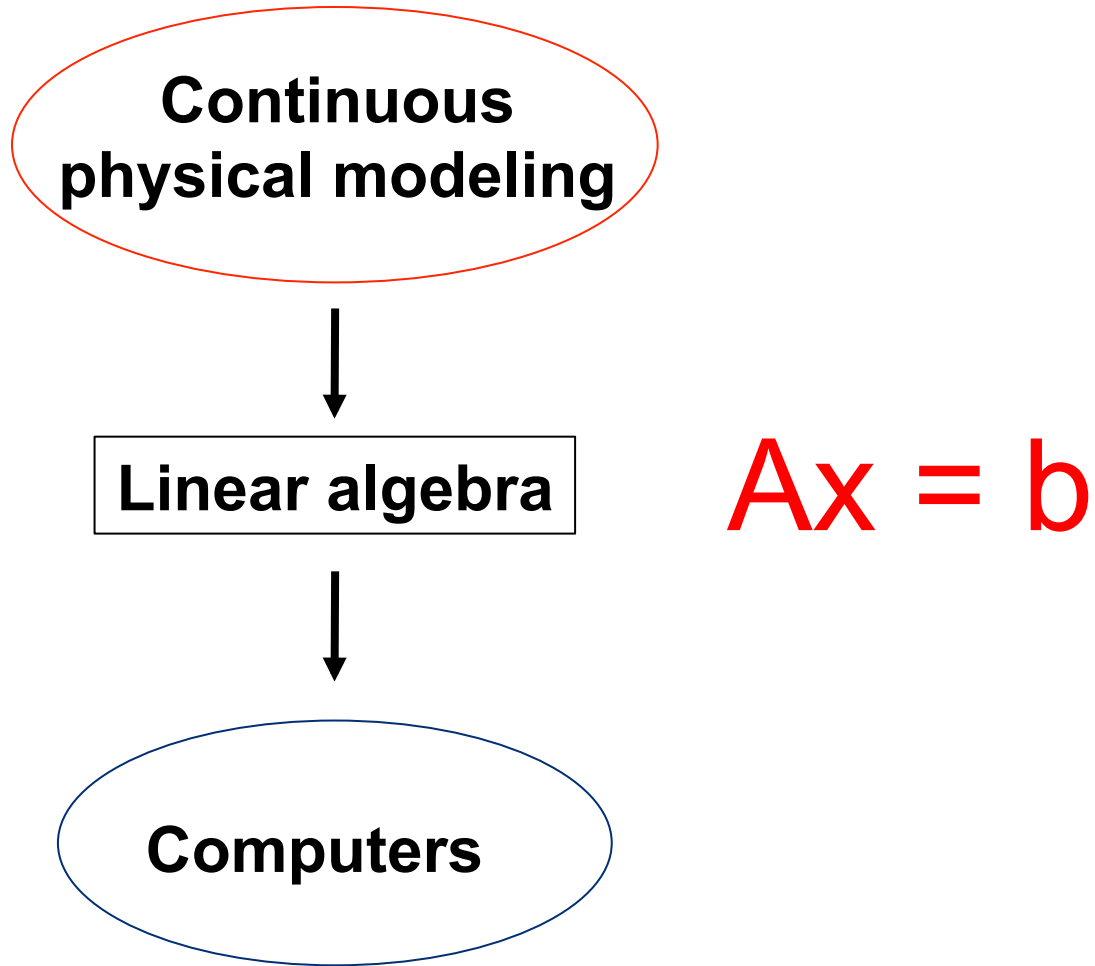


***CS240A: Conjugate Gradients
and the Model Problem***

The middleware of scientific computing



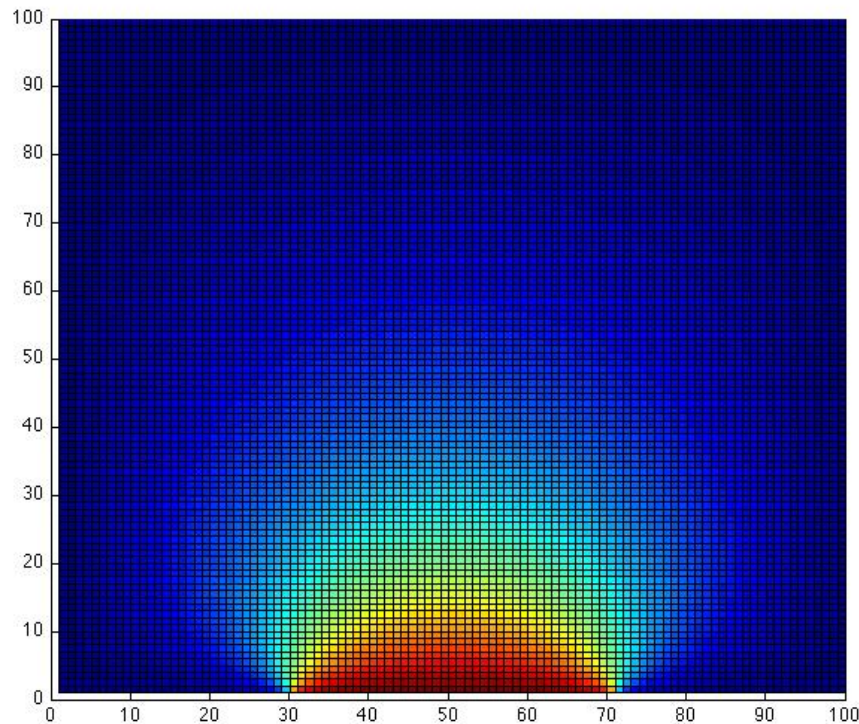
Example: The Temperature Problem

- A cabin in the snow
- Wall temperature is 0° , except for a radiator at 100°
- What is the temperature in the interior?



Example: The Temperature Problem

- A cabin in the snow (a square region ☺)
- Wall temperature is 0° , except for a radiator at 100°
- What is the temperature in the interior?



The physics: Poisson's equation

$$\nabla^2 u(x, y) \equiv \frac{\partial^2 u}{\partial x^2}(x, y) + \frac{\partial^2 u}{\partial y^2}(x, y) = f(x, y)$$

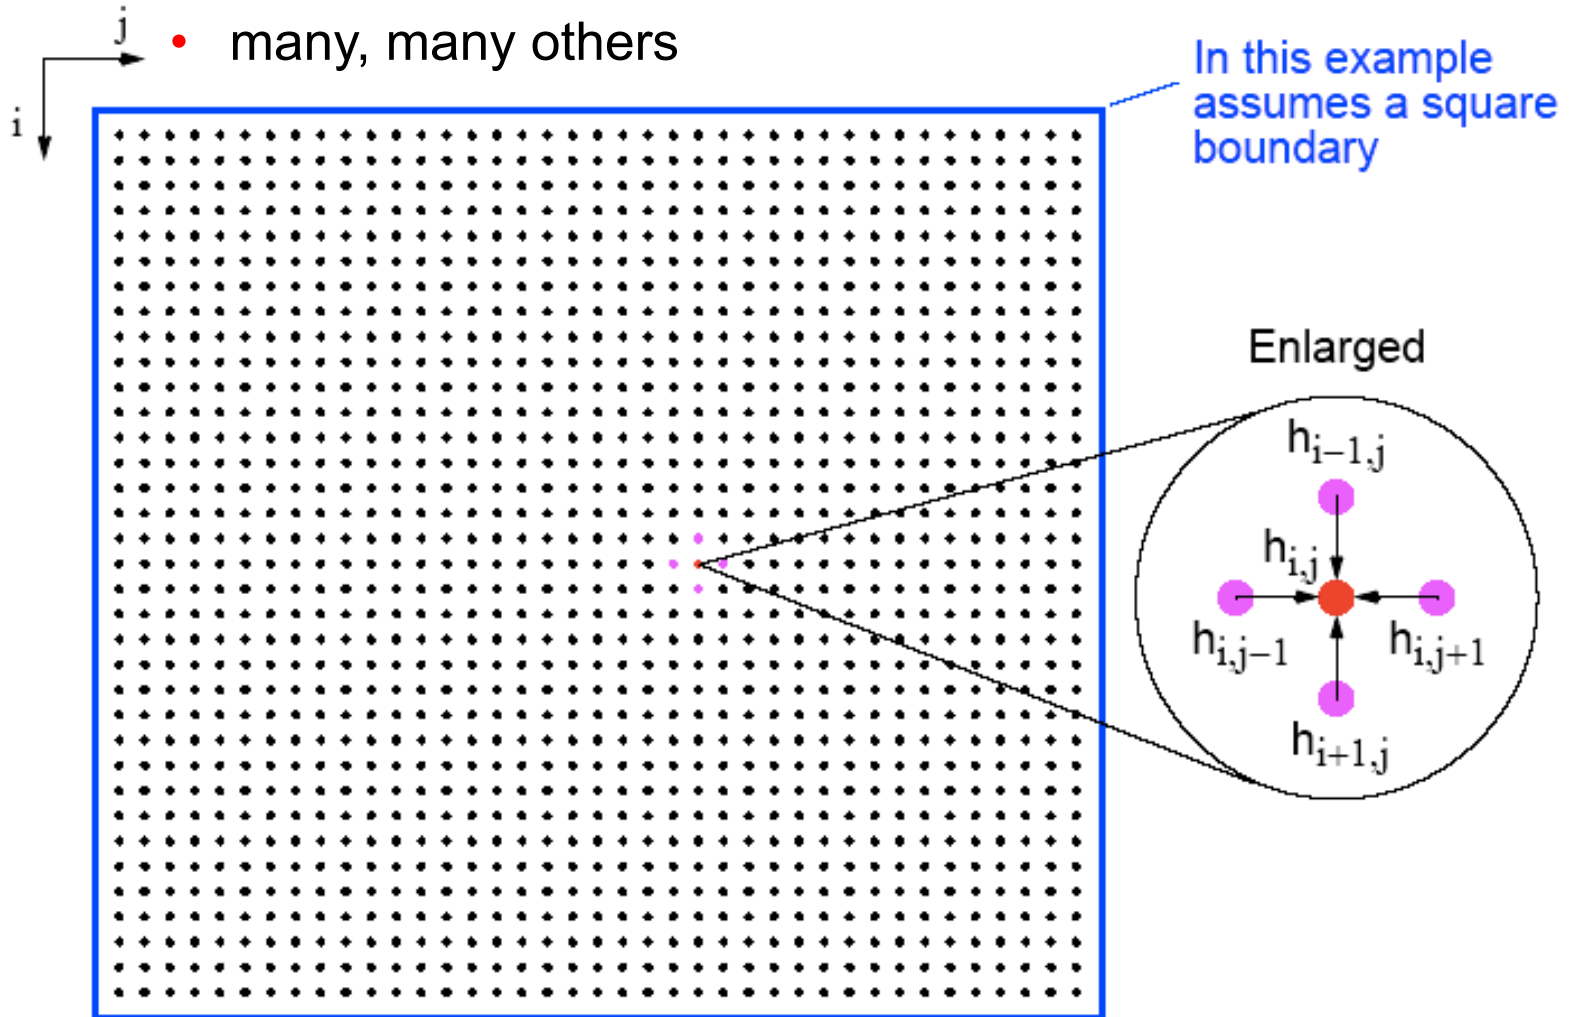
for $(x, y) \in R = \{ (x, y) \mid a < x < b, \ c < y < d \}$, and

$$u(x, y) = g(x, y)$$

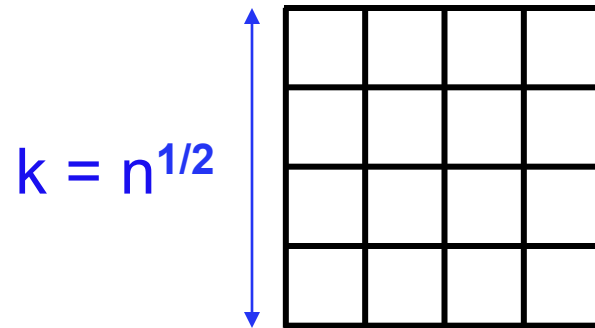
for (x, y) on the boundary of R .

Many Physical Models Use Stencil Computations

- PDE models of heat, fluids, structures, ...
- Weather, airplanes, bridges, bones, ...
- Game of Life
- many, many others



Model Problem: Solving Poisson's equation for temperature



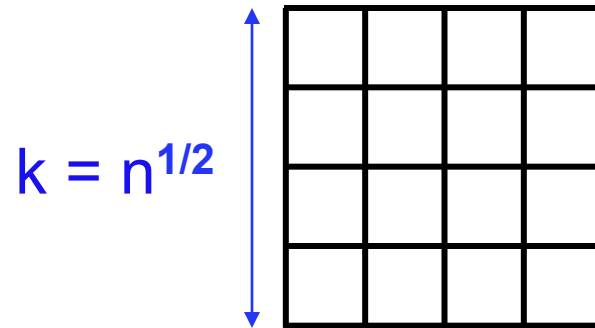
- Discrete approximation to Poisson's equation:

$$t(i) = \frac{1}{4} (t(i-k) + t(i-1) + t(i+1) + t(i+k))$$

- Intuitively:

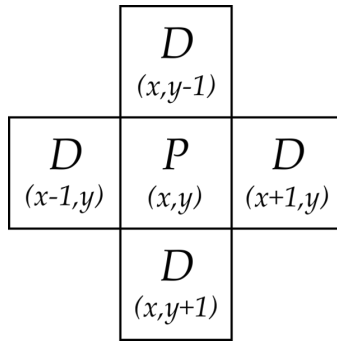
Temperature at a point is the average
of the temperatures at surrounding points

Model Problem: Solving Poisson's equation for temperature

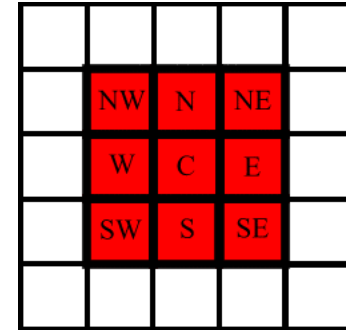


- For each i from 1 to n , except on the boundaries:
$$-t(i-k) - t(i-1) + 4t(i) - t(i+1) - t(i+k) = 0$$
- n equations in n unknowns: $A*t = b$
- Each row of A has at most 5 nonzeros
- In three dimensions, $k = n^{1/3}$ and each row has at most 7 nzs

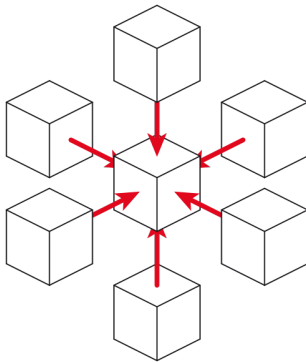
Examples of stencils



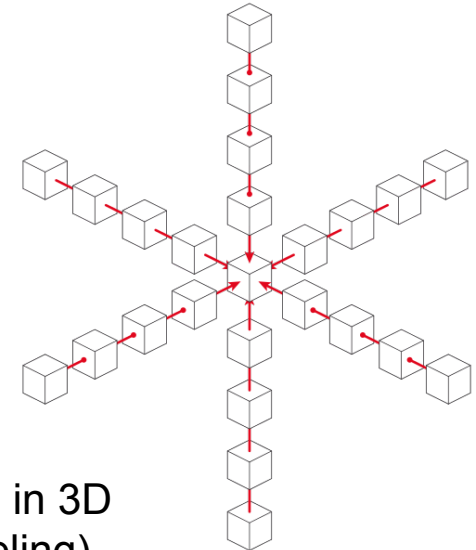
5-point stencil in 2D
(temperature problem)



9-point stencil in 2D
(game of Life)



7-point stencil in 3D
(3D temperature problem)



25-point stencil in 3D
(seismic modeling)

... and many more

A Stencil Computation Solves a System of Linear Equations

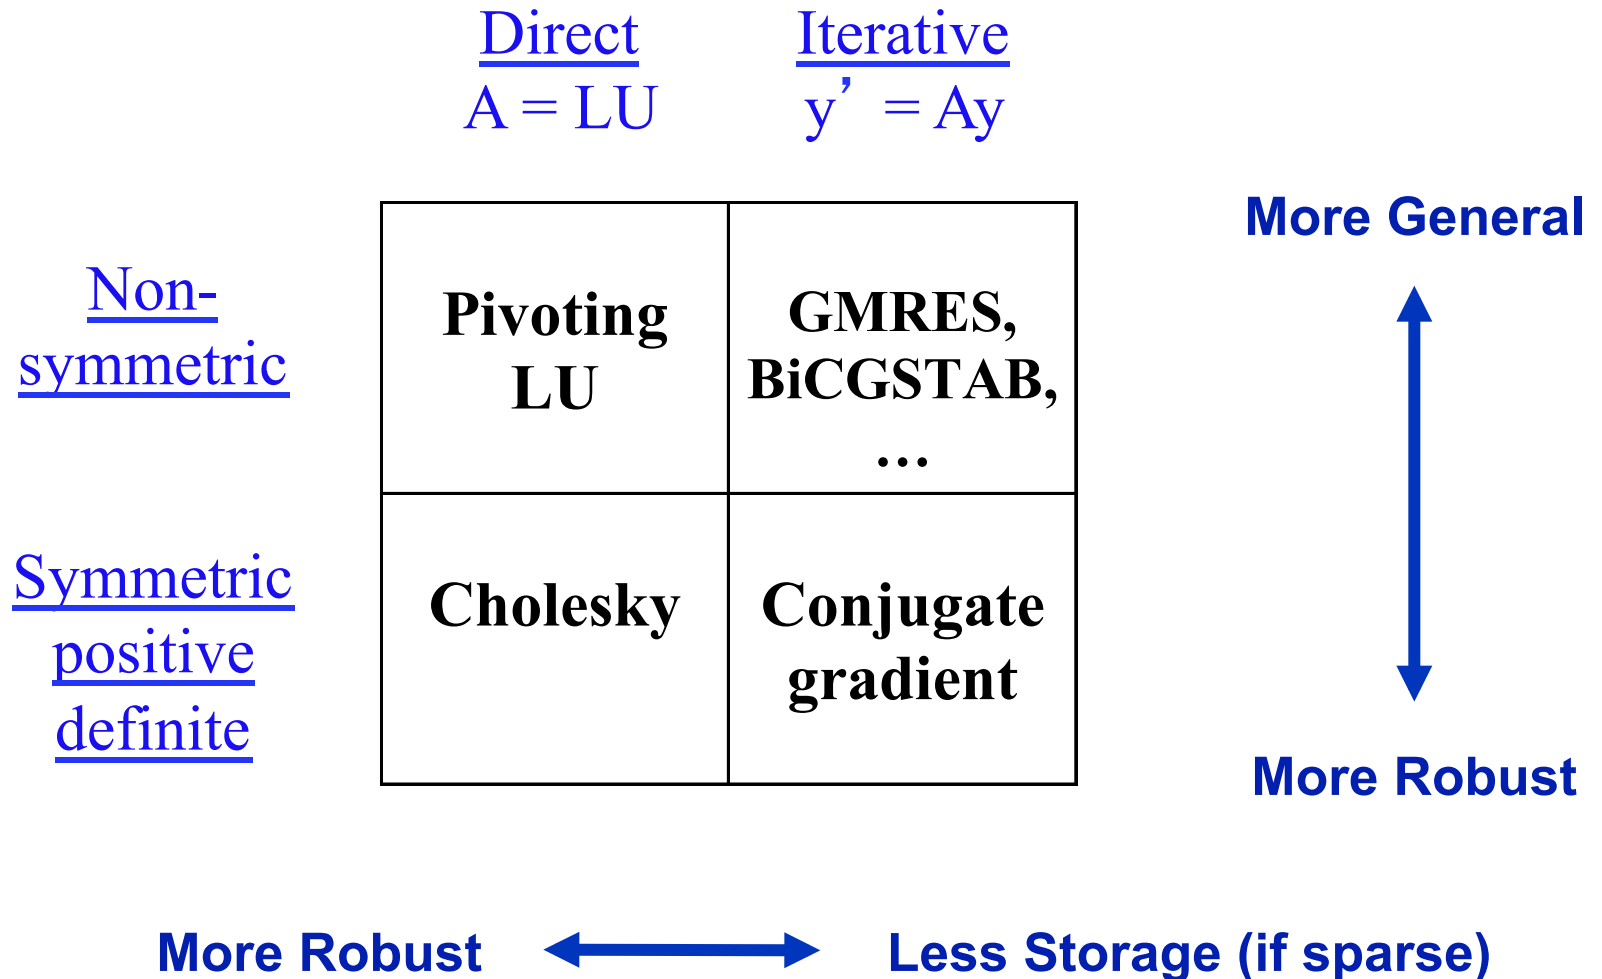
- Solve $Ax = b$ for x
- Matrix A , right-hand side vector b , unknown vector x
- A is *sparse*: most of the entries are 0

The diagram illustrates the system of linear equations $Ax = b$ for a stencil computation. The matrix A is shown as a large vertical column with a dashed diagonal line representing the main diagonal. A specific row, labeled " i th equation" in purple, is highlighted with a box. This row contains the coefficients $a_{i,i-n}$, $a_{i,i-1}$, $a_{i,i}$, $a_{i,i+1}$, and $a_{i,i+n}$ in red, with values 1, 1, -4, 1, and 1 in purple above them. The matrix is annotated with "Those equations with a boundary point on diagonal unnecessary for solution" pointing to the dashed diagonal and "To include boundary values and some zero entries (see text)" pointing to the vector x .

The vector x is shown as a vertical column with elements $x_1, x_2, \dots, x_{N-1}, x_N$. The right-hand side vector b is shown as a vertical column with zeros. The equation is represented as:

$$A \times x = b$$

The Landscape of $Ax=b$ Solvers



CS 240A: Solving $Ax = b$ in parallel

- Dense A: Gaussian elimination with partial pivoting (LU)
 - See Jim Demmel's slides
 - Same flavor as matrix * matrix, but more complicated
- Sparse A: Iterative methods – Conjugate gradient, etc.
 - Sparse matrix times dense vector
- Sparse A: Gaussian elimination – Cholesky, LU, etc.
 - Graph algorithms
- Sparse A: Preconditioned iterative methods and multigrid
 - Mixture of lots of things

CS 240A: Solving $Ax = b$ in parallel

- Dense A: Gaussian elimination with partial pivoting
 - See Jim Demmel's slides
 - Same flavor as matrix * matrix, but more complicated
- Sparse A: Iterative methods – Conjugate gradient etc.
 - Sparse matrix times dense vector
- Sparse A: Gaussian elimination – Cholesky, LU, etc.
 - Graph algorithms
- Sparse A: Preconditioned iterative methods and multigrid
 - Mixture of lots of things

Conjugate gradient iteration for $Ax = b$

$x_0 = 0$ approx solution

$r_0 = b$ residual = $b - Ax$

$d_0 = r_0$ search direction

for $k = 1, 2, 3, \dots$

$x_k = x_{k-1} + \dots$ new approx solution

$r_k = \dots$ new residual

$d_k = \dots$ new search direction

Conjugate gradient iteration for $Ax = b$

$x_0 = 0$ approx solution

$r_0 = b$ residual = $b - Ax$

$d_0 = r_0$ search direction

for $k = 1, 2, 3, \dots$

$\alpha_k = \dots$ step length

$x_k = x_{k-1} + \alpha_k d_{k-1}$ new approx solution

$r_k = \dots$ new residual

$d_k = \dots$ new search direction

Conjugate gradient iteration for $Ax = b$

$x_0 = 0$ approx solution

$r_0 = b$ residual = $b - Ax$

$d_0 = r_0$ search direction

for $k = 1, 2, 3, \dots$

$\alpha_k = (r_{k-1}^T r_{k-1}) / (d_{k-1}^T A d_{k-1})$ step length

$x_k = x_{k-1} + \alpha_k d_{k-1}$ new approx solution

$r_k = \dots$ new residual

$d_k = \dots$ new search direction

Conjugate gradient iteration for $Ax = b$

$x_0 = 0$ approx solution

$r_0 = b$ residual = $b - Ax$

$d_0 = r_0$ search direction

for $k = 1, 2, 3, \dots$

$\alpha_k = (r_{k-1}^T r_{k-1}) / (d_{k-1}^T A d_{k-1})$ step length

$x_k = x_{k-1} + \alpha_k d_{k-1}$ new approx solution

$r_k = \dots$ new residual

$\beta_k = (r_k^T r_k) / (r_{k-1}^T r_{k-1})$

$d_k = r_k + \beta_k d_{k-1}$ new search direction

Conjugate gradient iteration for $Ax = b$

$x_0 = 0$ approx solution

$r_0 = b$ residual = $b - Ax$

$d_0 = r_0$ search direction

for $k = 1, 2, 3, \dots$

$\alpha_k = (r_{k-1}^T r_{k-1}) / (d_{k-1}^T A d_{k-1})$ step length

$x_k = x_{k-1} + \alpha_k d_{k-1}$ new approx solution

$r_k = r_{k-1} - \alpha_k A d_{k-1}$ new residual

$\beta_k = (r_k^T r_k) / (r_{k-1}^T r_{k-1})$

$d_k = r_k + \beta_k d_{k-1}$ new search direction

Conjugate gradient iteration to solve $A^*x=b$

$x_0 = 0, \quad r_0 = b, \quad d_0 = r_0$ (these are all vectors)

for $k = 1, 2, 3, \dots$

$$\alpha_k = (r_{k-1}^T r_{k-1}) / (d_{k-1}^T A d_{k-1}) \quad \text{step length}$$

$$x_k = x_{k-1} + \alpha_k d_{k-1} \quad \text{approximate solution}$$

$$r_k = r_{k-1} - \alpha_k A d_{k-1} \quad \text{residual} = b - A x_k$$

$$\beta_k = (r_k^T r_k) / (r_{k-1}^T r_{k-1}) \quad \text{improvement}$$

$$d_k = r_k + \beta_k d_{k-1} \quad \text{search direction}$$

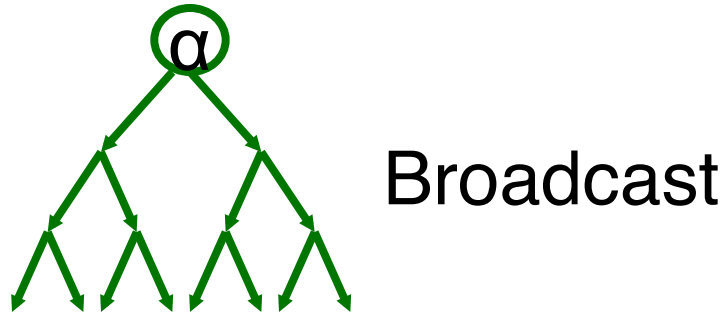
- One matrix-vector multiplication per iteration
- Two vector dot products per iteration
- Four n-vectors of working storage

Vector and matrix primitives for CG

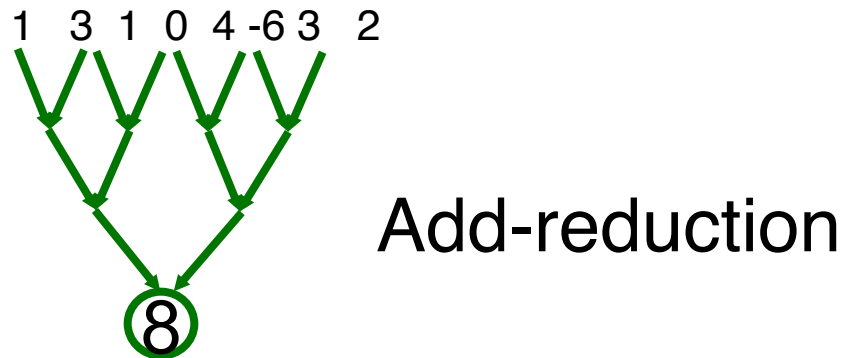
- **DAXPY:** $v = \alpha * v + \beta * w$ (vectors v, w ; scalars α, β)
 - **Broadcast** the scalars α and β , then independent $*$ and $+$
 - **comm volume** = $2p$, **span** = $\log n$
- **DDOT:** $\alpha = v^T * w = \sum_j v[j] * w[j]$ (vectors v, w ; scalar α)
 - Independent $*$, then $+$ **reduction**
 - **comm volume** = p , **span** = $\log n$
- **Matvec:** $v = A * w$ (matrix A , vectors v, w)
 - The hard part
 - But all you need is a subroutine to compute v from w
 - Sometimes you don't need to store A (e.g. temperature problem)
 - Usually you do need to store A , but it's *sparse* ...

Broadcast and reduction

- **Broadcast** of 1 value to p processors in $\log p$ time



- **Reduction** of p values to 1 in $\log p$ time
- Takes advantage of associativity in $+$, $*$, \min , \max , etc.

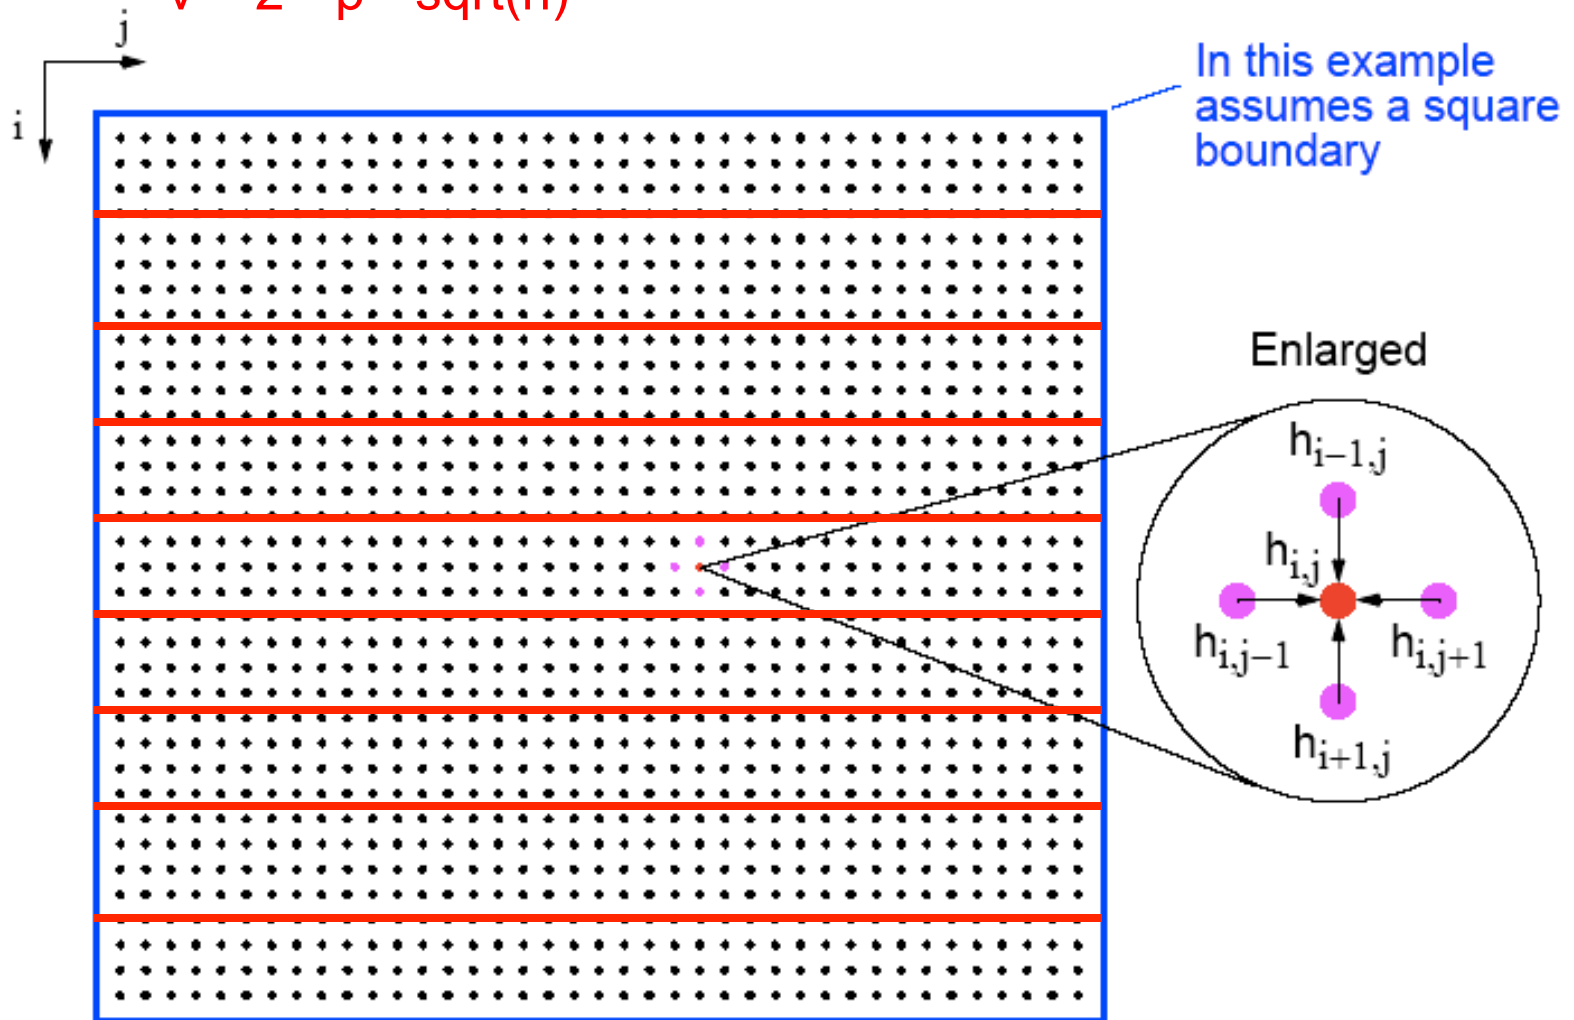


Where's the data (temperature problem)?

- The matrix A: **Nowhere!!**
- The vectors x, b, r, d:
 - Each vector is one value per stencil point
 - Divide stencil points among processors, **n/p** points each
- How do you divide up the **sqrt(n)** by **sqrt(n)** region of points?
- Block row (or block col) layout: **$v = 2 * p * \text{sqrt}(n)$**
- 2-dimensional block layout: **$v = 4 * \text{sqrt}(p) * \text{sqrt}(n)$**

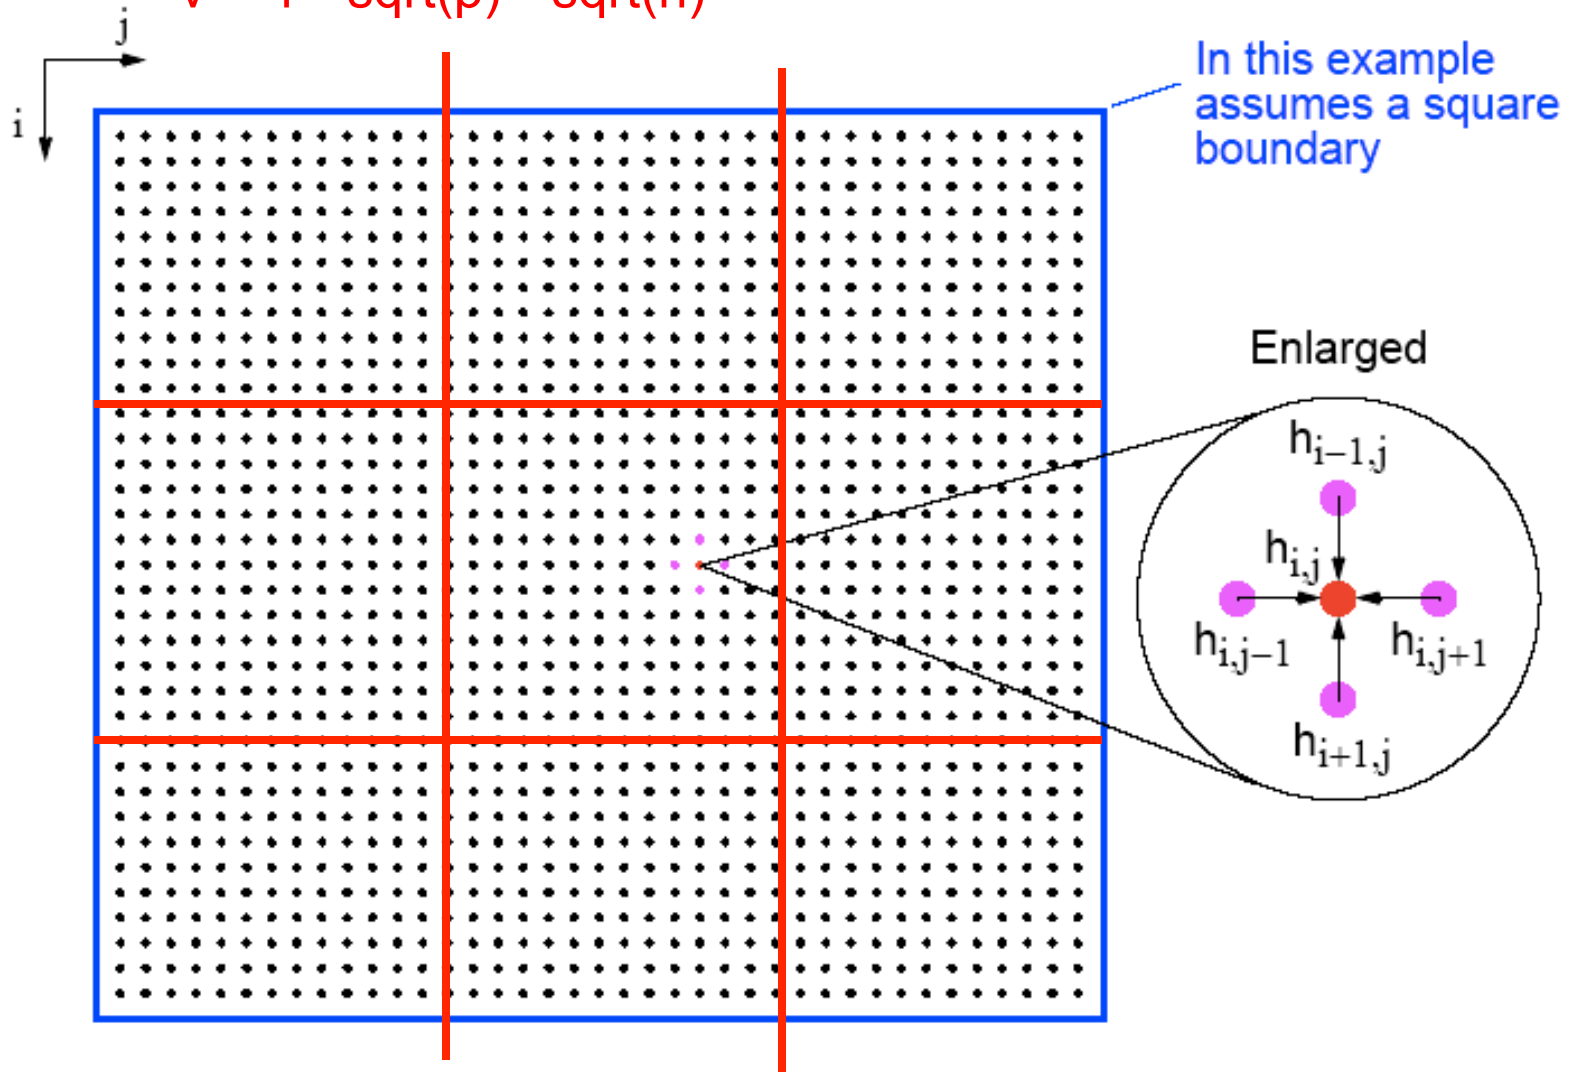
How do you partition the $\text{sqrt}(n)$ by $\text{sqrt}(n)$ stencil points?

- First version: number the grid by rows
- Leads to a block row decomposition of the region
- $v = 2 * p * \text{sqrt}(n)$



How do you partition the $\text{sqrt}(n)$ by $\text{sqrt}(n)$ stencil points?

- Second version: 2D block decomposition
- Numbering is a little more complicated
- $v = 4 * \text{sqrt}(p) * \text{sqrt}(n)$



Where's the data (temperature problem)?

- The matrix A: **Nowhere!!**
- The vectors x, b, r, d:
 - Each vector is one value per stencil point
 - Divide stencil points among processors, **n/p** points each
- How do you divide up the **\sqrt{n}** by **\sqrt{n}** region of points?
- Block row (or block col) layout: **$v = 2 * p * \sqrt{n}$**
- 2-dimensional block layout: **$v = 4 * \sqrt{p} * \sqrt{n}$**

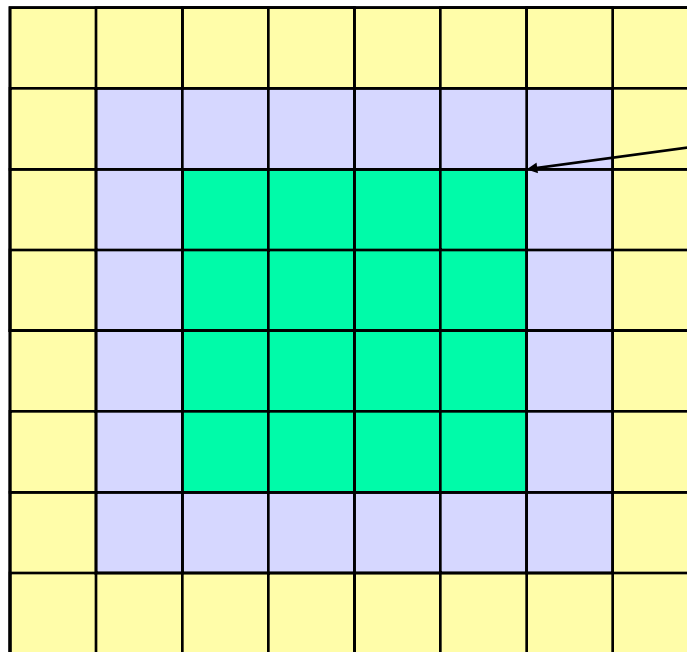
Detailed complexity measures for data movement I: Latency/Bandwidth Model

Moving data between processors by message-passing

- Machine parameters:
 - α latency (message startup time in seconds)
 - β inverse bandwidth (in seconds per word)
 - between nodes of Triton, $\alpha \sim 2.2 \times 10^{-6}$ and $\beta \sim 6.4 \times 10^{-9}$
- Time to send & recv or bcast a message of w words: $\alpha + w\beta$
- t_{comm} total communication time
- t_{comp} total computation time
- Total parallel time: $t_p = t_{\text{comp}} + t_{\text{comm}}$

Ghost Nodes in Stencil Computations

Comm cost = α * (#messages) + β * (total size of messages)



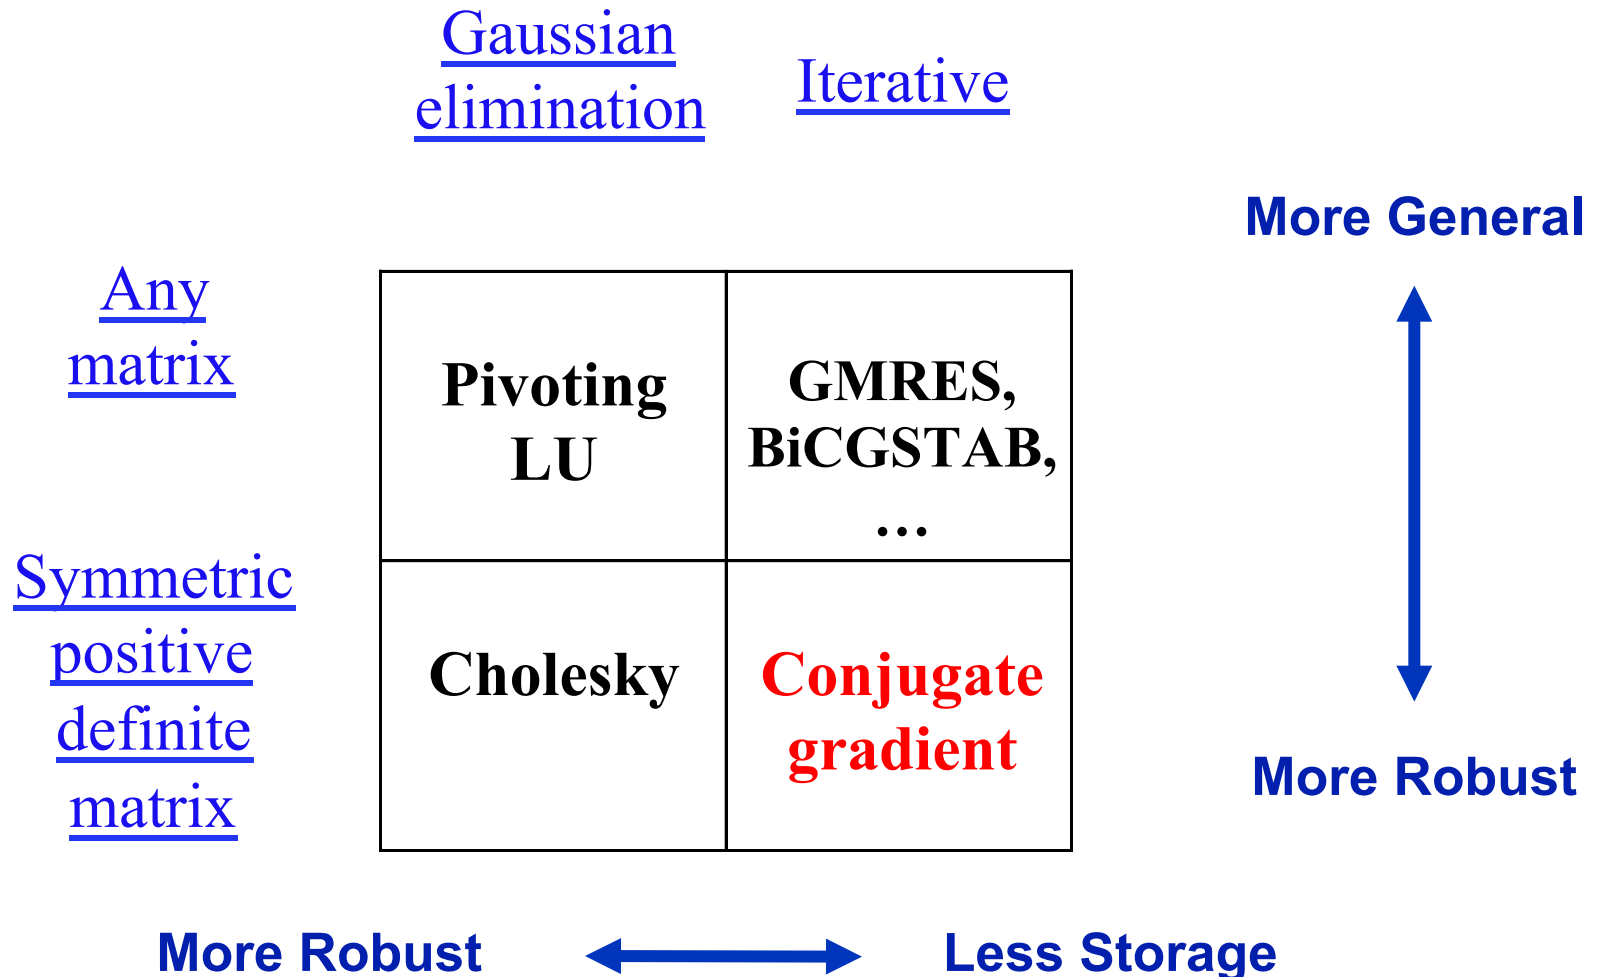
Green = my interior nodes

Blue = my boundary nodes

Yellow
= neighbors' boundary nodes
= my "ghost nodes"

- Keep a ghost copy of neighbors' boundary nodes
- Communicate **every second iteration**, not every iteration
- Reduces #messages, **not** total size of messages
- Costs extra memory and computation
- Can also use more than one layer of ghost nodes

The Landscape of $Ax = b$ Algorithms



Conjugate gradient in general

- CG can be used to solve *any* system $Ax = b$, if ...

Conjugate gradient in general

- CG can be used to solve *any* system $Ax = b$, if ...
- The matrix A is *symmetric* ($a_{ij} = a_{ji}$) ...
- ... and *positive definite* (all eigenvalues > 0).

Conjugate gradient in general

- CG can be used to solve *any* system $Ax = b$, if ...
- The matrix A is *symmetric* ($a_{ij} = a_{ji}$) ...
- ... and *positive definite* (all eigenvalues > 0).
- Symmetric positive definite matrices occur a lot in scientific computing & data analysis!

Conjugate gradient in general

- CG can be used to solve *any* system $Ax = b$, if ...
- The matrix A is *symmetric* ($a_{ij} = a_{ji}$) ...
- ... and *positive definite* (all eigenvalues > 0).
- Symmetric positive definite matrices occur a lot in scientific computing & data analysis!
- But usually the matrix isn't just a stencil.
- Now we do need to store the matrix A . Where's the data?

Conjugate gradient in general

- CG can be used to solve *any* system $Ax = b$, if ...
- The matrix A is *symmetric* ($a_{ij} = a_{ji}$) ...
- ... and *positive definite* (all eigenvalues > 0).
- Symmetric positive definite matrices occur a lot in scientific computing & data analysis!
- But usually the matrix isn't just a stencil.
- Now we do need to store the matrix A . Where's the data?
- The key is to use graph data structures and algorithms.

Vector and matrix primitives for CG

- **DAXPY:** $v = \alpha * v + \beta * w$ (vectors v, w ; scalars α, β)
 - **Broadcast** the scalars α and β , then independent $*$ and $+$
 - **comm volume** = $2p$, **span** = $\log n$
- **DDOT:** $\alpha = v^T * w = \sum_j v[j] * w[j]$ (vectors v, w ; scalar α)
 - Independent $*$, then $+$ **reduction**
 - **comm volume** = p , **span** = $\log n$
- **Matvec:** $v = A * w$ (matrix A , vectors v, w)
 - The hard part
 - But all you need is a subroutine to compute v from w
 - Sometimes you don't need to store A (e.g. temperature problem)
 - Usually you do need to store A , but it's *sparse* ...

Conjugate gradient: Krylov subspaces

- Eigenvalues: $Av = \lambda v$ $\{ \lambda_1, \lambda_2, \dots, \lambda_n \}$

- Cayley-Hamilton theorem:

$$(A - \lambda_1 I) \cdot (A - \lambda_2 I) \cdot \dots \cdot (A - \lambda_n I) = 0$$

$$\text{Therefore } \sum_{0 \leq i \leq n} c_i A^i = 0 \text{ for some } c_i$$

$$\text{so } A^{-1} = \sum_{1 \leq i \leq n} (-c_i/c_0) A^{i-1}$$

- Krylov subspace:

$$\text{Therefore if } Ax = b, \text{ then } x = A^{-1} b \text{ and} \\ x \in \text{span} (b, Ab, A^2b, \dots, A^{n-1}b) = K_n(A, b)$$

Conjugate gradient: Orthogonal sequences

- Krylov subspace: $K_i(A, b) = \text{span}(b, Ab, A^2b, \dots, A^{i-1}b)$
- Conjugate gradient algorithm:
 - for $i = 1, 2, 3, \dots$
 - find $x_i \in K_i(A, b)$
 - such that $r_i = (b - Ax_i) \perp K_i(A, b)$
- Notice $r_i \in K_{i+1}(A, b)$, so $r_i \perp r_j$ for all $j < i$
- Similarly, the “directions” are A -orthogonal:
$$(x_i - x_{i-1})^T \cdot A \cdot (x_j - x_{j-1}) = 0$$
- The magic: Short recurrences. . .
 - A is symmetric \Rightarrow can get next residual and direction from the previous one, without saving them all.

Conjugate gradient: Convergence

- In exact arithmetic, CG converges in n steps
(completely unrealistic!!)
- Accuracy after k steps of CG is related to:
 - consider polynomials of degree k that are equal to 1 at 0.
 - how small can such a polynomial be at all the eigenvalues of A ?
- Thus, eigenvalues close together are good.
- Condition number: $\kappa(A) = \|A\|_2 \|A^{-1}\|_2 = \lambda_{\max}(A) / \lambda_{\min}(A)$
- Residual is reduced by a constant factor by $O(\kappa^{1/2}(A))$ iterations of CG.

Other Krylov subspace methods

- Nonsymmetric linear systems:
 - GMRES:
for $i = 1, 2, 3, \dots$
find $x_i \in K_i(A, b)$ such that $r_i = (Ax_i - b) \perp K_i(A, b)$
But, no short recurrence \Rightarrow save old vectors \Rightarrow lots more space
(Usually “restarted” every k iterations to use less space.)
 - BiCGStab, QMR, etc.:
Two spaces $K_i(A, b)$ and $K_i(A^T, b)$ w/ mutually orthogonal bases
Short recurrences $\Rightarrow O(n)$ space, but less robust
 - Convergence and preconditioning more delicate than CG
 - Active area of current research
- Eigenvalues: Lanczos (symmetric), Arnoldi (nonsymmetric)

Conjugate gradient iteration

$$x_0 = 0, \quad r_0 = b, \quad d_0 = r_0$$

for $k = 1, 2, 3, \dots$

$$\alpha_k = (r_{k-1}^T r_{k-1}) / (d_{k-1}^T A d_{k-1}) \quad \text{step length}$$

$$x_k = x_{k-1} + \alpha_k d_{k-1} \quad \text{approx solution}$$

$$r_k = r_{k-1} - \alpha_k A d_{k-1} \quad \text{residual}$$

$$\beta_k = (r_k^T r_k) / (r_{k-1}^T r_{k-1}) \quad \text{improvement}$$

$$d_k = r_k + \beta_k d_{k-1} \quad \text{search direction}$$

- One matrix-vector multiplication per iteration
- Two vector dot products per iteration
- Four n-vectors of working storage