Avoiding Communication in SpMV-like Computations

Marghoob Mohiyuddin Mark Hoemmen James Demmel
Katherine Yelick
marghoob@eecs.berkeley.edu

Department of Electrical Engineering and Computer Science
University of California at Berkeley

April 28, 2008

Outline

Communication is Costly, Computation is Cheap

- Gap between computational capability and communication cost increasing exponentially:
  - Floating-point time $\ll$ 1/network BW $\ll$ network latency
  - Floating-point time $\ll$ 1/memory BW $\ll$ memory latency
- Applications need to be designed with this gap in mind
  - Communication hiding not enough (speedup $\leq$ 2x)
    - Latency can be dealt with by overlap, but limited by the amount of computation
  - Communication avoiding:
    - Trade off communication with computation.
    - Arbitrary speedups possible

The Akx Kernel

- Given $n \times n$ sparse matrix A, vector x, integer $k > 0$,
- Compute the k vectors $Ax, A^2x, \ldots, A^kx$ efficiently.
  - Parallel and sequential algorithms.
- Arises in Krylov Subspace Methods.
  - Need to look at the linear subspace spanned by $[x, Ax, A^2x, \ldots, A^kx]$.

Setup

- Matrix A and vector x divided in to p row blocks.
- Parallel machine:
  - Each proc. operates on a separate block.
  - Interproc. communication for remote dependencies.
- Sequential machine:
  - Each block stored contiguously in slow memory.
  - Algorithm operates on a block-by-block basis.
- Output vectors computed on a per block basis.

Naïve Algorithm: 9-pt Operator for 3 Iterations

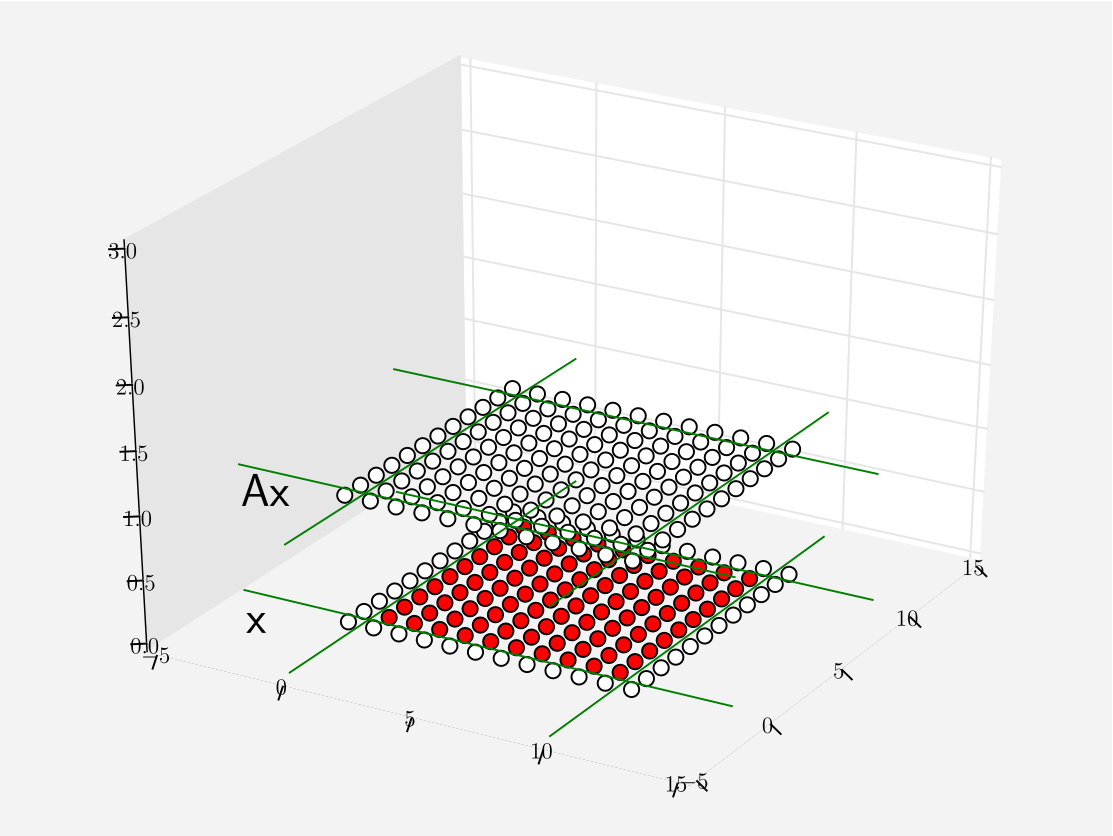

1. Elements of x arranged in a 2D mesh.
   - Matrix A defines the nearest neighbor connections
2. 9-point operator computes a function of an entry of x and its 8 neighbors

Naïve Algorithm: 9-pt Operator for 3 iterations

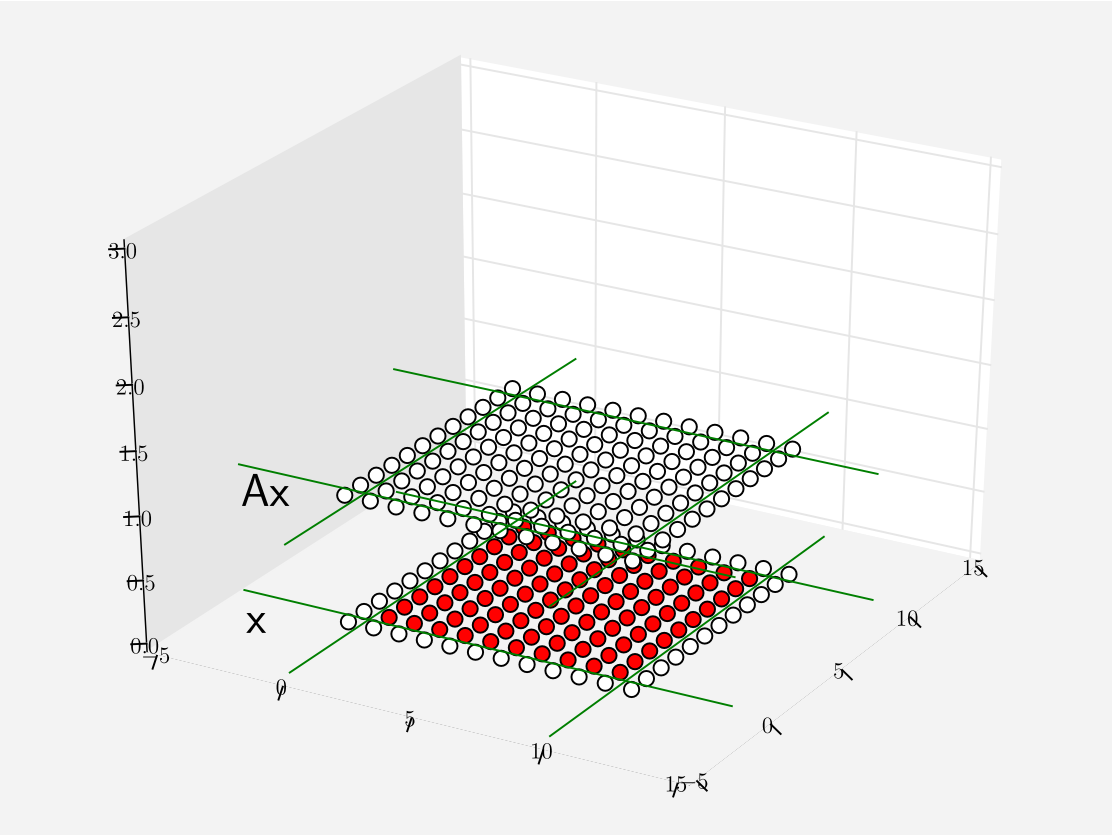

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

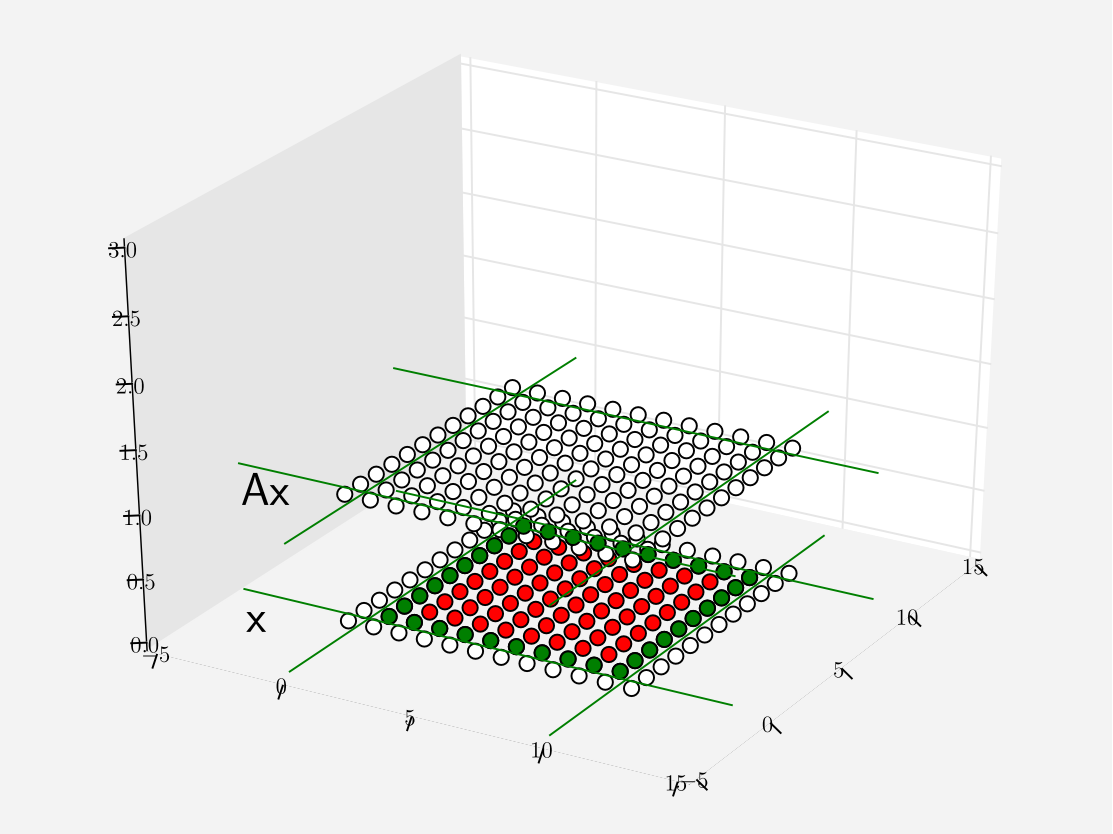

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

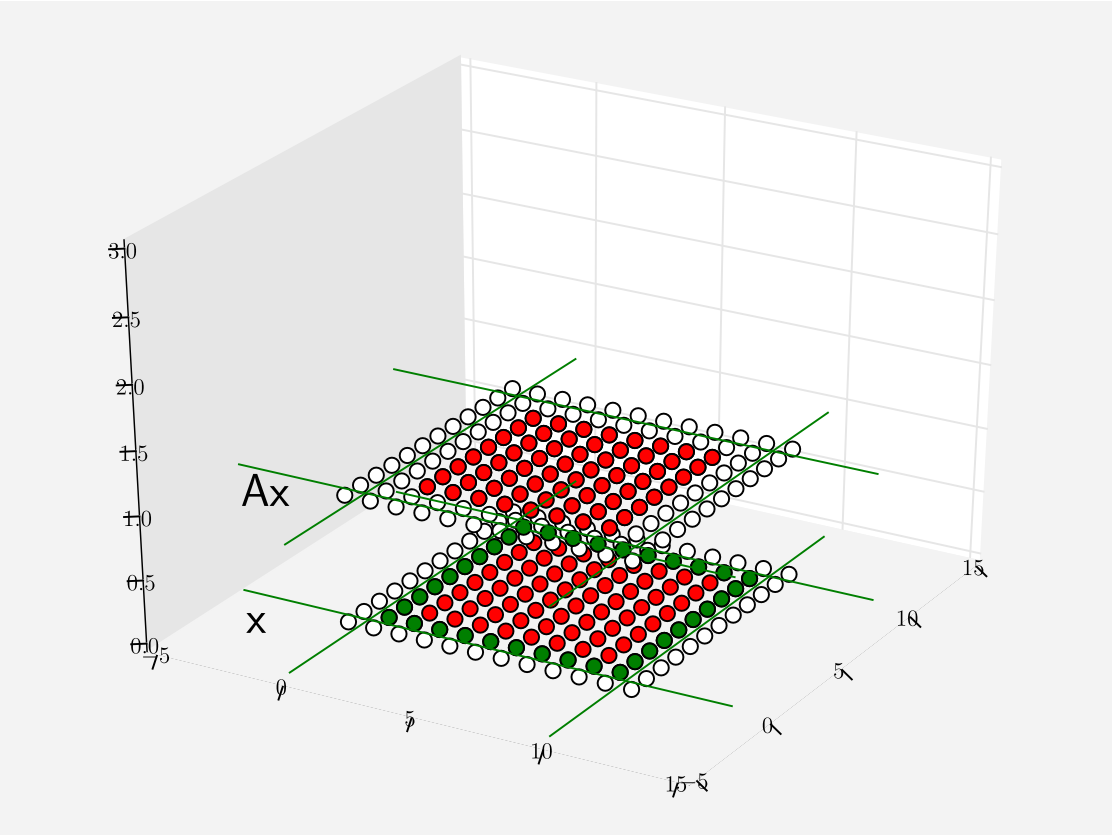

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

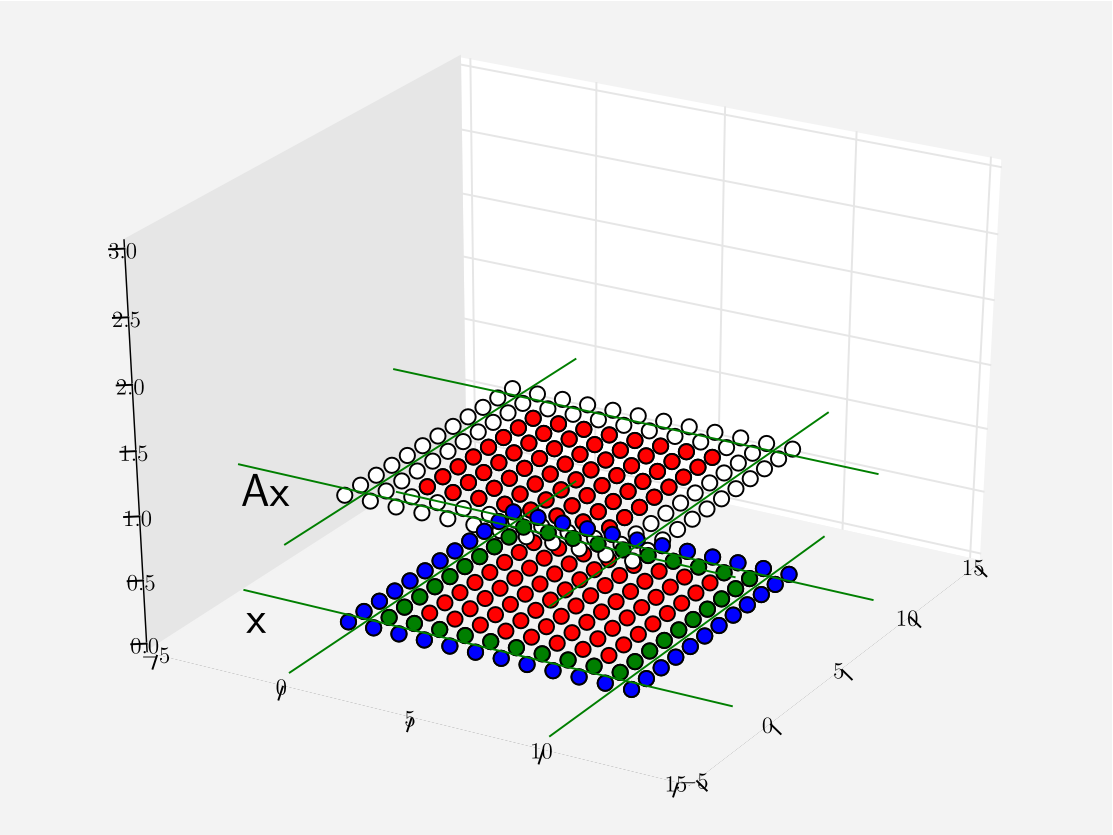

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

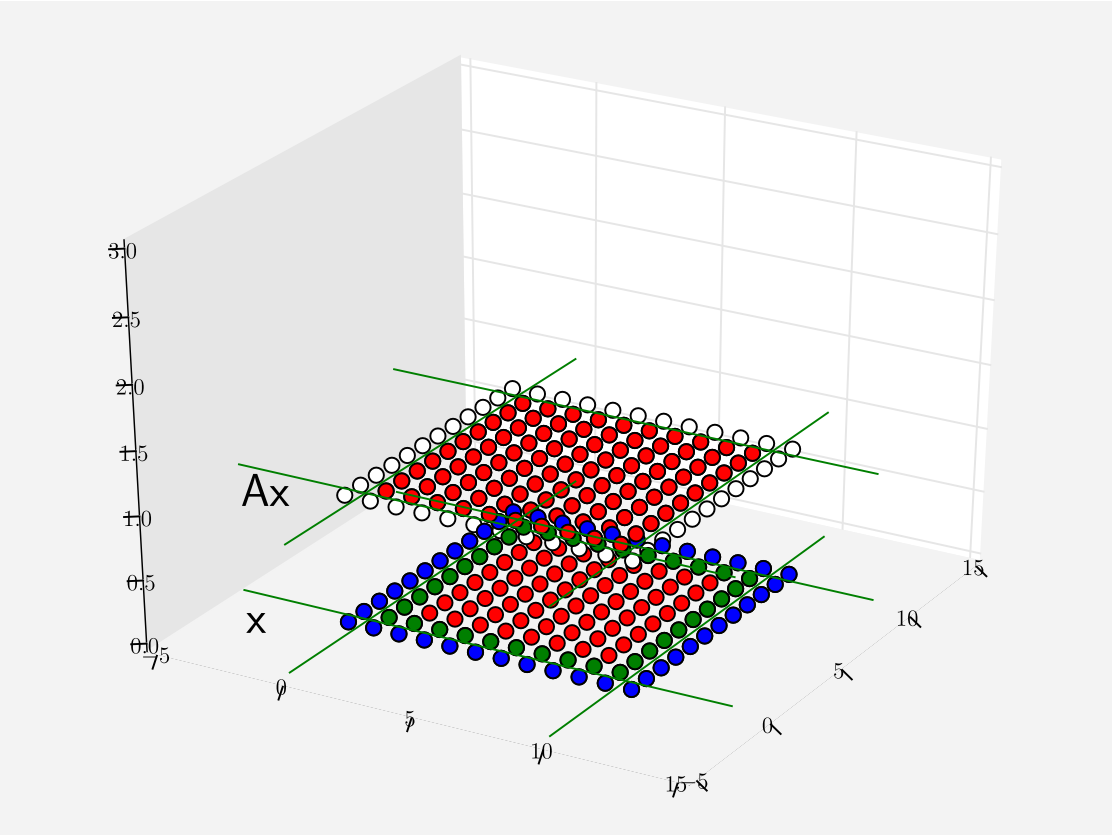

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

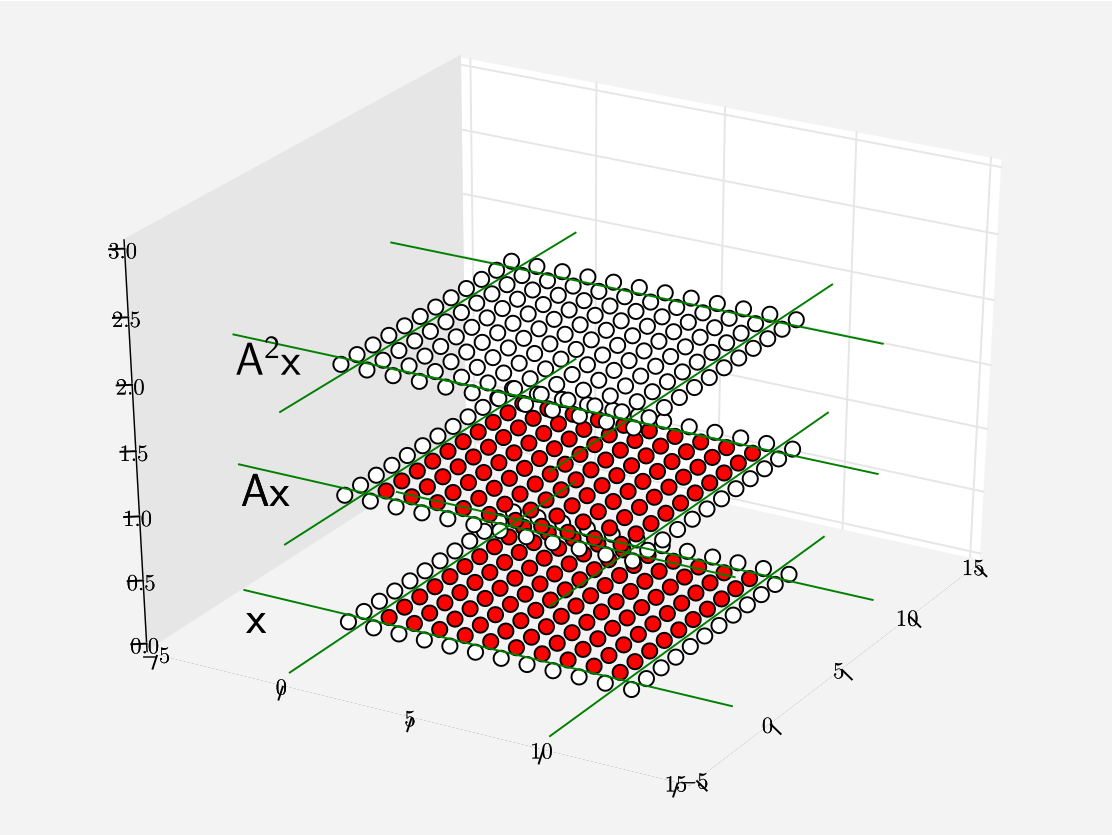

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

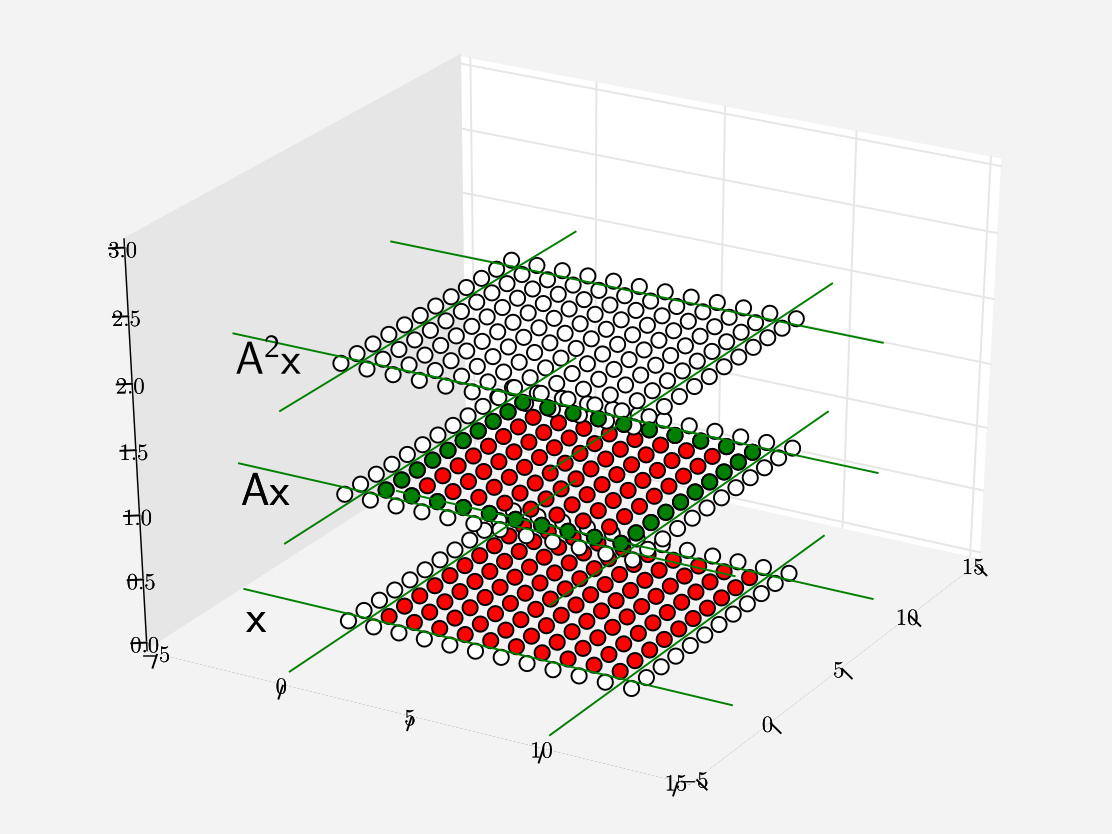

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

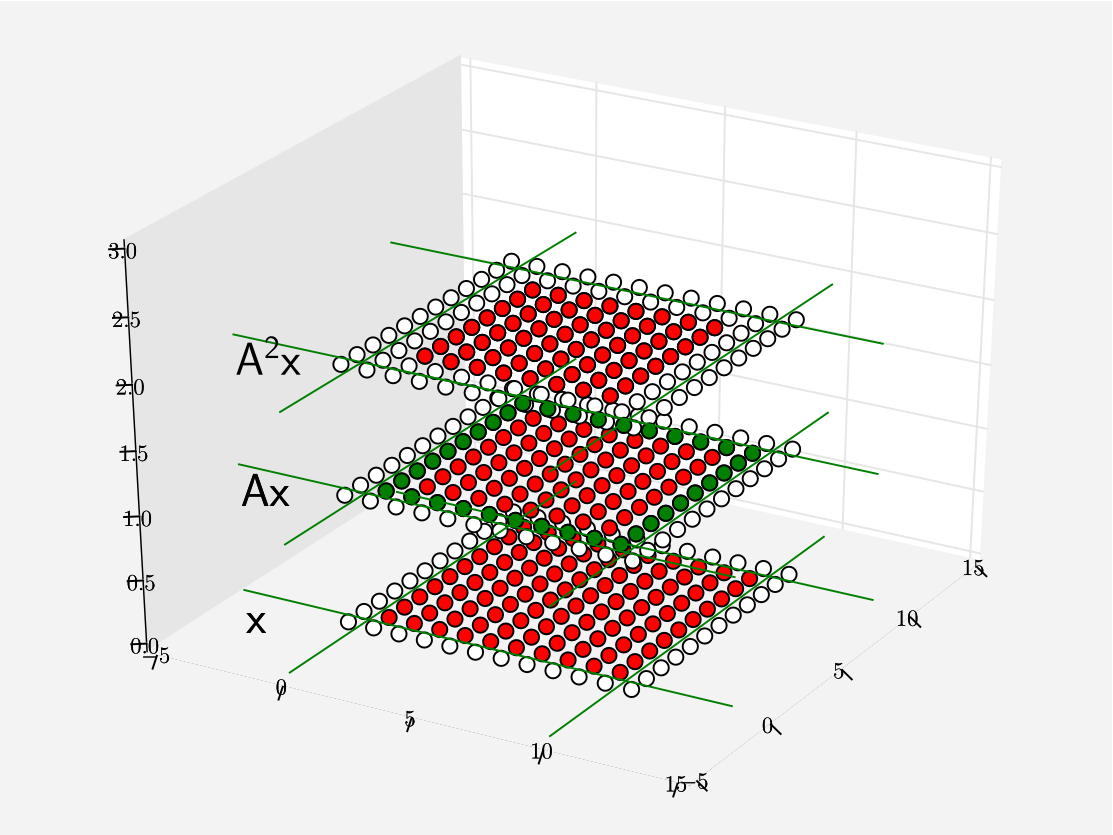

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

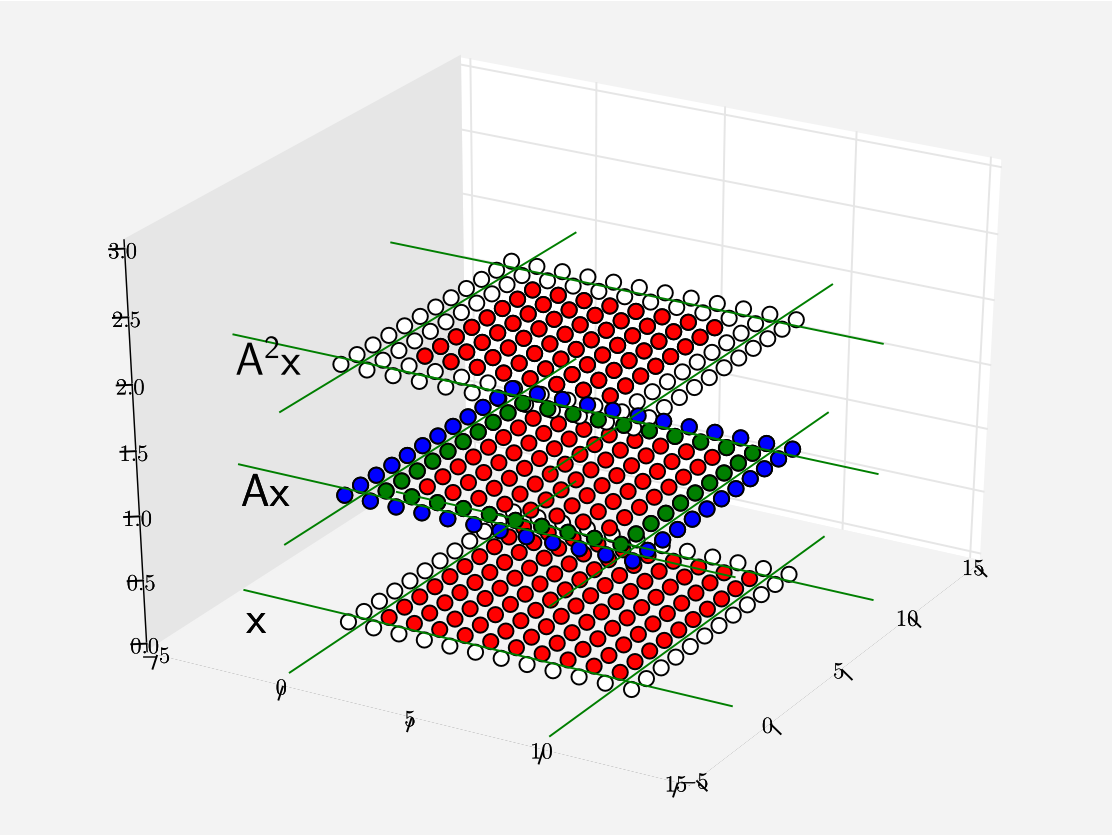

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

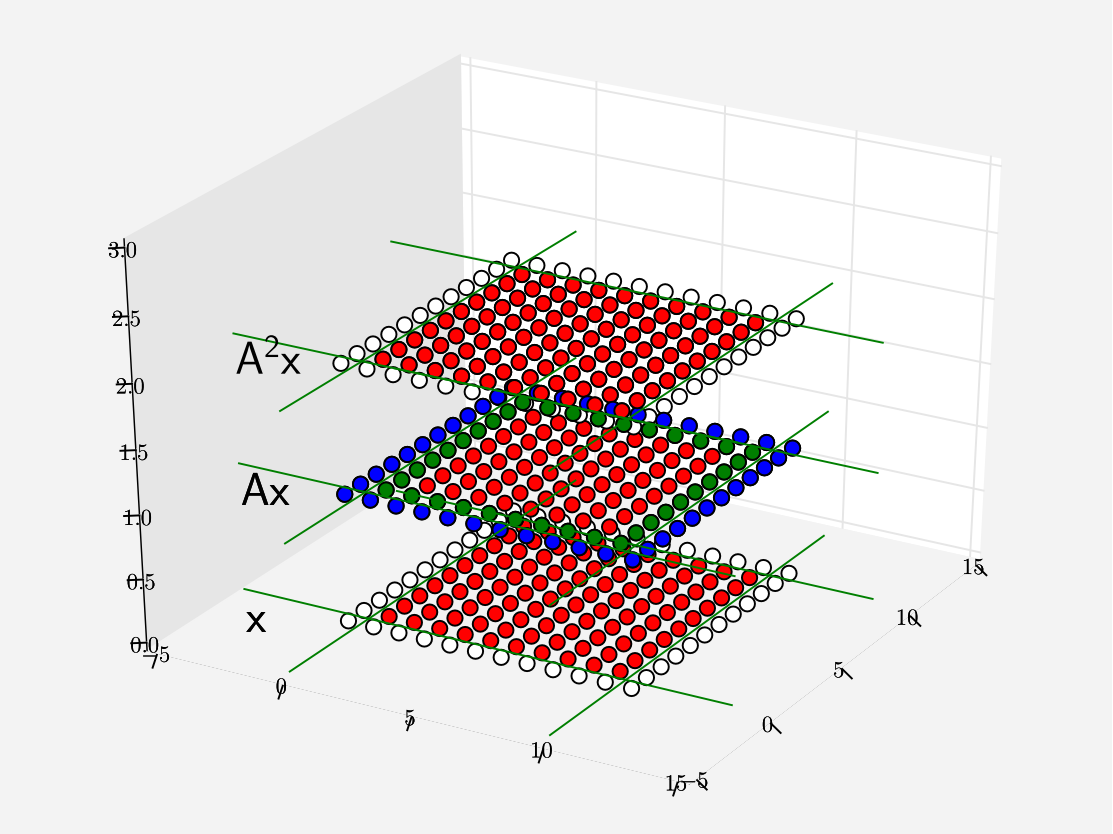

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

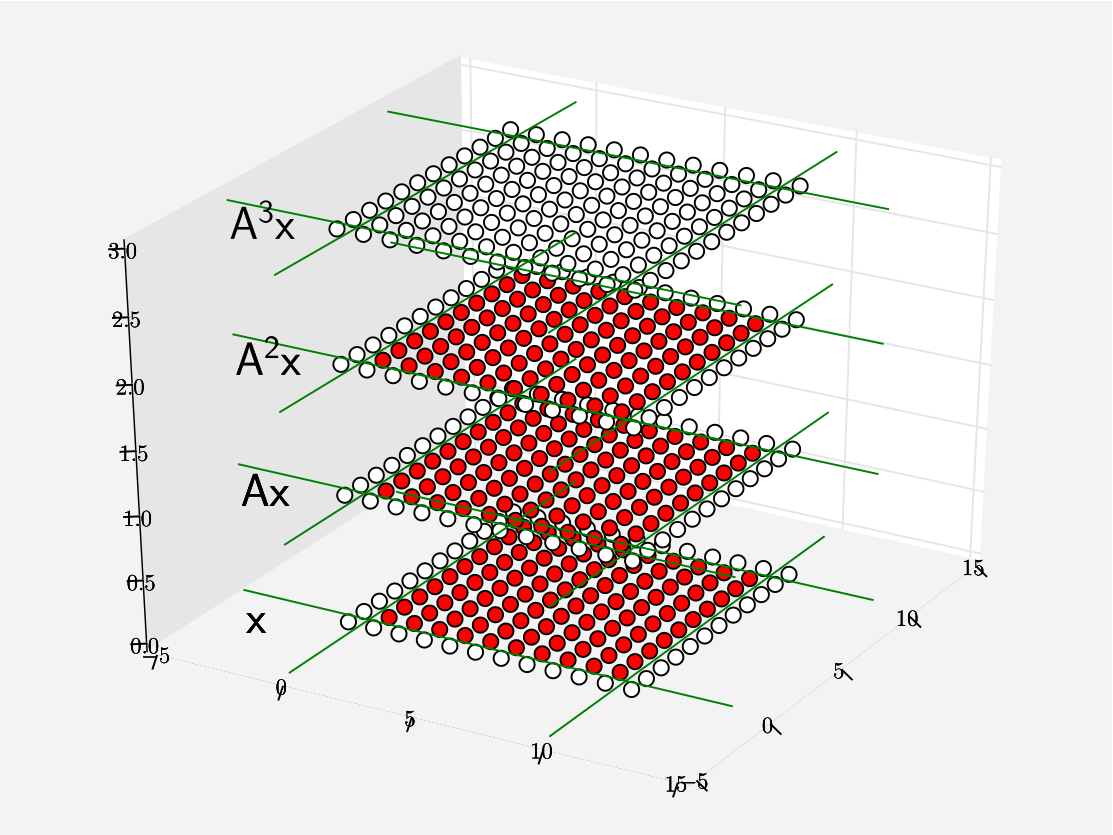

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

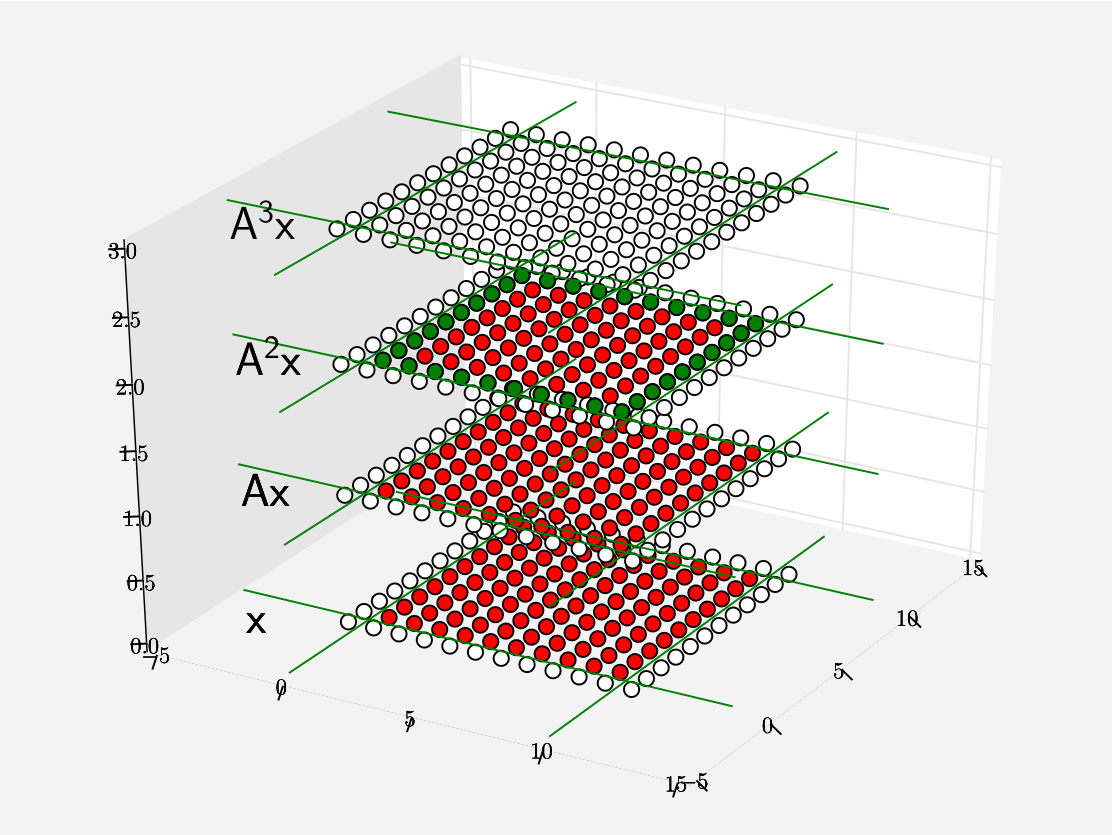

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

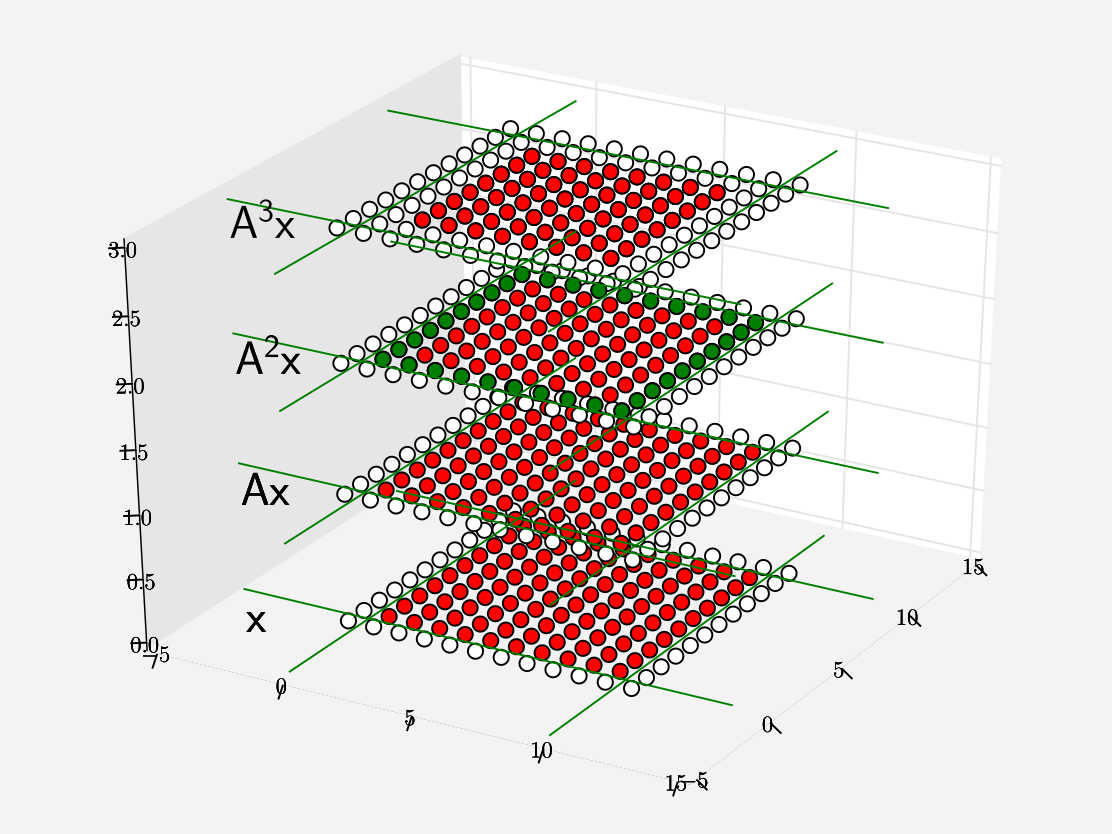

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

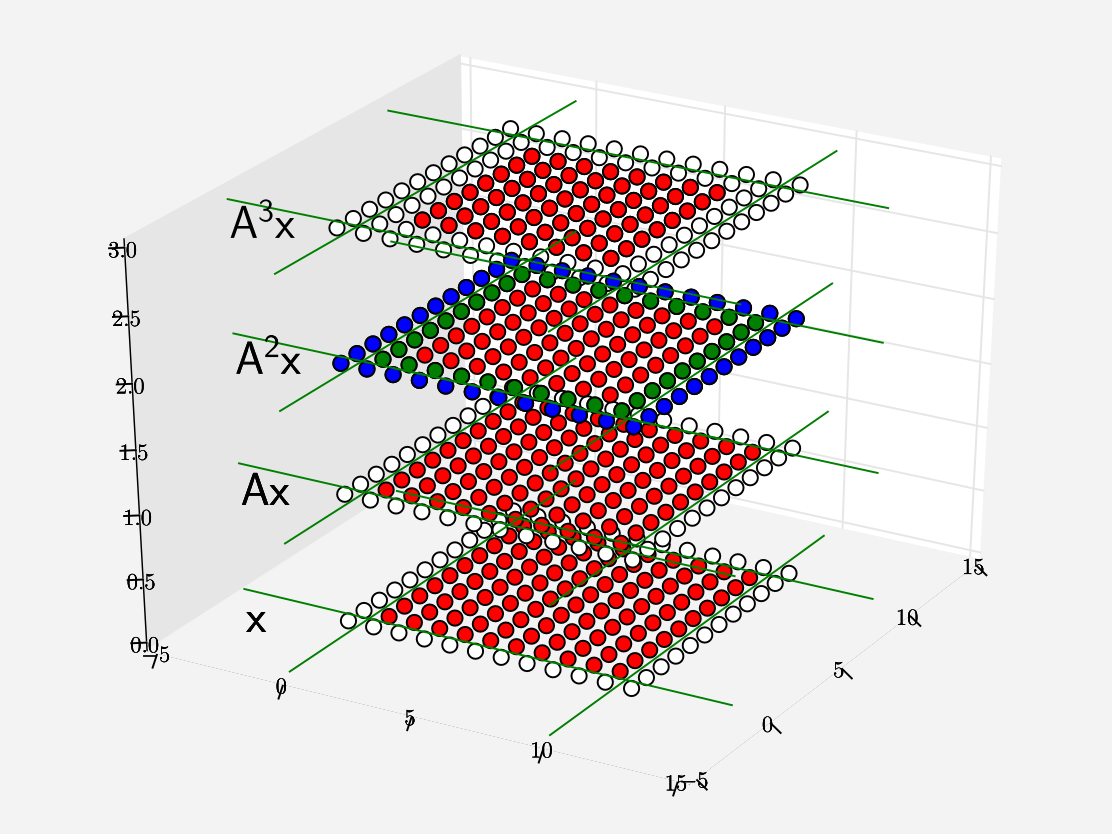

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

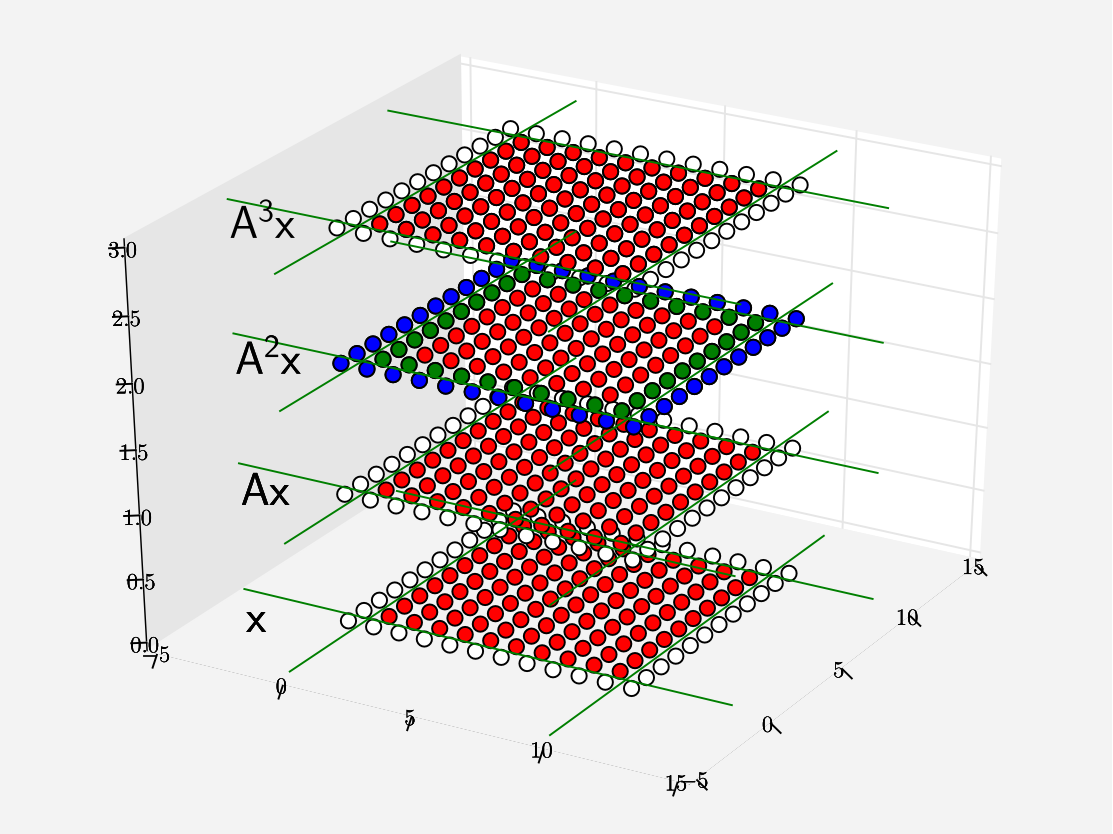

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm: 9-pt Operator for 3 iterations

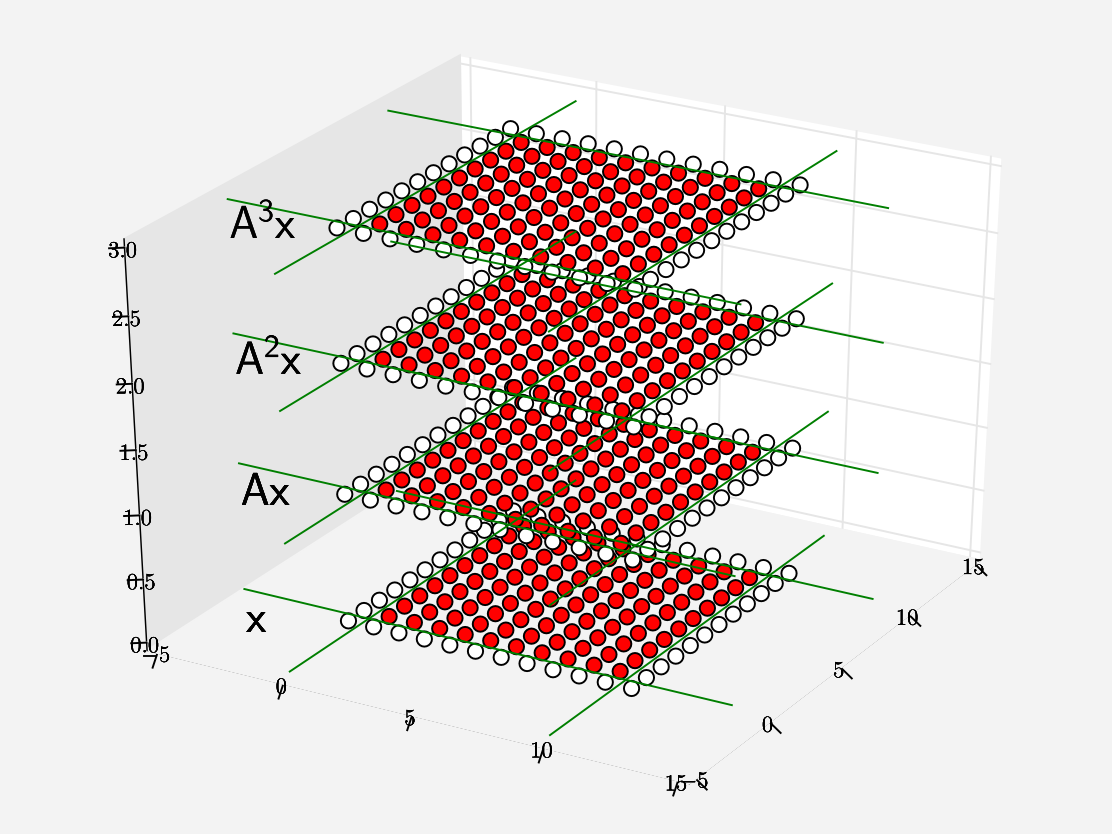

1. Entries available marked red.
2. Computing Ax: Send entries needed by other procs (green).
3. Computing Ax: Compute locally dependent entries.
4. Computing Ax: Receive entries from other procs (blue).
5. Computing Ax: Compute remaining entries of Ax.
6. Compute A^2x as $A(Ax)$.
7. Compute A^3x as $A(A^2x)$.

Naïve Algorithm

For $i = 1, \ldots, k$,

Compute $A^i x$ using the product of A and $A^{i-1}x$

- Fetch required entries of $A^{i-1}x$ from other procs/blocks.

- $O(k)$ messages between any two procs/blocks.
  - Latency cost is k times the minimum, namely O(1).
  - Objective: O(1) messages between any 2 blocks/procs.
- Sequential machine: A and x read k times from slow to fast memory.
  - Bandwidth cost is k times the minimum.
  - Objective: Read A and x at most once.
- Minimum number of floating-point operations performed.
  - Objective: Minimize number of extra flops.

Improving Naïve Algorithm: 9-point operator for 3 iterations

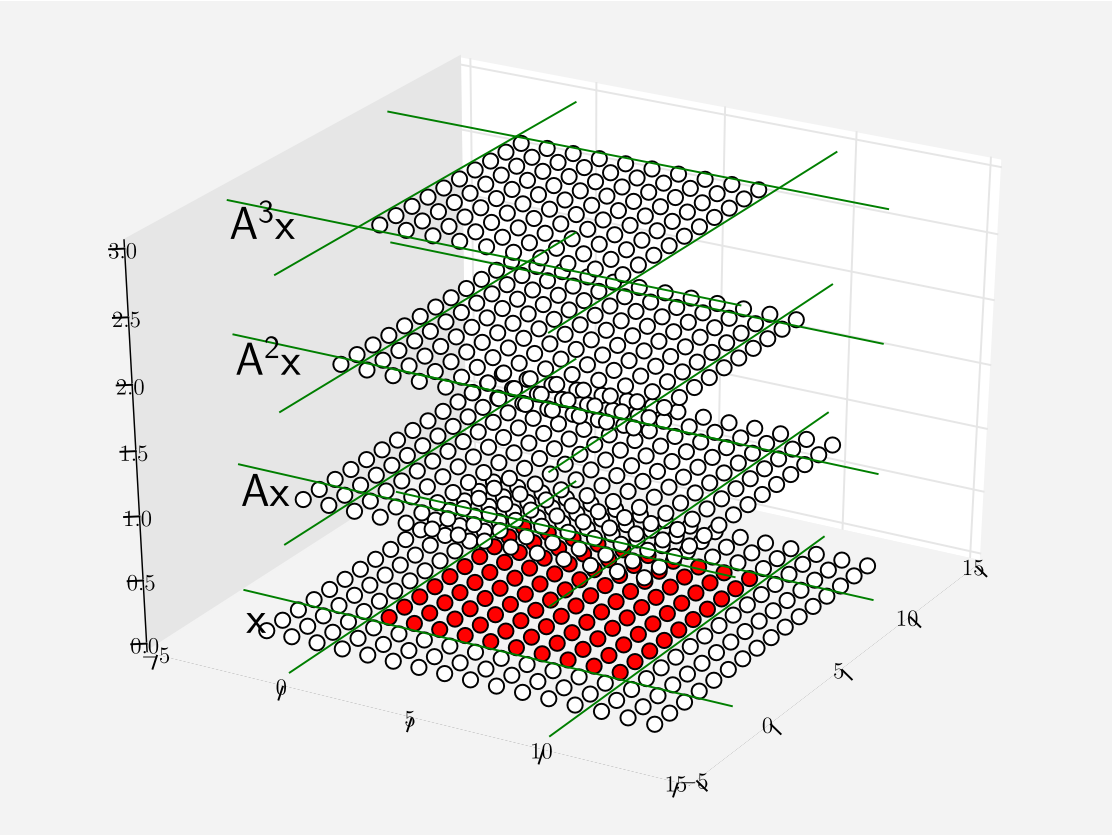

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

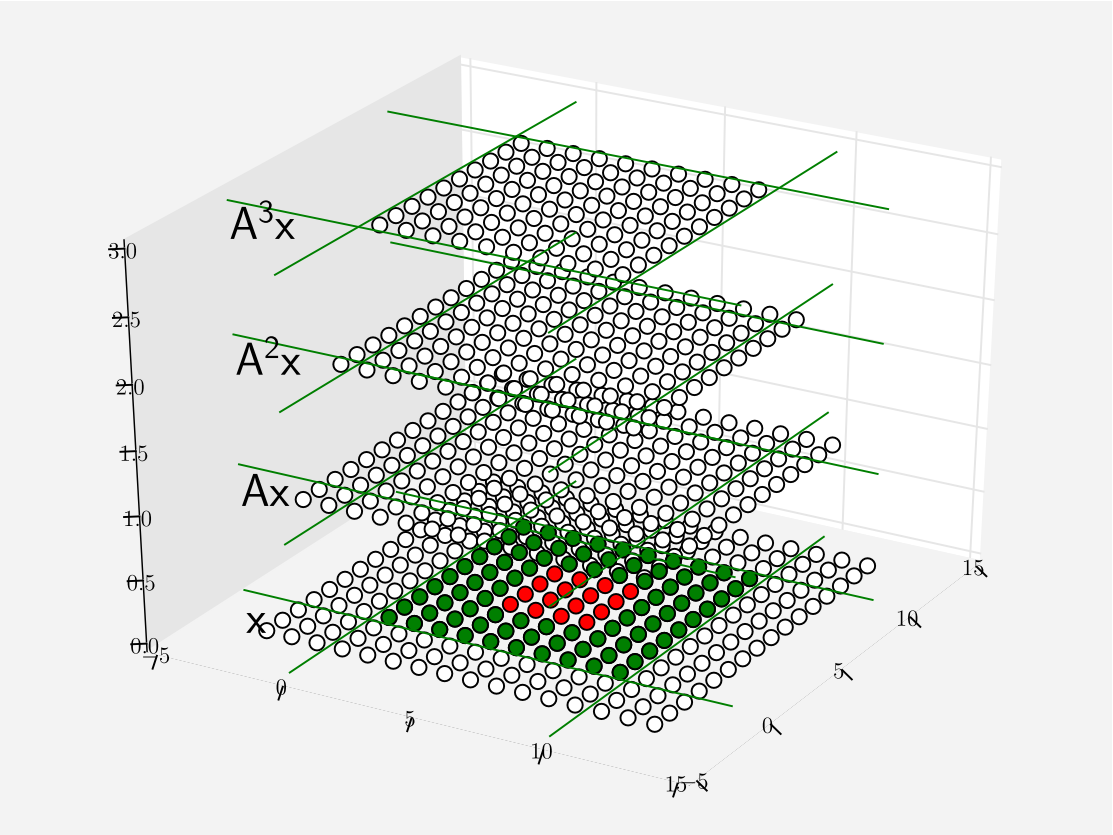

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

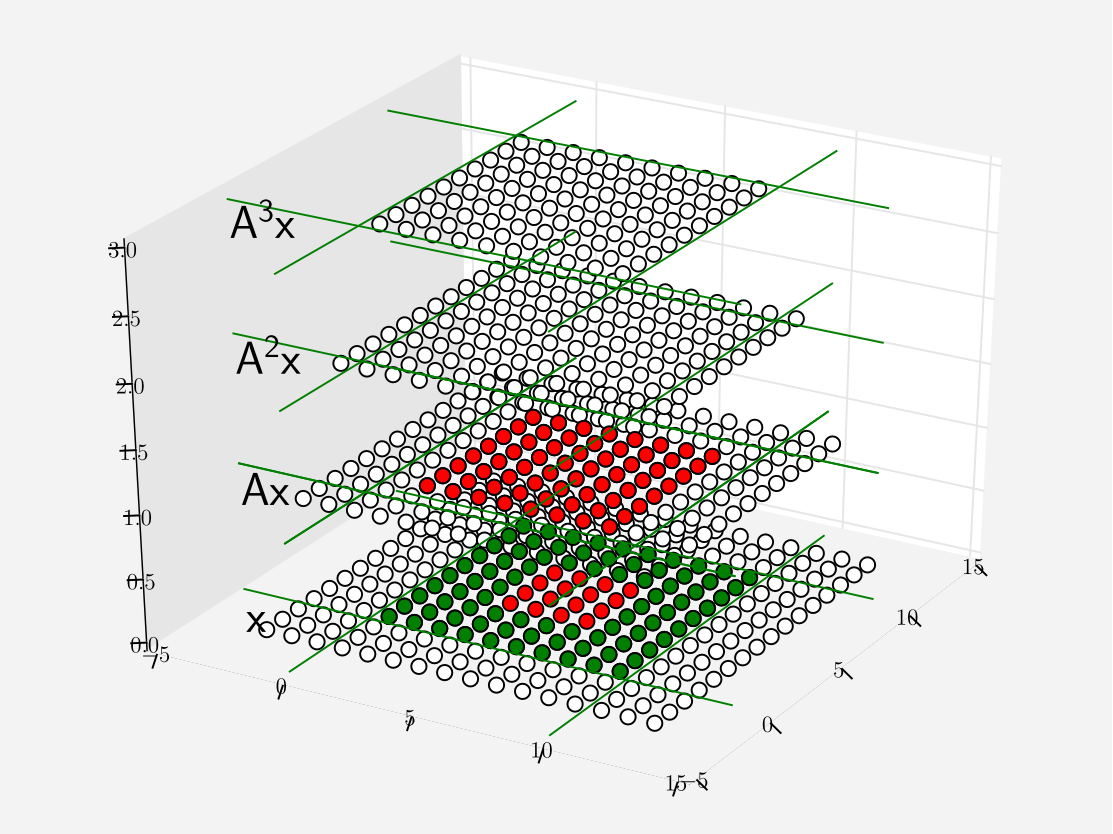

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

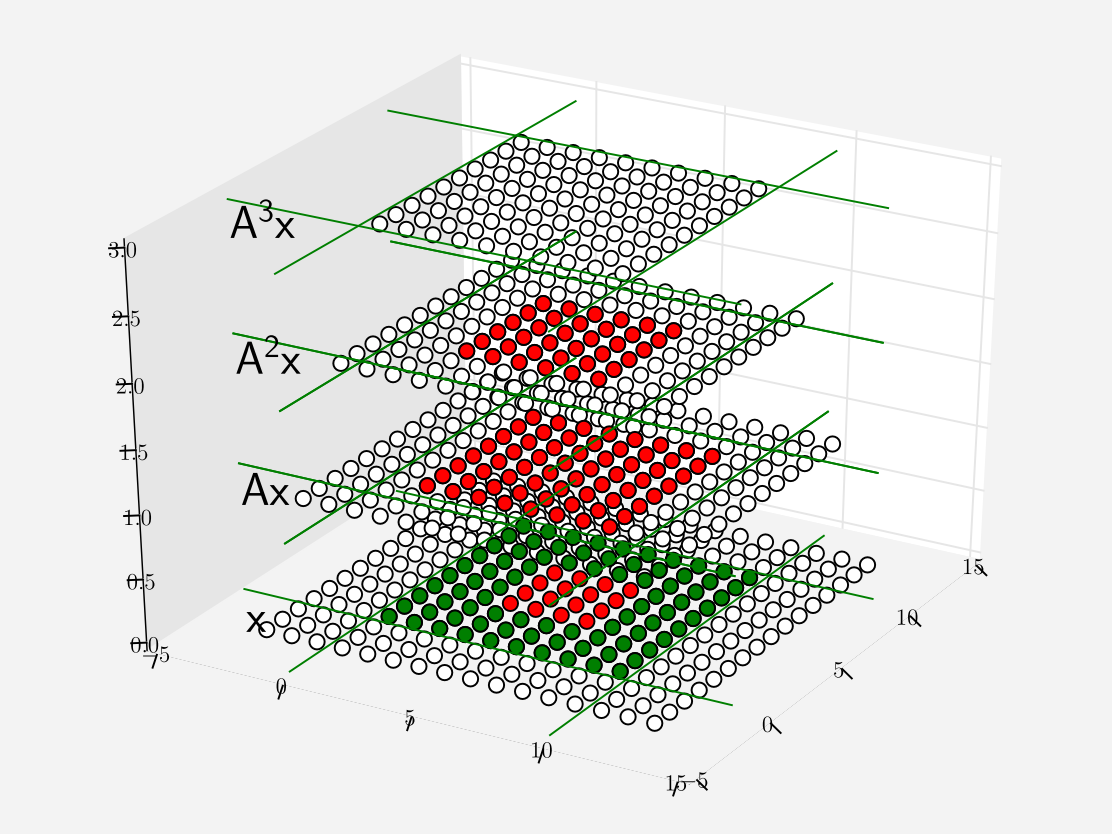

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

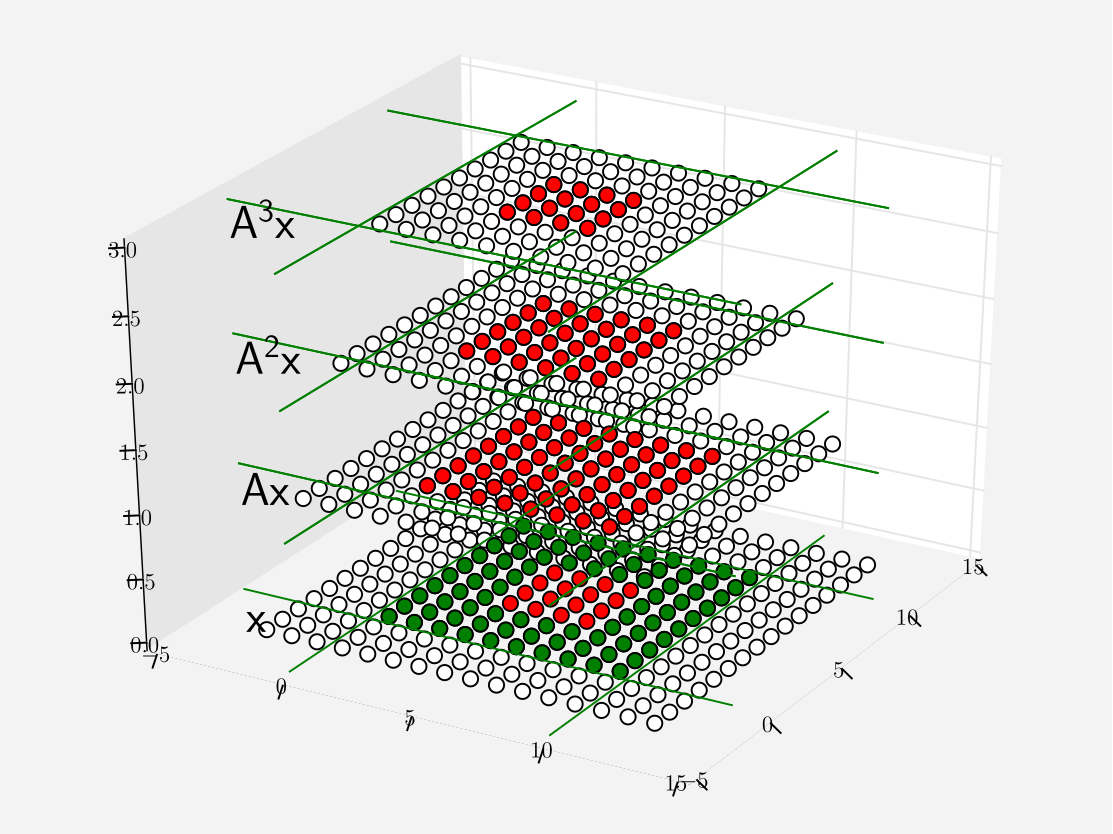

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

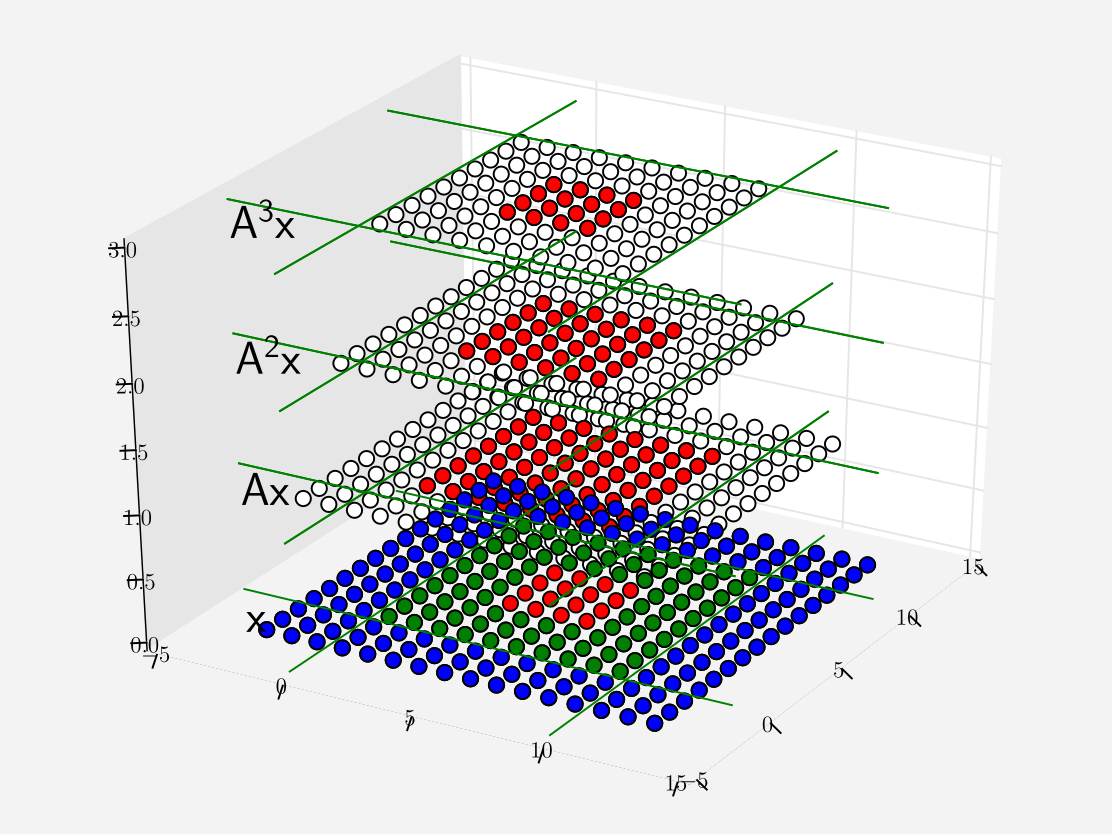

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

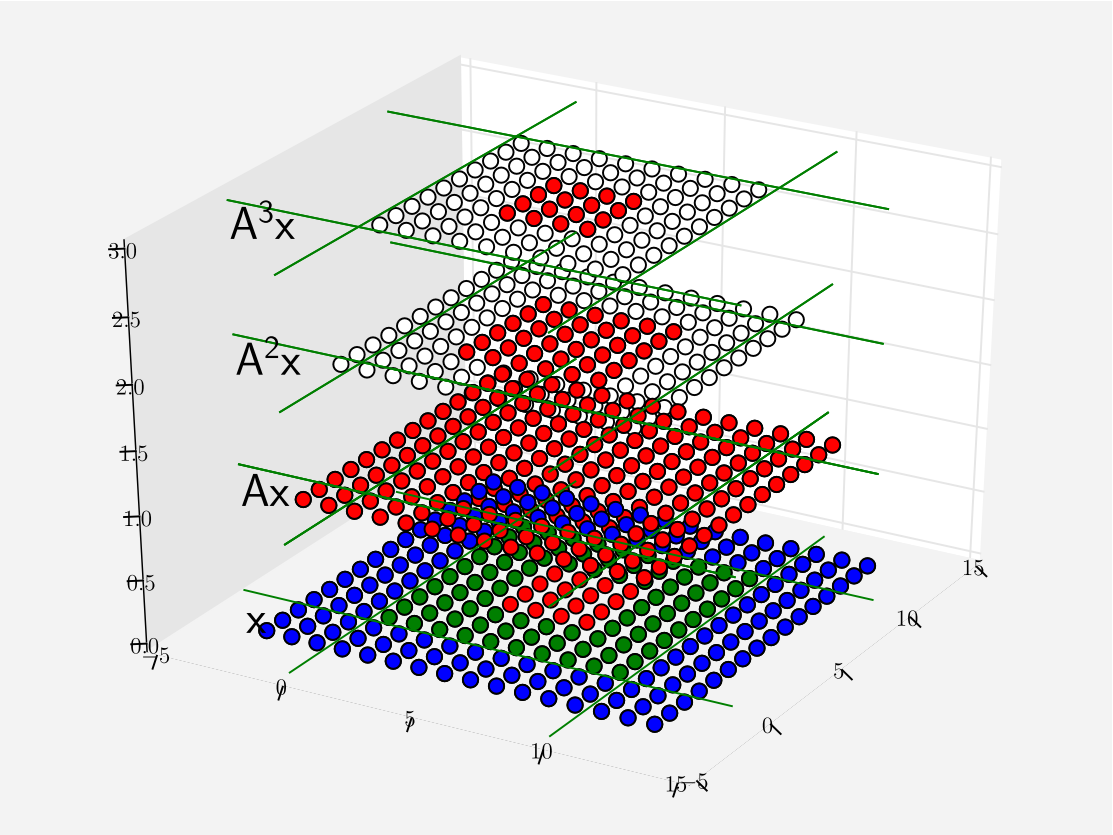

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

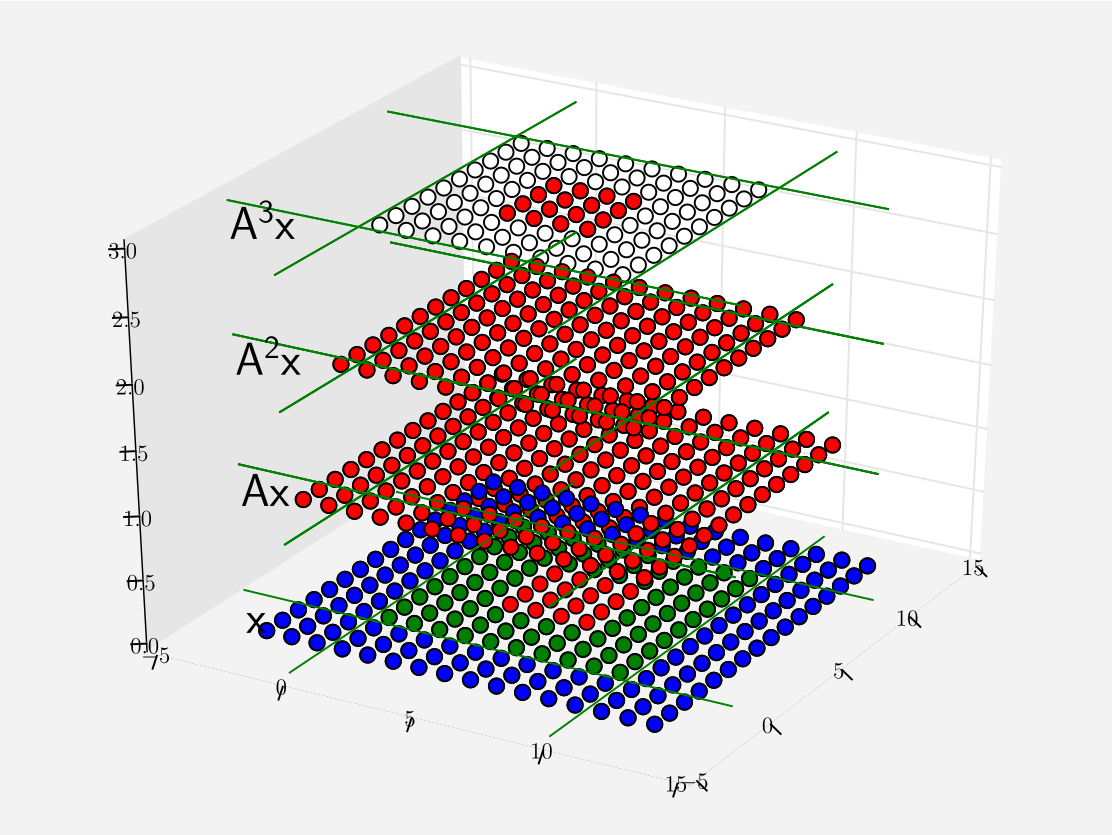

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

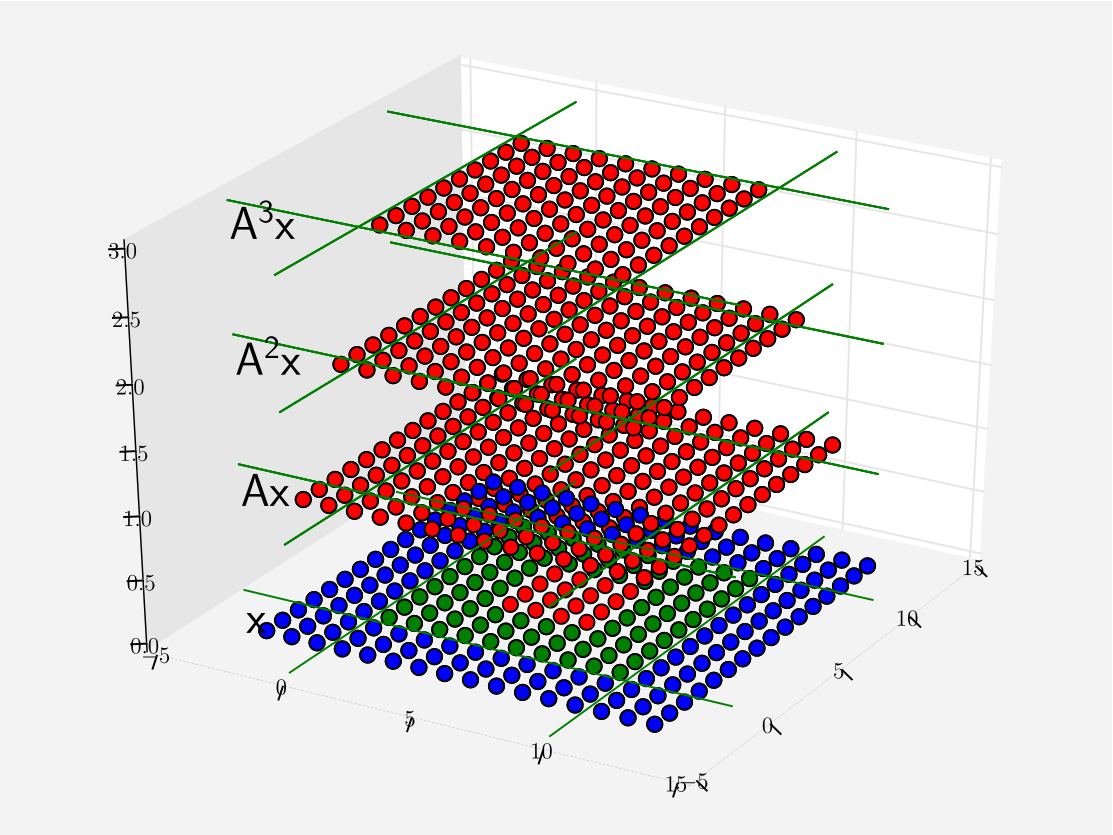

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Improving Naïve Algorithm: 9-point operator for 3 iterations

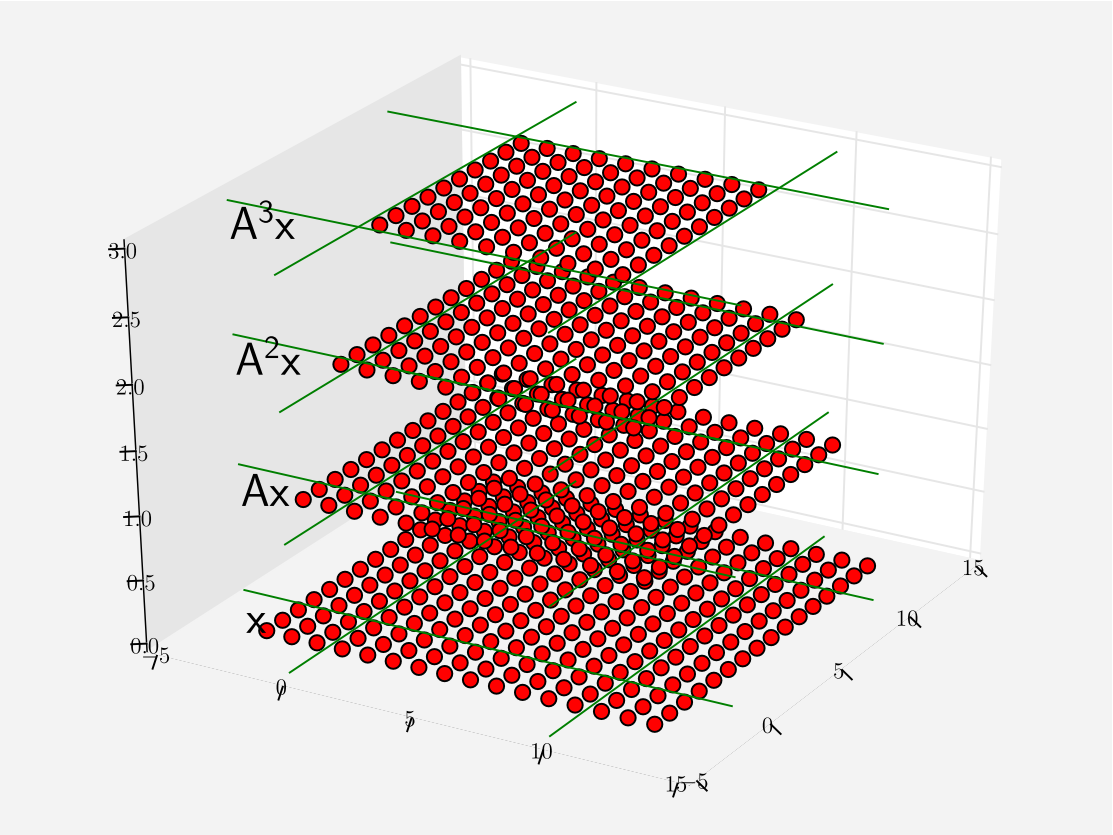

1. Entries available marked red.
2. Send entries needed by other procs.
3. Compute locally computable entries.
4. Receive entries from other procs (blue).
5. Compute remotely dependent entries.

Parallel Algorithm 1 (PA1)

1. The required entries of x are computed as dependencies on x of entries of $A^k x$ in block i.
2. Send the entries of x needed by other procs.
3. Compute the locally computable entries of $A^j x$ for $1 \leq j \leq k$.
4. Receive the entries of x needed from other procs.
5. Compute the remaining entries of $A^j x$ for $1 \leq j \leq k$.

- $O(1)$ messages between any 2 procs.
- Redundant computations.
- Applicable to arbitrary sparse matrices
  - Works well when the matrix is "well partitioned"
    $\Rightarrow$ partitions have small "surface to volume ratio"

Improving PA1: 9-pt operator for 3 iterations

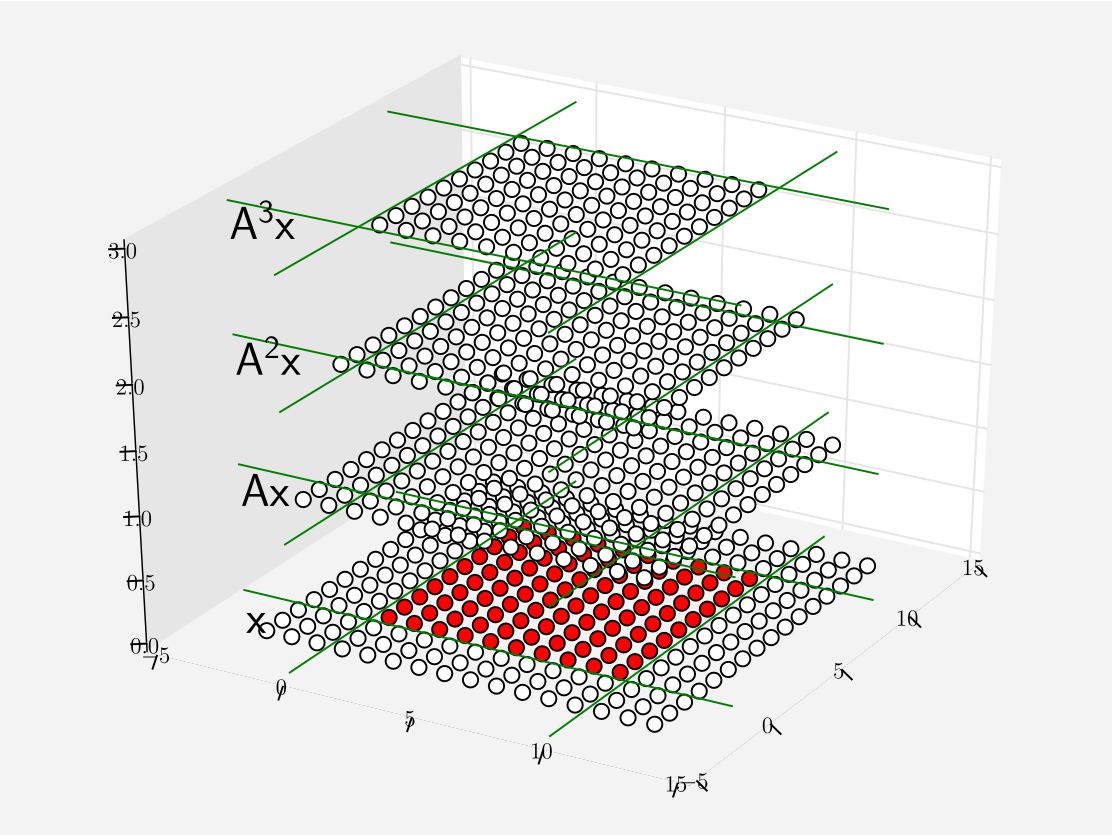

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

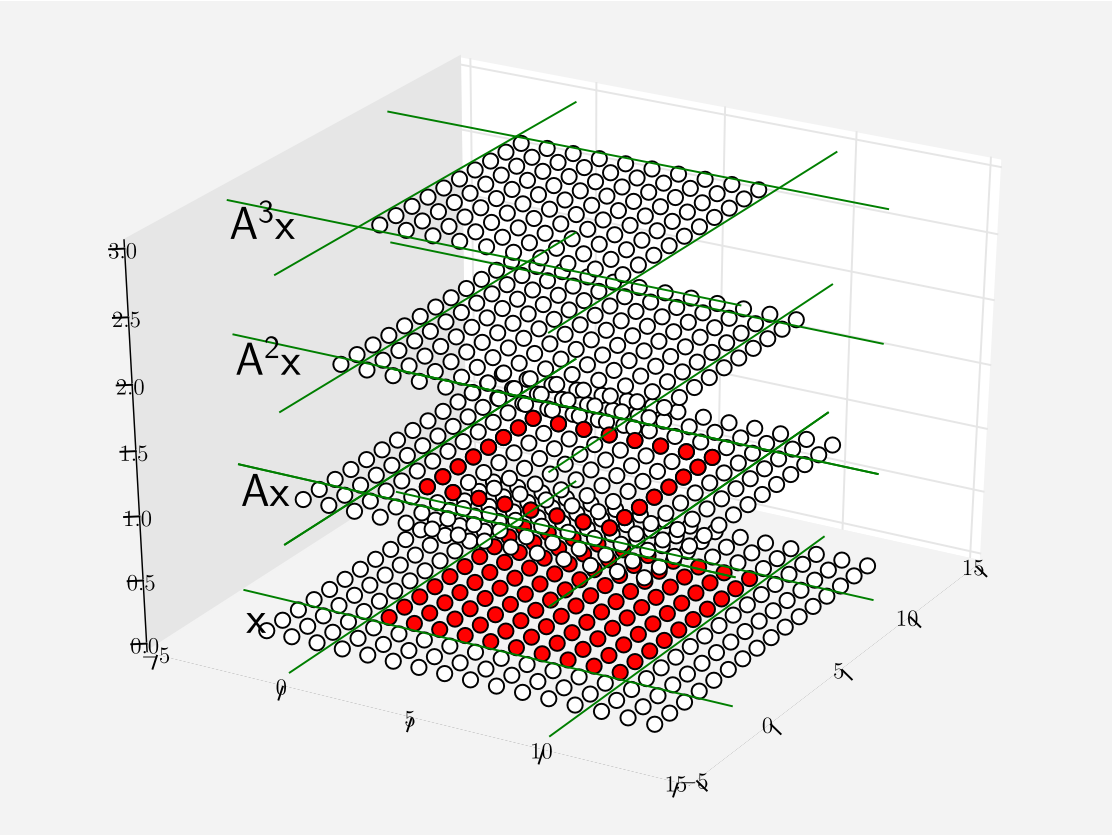

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

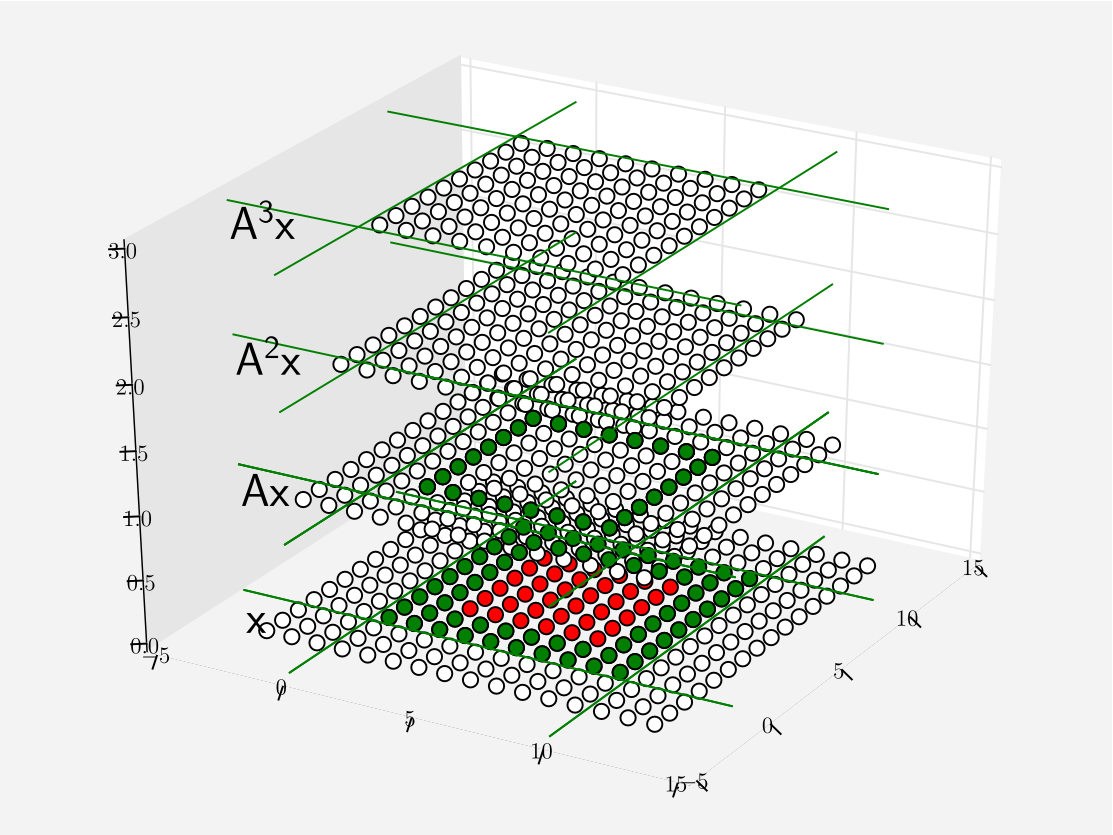

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

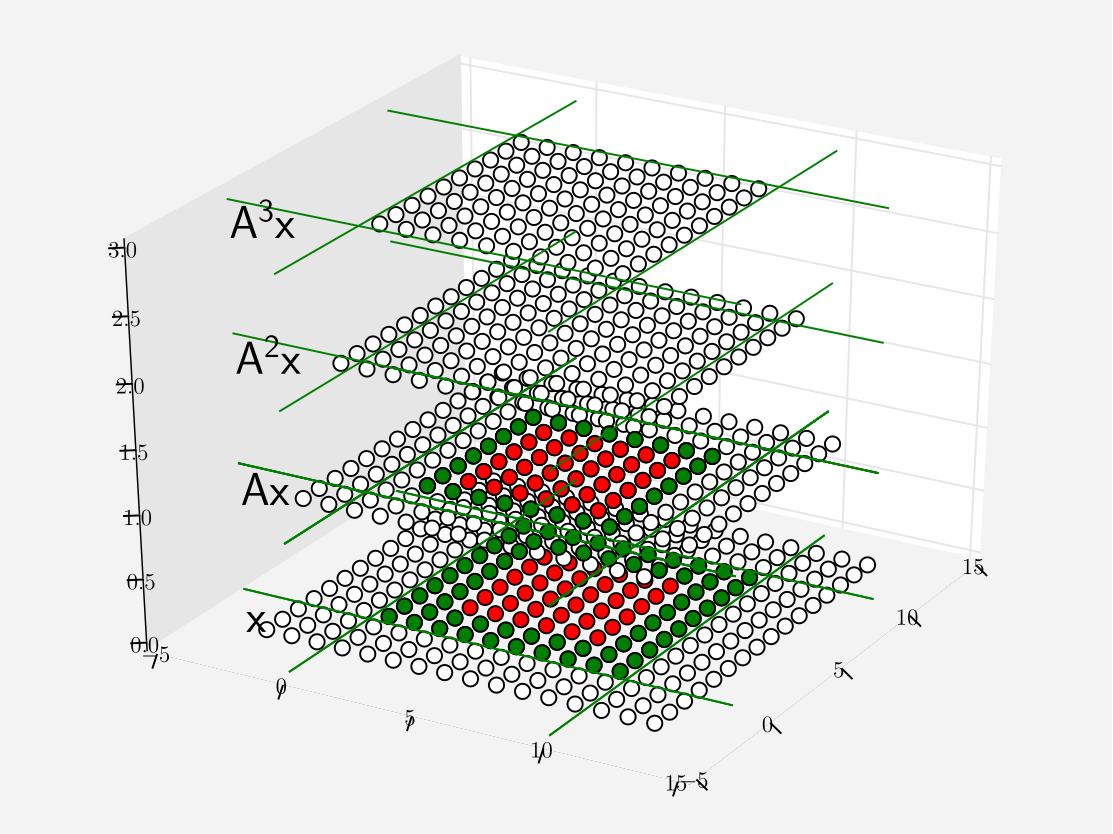

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

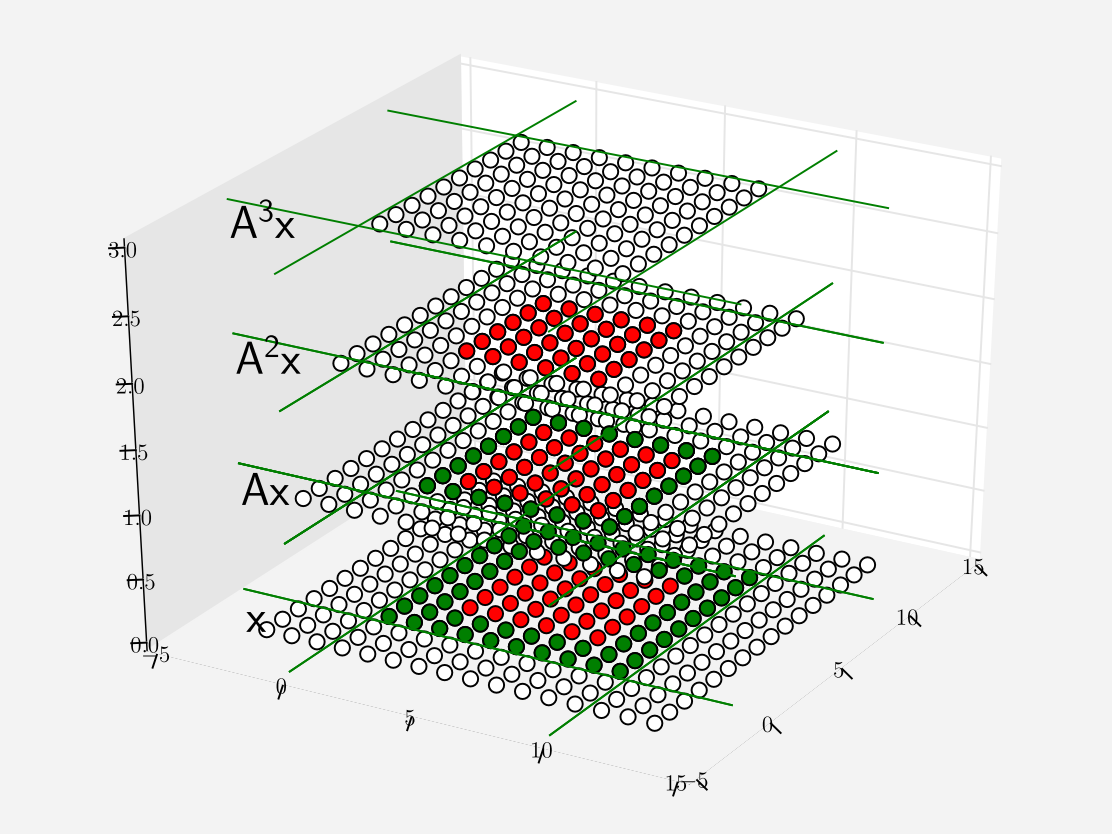

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

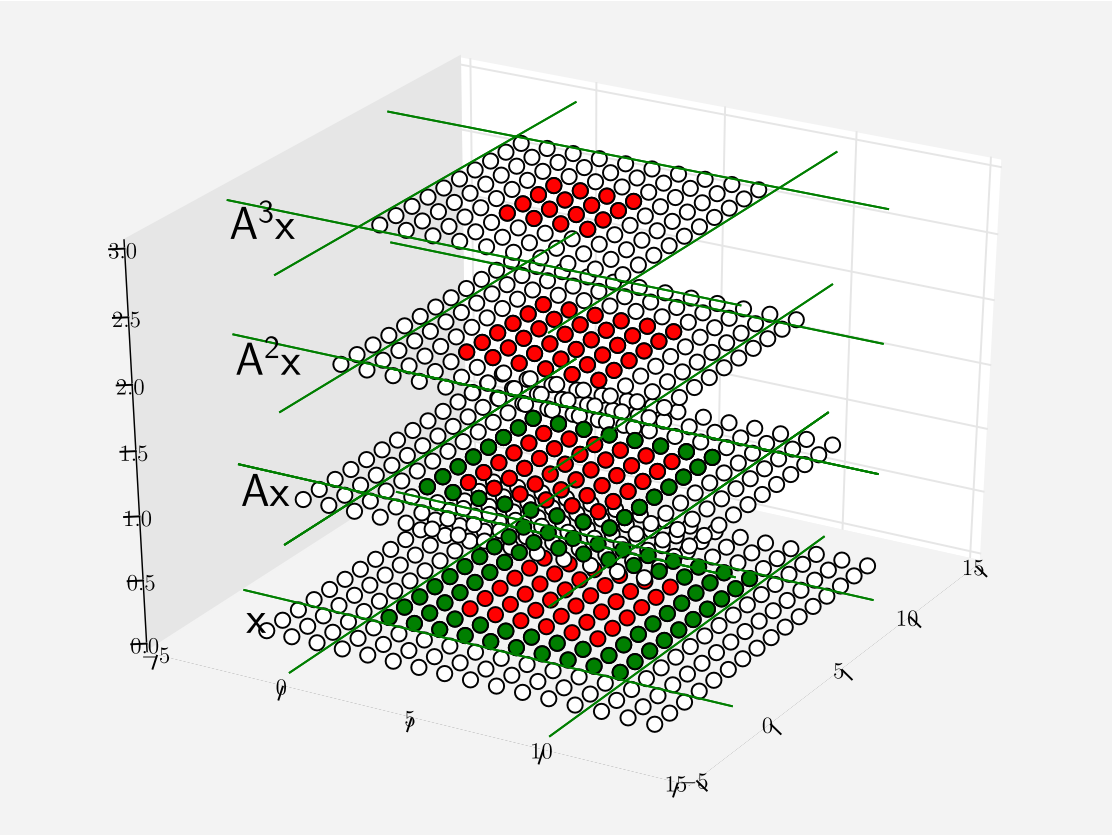

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

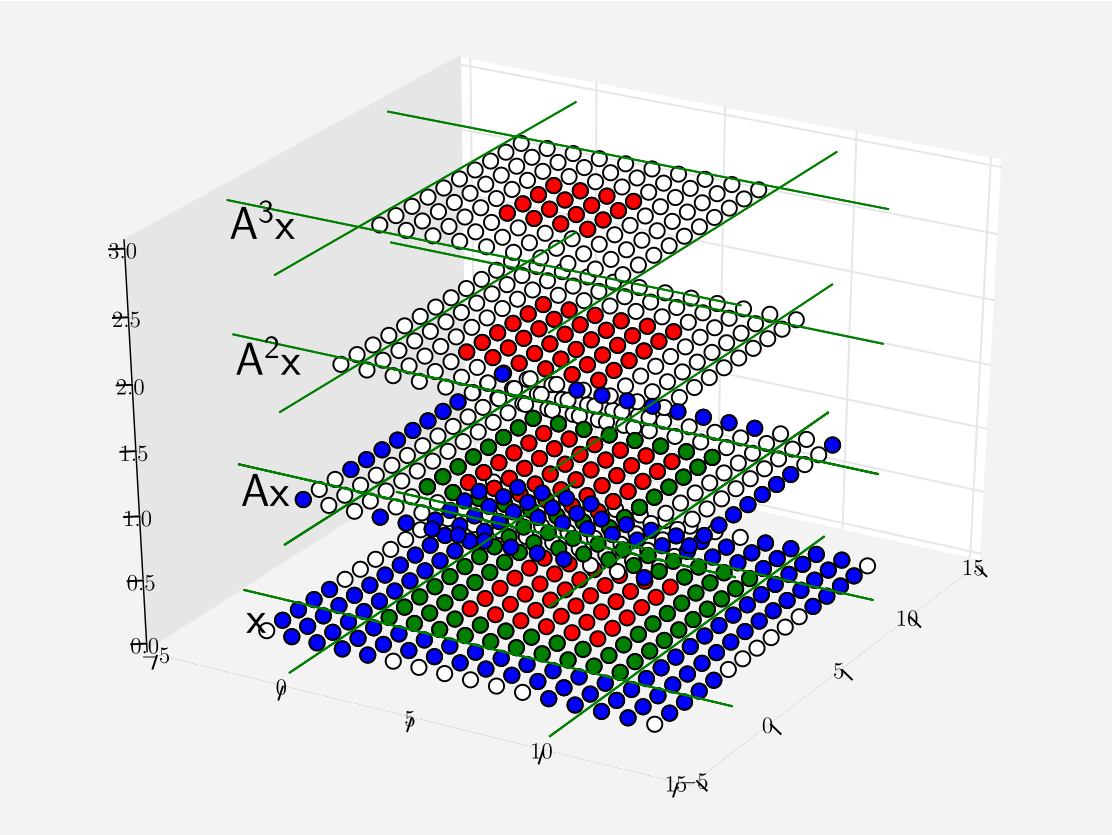

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

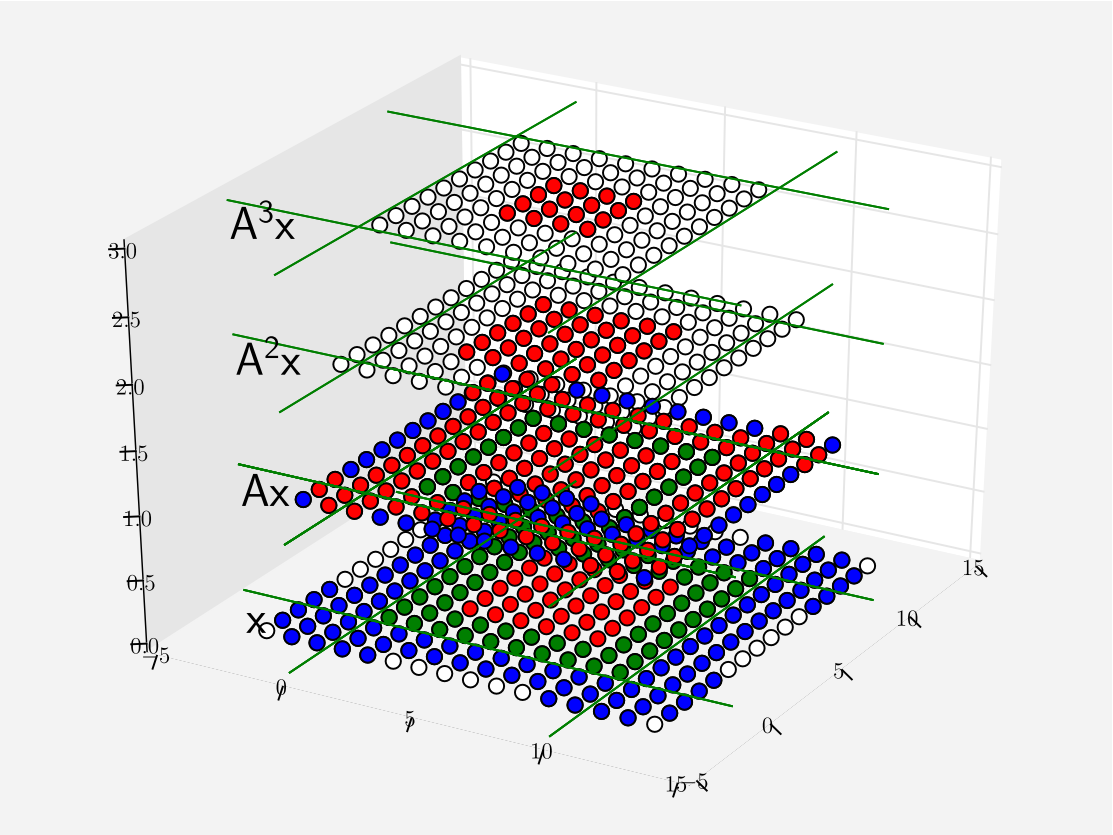

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

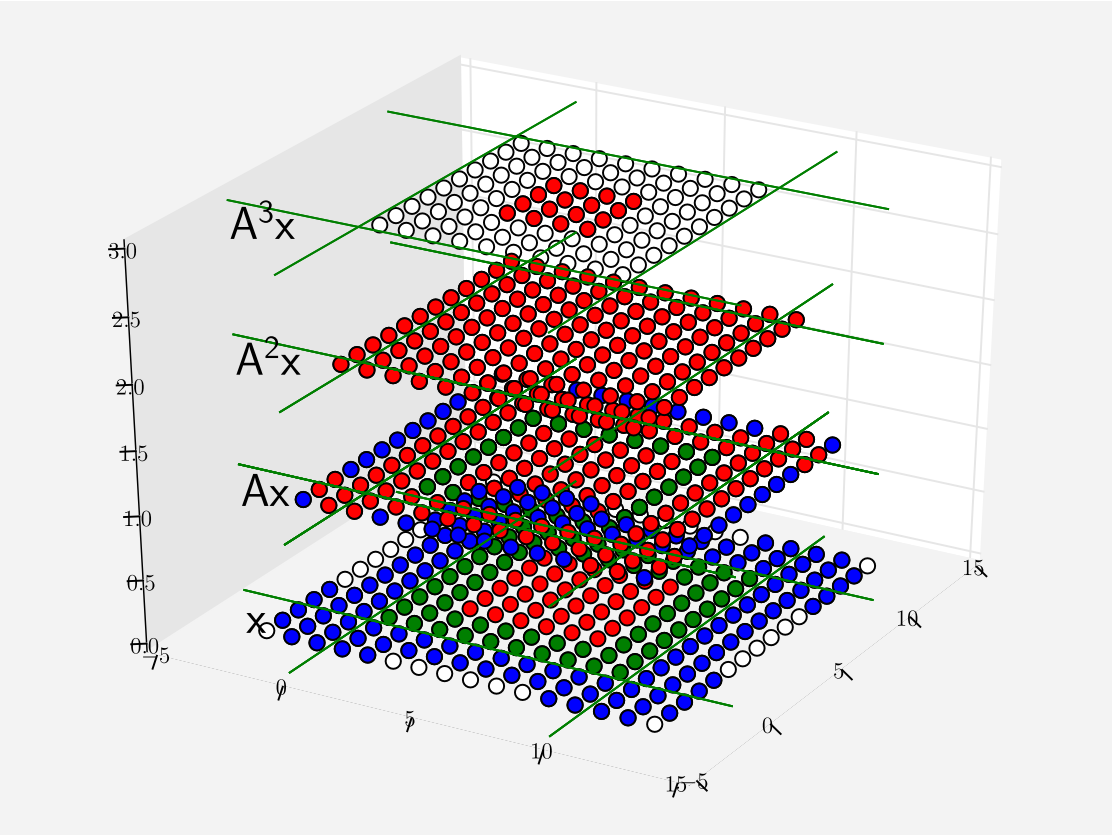

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Improving PA1: 9-pt operator for 3 iterations

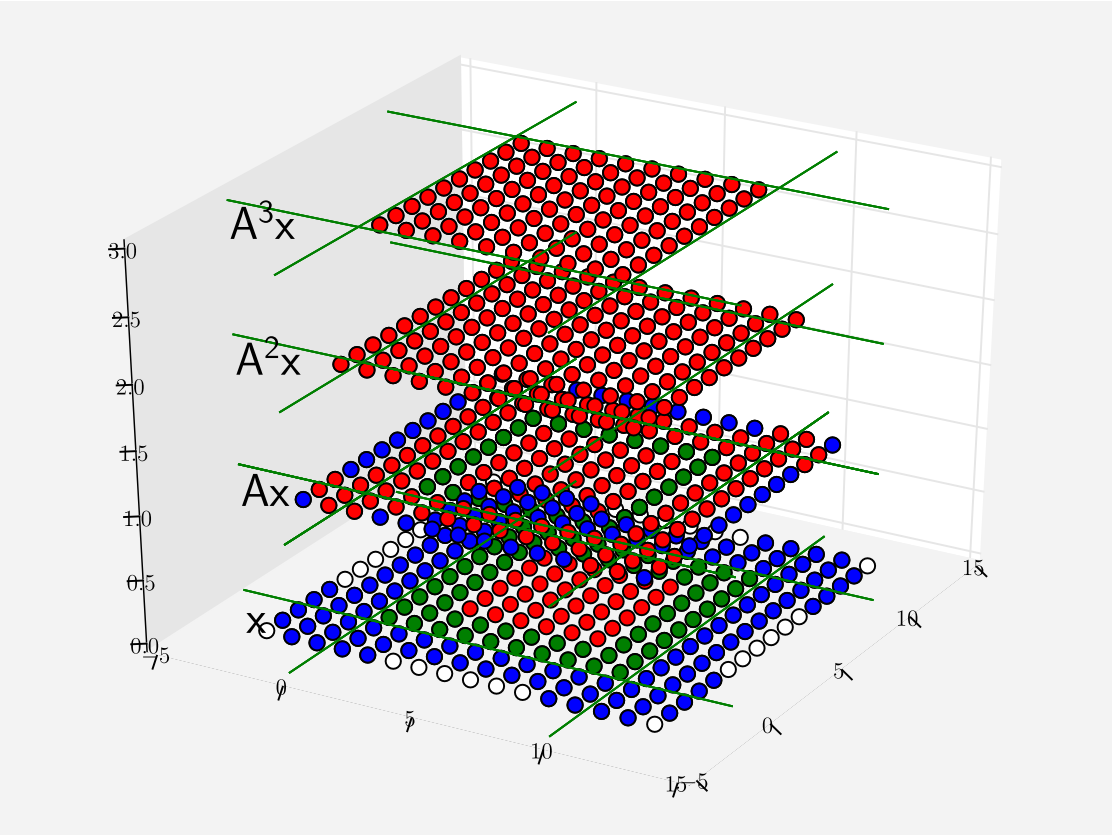

1. Entries available marked red.
2. Compute entries needed by other procs.
3. Send entries to other procs (green).
4. Compute remaining locally computable.
5. Receive entries from other procs (blue).
6. Compute remotely dependent entries.

Parallel Algorithm 2 (PA2)

1. Compute the locally computable entries of $x, Ax, \ldots, A^k x$ needed by other procs. These are the entries on the boundary of locally computable and remotely dependent entries.
2. Send the entries of $x, Ax, \ldots, A^k x$ needed by other procs.
3. Compute remaining locally computable entries of $[Ax, \ldots, A^k x]$.
4. Receive the entries of $x, Ax, \ldots, A^k x$ needed from other procs.
5. Compute the remaining entries of $A^j x$ for $1 \leq j \leq k$ using the already computed entries and the fetched entries.

- $O(1)$ messages between any 2 procs..
- Fewer flops than PA1.
  - Half as much redundant computation as PA1
    - Locally computable entries computed on only their host proc.

Sequential Algorithm

1. For $i = 1, \ldots, p$, (p= number of blocks of A and x)
2. Load block i from slow memory to fast memory.
3. Load parts of x needed from other blocks in to fast memory.
4. Compute the local entries of $Ax, \ldots, A^k x$.
5. Store the computed entries in to slow memory.

- Entries of x and A may be reordered to minimize the cost of accessing slow memory.
  - Minimizing number of slow memory accesses.
    - 2 kinds of Travelling Salesman Problems: one for ordering block entries, and other for the ordering of blocks
  - Minimizing number of entries fetched from slow memory:
    - Extra entries fetched to reduce the number of slow memory accesses.

Sequential Algorithm: Ordering Example for 9-pt Operator

Block level ordering

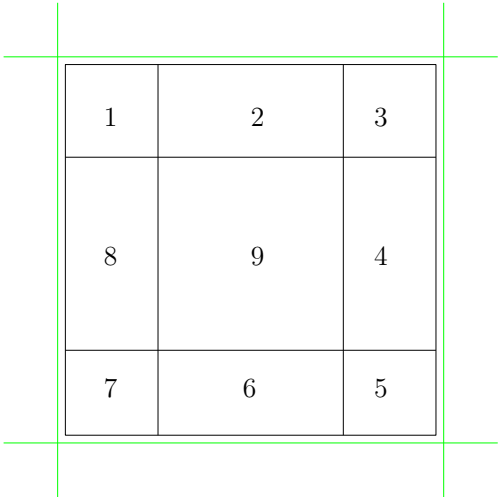

- Left block needs 2 accesses to fetch the entries in 1, 7, 8.
- Other blocks need 1 access to fetch their needed entries.

Global ordering

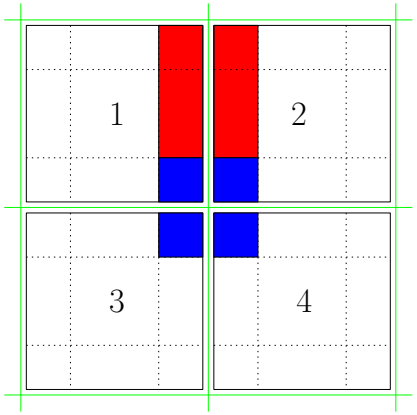

- Computing block 2 after block 1 $\Rightarrow$ colored regions of x do not need to be fetched
- Computing block 4 after block 1 $\Rightarrow$ only the blue regions of x do not need to fetched

Modeled Speedup for Sequential Out-of-Core Algorithm on 9-pt Operator

OOC: 2D 9-pt stencil

$p_{max} = 10^{12}$, Lat=5.7 ms, Bw=62.5 MBytes/s, flops/s=500 MFlops/s, mem=4 GBytes

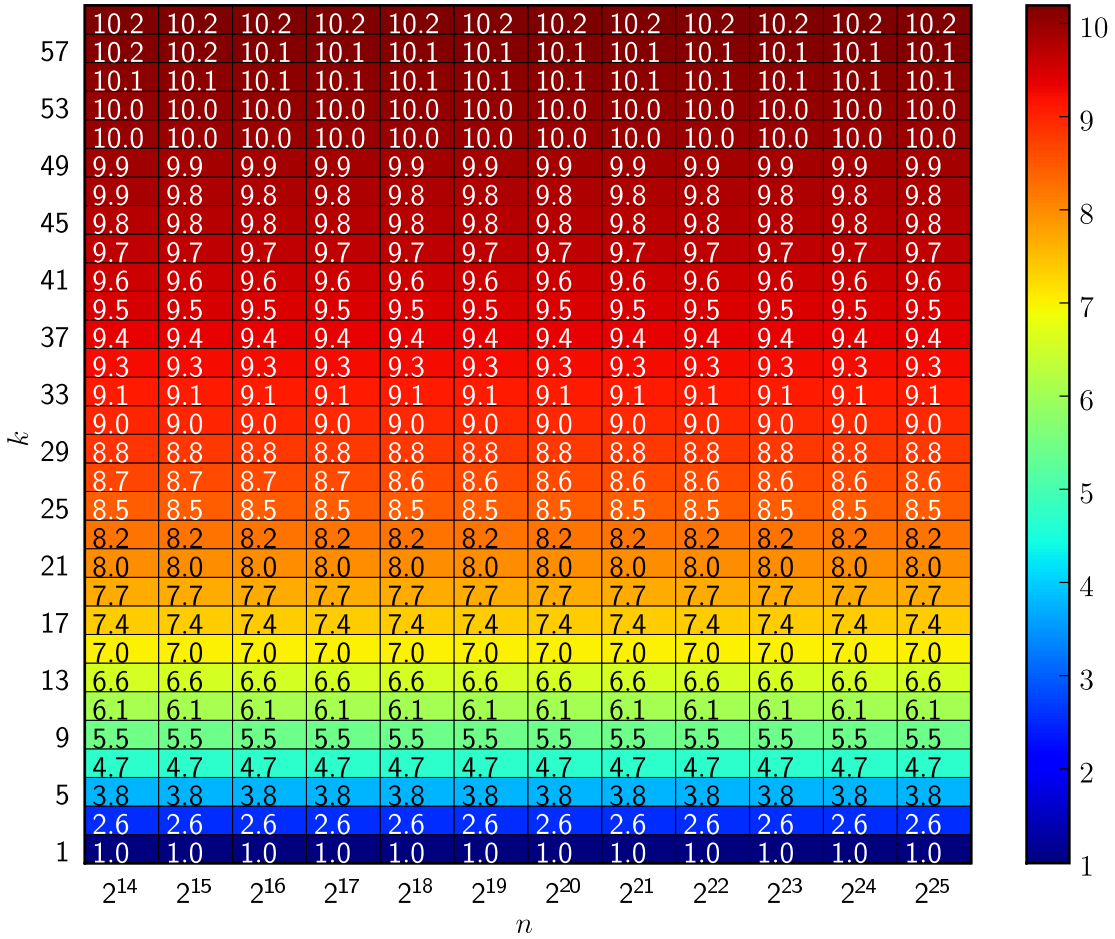

- 500 MFlops/s, mem = 4 GBytes, lat = 5.7 ms, bw = 62.5 MBytes/s.
- No. of blocks p ($1 \leq p \leq p_{max}$) chosen for best perf.
- Speedups across whole range of problem sizes (at least 10x)
  - Reading A and x always costs bw. $\Rightarrow$ speedups always possible.

Modeled Speedup for Sequential Out-of-Core Algorithm on 27-point Operator

OOC: 3D 27-pt stencil

$p_{max} = 10^{12}$, Lat=5.7 ms, Bw=62.5 MBytes/s, flops/s=500 MFlops/s, mem=4 GBytes

k \ n	2^8	2^9	2^{10}	2^{11}	2^{12}	2^{13}	2^{14}	2^{15}	2^{16}	2^{17}
57	9.24									
	9.30									
	9.34									
53	9.39			4.71	4.71	4.71	4.67	4.66	4.65	4.65
	9.42	5.47	5.46	5.13	4.96	4.96	4.92	4.90	4.90	4.90
49	9.46	5.57	5.56	5.22	5.22	5.22	5.18	5.16	5.16	5.16
	9.48	5.67	5.65	5.66	5.48	5.48	5.44	5.42	5.41	5.41
45	9.50	5.76	5.75	5.75	5.75	5.66	5.66	5.66	5.66	5.66
	9.51	7.42	6.56	6.19	6.01	5.92	5.92	5.90	5.90	5.90
41	9.51	7.47	6.64	6.28	6.27	6.18	6.14	6.14	6.13	6.13
	9.50	7.53	6.71	6.36	6.35	6.35	6.35	6.35	6.35	6.35
37	9.47	7.57	6.78	6.78	6.60	6.60	6.56	6.56	6.56	6.56
	9.43	7.60	6.84	6.84	6.84	6.75	6.75	6.75	6.75	6.75
33	9.38	7.63	7.61	7.24	7.06	6.97	6.93	6.93	6.92	6.92
	9.31	7.64	7.62	7.27	7.09	7.09	7.09	7.09	7.08	7.08
29	9.21	7.63	7.61	7.28	7.27	7.27	7.23	7.21	7.20	7.20
	9.09	7.60	7.58	7.59	7.42	7.34	7.30	7.30	7.30	7.30
25	8.94	7.55	7.53	7.54	7.38	7.38	7.38	7.36	7.36	7.36
	8.76	7.47	7.45	7.46	7.46	7.39	7.39	7.39	7.39	7.39
21	8.54	8.53	7.90	7.62	7.48	7.41	7.38	7.36	7.36	7.36
	8.28	8.26	7.70	7.45	7.32	7.32	7.29	7.29	7.29	7.28
17	7.95	7.94	7.44	7.22	7.22	7.16	7.16	7.14	7.14	7.14
	7.57	7.55	7.12	6.93	6.92	6.92	6.92	6.92	6.92	6.91
13	7.10	7.09	6.72	6.73	6.64	6.60	6.60	6.60	6.60	6.60
	6.53	6.52	6.23	6.23	6.23	6.20	6.18	6.17	6.17	6.17
9	5.84	5.83	5.61	5.62	5.62	5.61	5.60	5.60	5.60	5.60
	5.00	4.99	4.99	4.92	4.88	4.86	4.85	4.85	4.85	4.85
5	3.96	3.95	3.95	3.91	3.89	3.89	3.89	3.89	3.89	3.89
	2.65	2.65	2.65	2.64	2.64	2.63	2.63	2.63	2.63	2.63
1	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

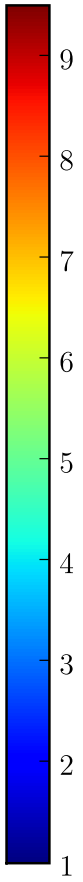

- Speedups across whole range of problem sizes (at least 7x).
- Speedups decrease after a certain k.

Modeled Speedup for Sequential Algorithm with On-Chip Cache as Fast Memory

- 9-pt. operator on $n \times n$ mesh, 27-pt. operator on $n \times n \times n$ mesh.
- Core performance = 2 GFlops/s, memory=8 MBytes, lat=200 ns, bw=5 GBytes/s
  - Models single core, single socket of a quad-core Intel Clovertown chip

Matrix	Range of n	(Range of) Modeled Speedup
9-pt operator	2^8 to 2^{19}	[2.45, 2.58]
27-pt operator	2^8 to 2^{12}	[1.34, 1.36]

- Speedups always possible across all problem sizes

Performance Model: Parallel Algorithm (PA2)

- 9-pt. operator on $n \times n$ mesh, 27-pt. operator on $n \times n \times n$ mesh.

Peta: No. procs. = 8100, proc. performance = 50 GFlops/s, memory=500 GBytes, lat=10 μs, bw=4 GBytes/s

Matrix	Range of n	Max Modeled Speedup
9-pt operator	2^{10} to 2^{22}	6.9
27-pt operator	2^{9} to 2^{14}	1.02

- Speedups for small problem sizes (for 9-pt operator, $2^{10} \leq n \leq 2^{13}$).
- Other problem sizes computation bound, so not limited by communication (hidden by overlap with communication).

Performance Model: Parallel Algorithm (PA2)

Grid: No. procs. = 125, proc. performance = 1 TFlops/s, memory=10 TBytes, lat=100 ms, bw=320 MBytes/s

Matrix	Range of n	Max Modeled Speedup
9-pt operator	2^{10} to 2^{22}	22.22
27-pt operator	2^{9} to 2^{14}	4.41

- Small problem sizes run on 1 proc.
- Speedups for moderate problem sizes (for 9-pt operator, $2^{15} \leq n \leq 2^{19}$).
- Large problem sizes computation bound.

Implementation

- Parallel algorithm implemented in UPC (Unified Parallel C).
  - Works for general sparse matrices.
  - For PA1, entries of A and x reordered to make local computations as k invocations of Sparse Matrix Vector multiplication.
- Sequential (out-of-core) algorithm implemented in C.
  - Slow memory assumed to be disk.
  - Reordering done to minimize bandwidth cost for disk access using a randomized heuristic.

Results: Sequential Out-of-Core Algorithm

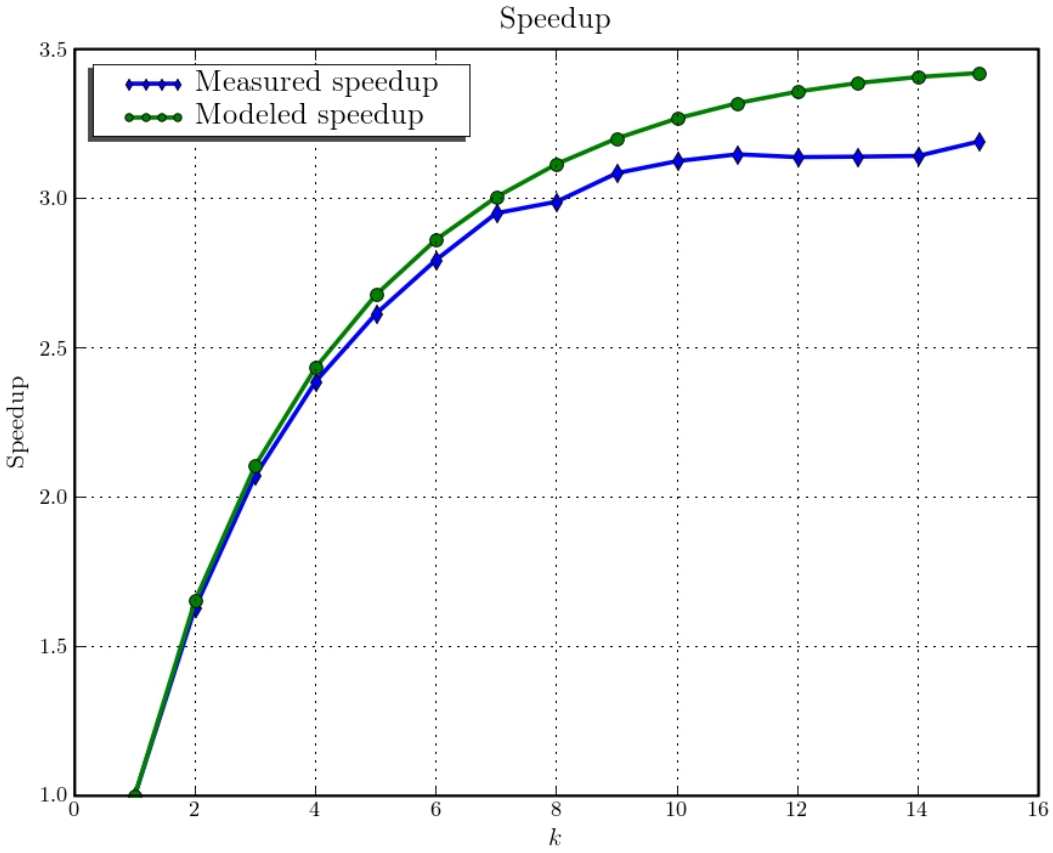

- Itanium II node with 5.2 GFlops peak flop rate.
- 27-point operator with 368^3 points partitioned in to $4^3 = 64$ blocks.
- Performance 6x slower than ideal machine (∞ DRAM).

Summary and Future Work

- Sequential and parallel communication avoiding algorithms for the Akx kernel.
  - Almost linear speedups possible.
  - Minimum latency cost.
  - Minimum bandwidth cost for the sequential algorithm.
- Performance modeling of the algorithms.
  - Parallel implementation expected to achieve speedups for moderate problem sizes.
  - Sequential implementation expected to achieve speedups across the whole range of problem sizes.
- Sequential implementation demonstrates speedup of 3x for a 27-point operator.
- Akx kernel part of a larger effort for communication avoiding iterative solvers
  - Extensions to polynomial bases
  - Incorporating preconditioning (matrix A multiplied by a preconditioner matrix M)

Thank You!

Modeled Speedup for Sequential Algorithm with Cache as Fast Memory on 9-pt Operator

CacheBlocked: 2D 9-pt stencil

$p_{max} = 10^{12}$, Lat=200 ns, Bw=5 GBytes/s, flops/s=2 GFlops/s, mem=8 MBytes

k \ n	2^8	2^9	2^{10}	2^{11}	2^{12}	2^{13}	2^{14}	2^{15}	2^{16}	2^{17}	2^{18}	2^{19}
49				1.53	1.50	1.50	1.50	1.49	1.49	1.49	1.49	1.49
45			1.56	1.56	1.56	1.56	1.56	1.56	1.56	1.56	1.56	1.56
		1.81	1.70	1.65	1.63	1.63	1.63	1.62	1.62	1.62	1.62	1.62
41	1.85	1.85	1.74	1.69	1.69	1.69	1.69	1.69	1.69	1.69	1.69	1.69
	1.89	1.89	1.78	1.78	1.78	1.77	1.76	1.76	1.76	1.76	1.76	1.76
37	1.92	1.92	1.92	1.87	1.84	1.83	1.83	1.83	1.83	1.82	1.82	1.82
	2.45	1.96	1.96	1.91	1.91	1.90	1.90	1.89	1.89	1.89	1.89	1.89
33	2.47	2.00	2.00	2.00	1.97	1.96	1.96	1.96	1.96	1.96	1.96	1.96
	2.49	2.04	2.04	2.04	2.04	2.03	2.03	2.03	2.02	2.02	2.02	2.02
29	2.51	2.28	2.18	2.13	2.10	2.09	2.09	2.09	2.09	2.09	2.09	2.09
	2.53	2.31	2.21	2.16	2.16	2.16	2.16	2.15	2.15	2.15	2.15	2.15
25	2.55	2.34	2.25	2.24	2.22	2.22	2.22	2.21	2.21	2.21	2.21	2.21
	2.56	2.37	2.28	2.28	2.28	2.28	2.27	2.27	2.27	2.27	2.27	2.27
21	2.57	2.39	2.39	2.35	2.33	2.33	2.32	2.32	2.32	2.32	2.32	2.32
	2.58	2.41	2.41	2.37	2.37	2.37	2.37	2.37	2.37	2.37	2.37	2.37
17	2.58	2.43	2.43	2.43	2.41	2.41	2.40	2.41	2.40	2.40	2.40	2.40
	2.57	2.44	2.44	2.44	2.44	2.44	2.44	2.43	2.43	2.43	2.43	2.43
13	2.55	2.55	2.49	2.46	2.45	2.45	2.45	2.45	2.45	2.45	2.45	2.45
	2.52	2.52	2.47	2.44	2.44	2.44	2.44	2.44	2.44	2.44	2.44	2.44
9	2.46	2.46	2.42	2.42	2.42	2.41	2.41	2.41	2.41	2.41	2.41	2.41
	2.36	2.36	2.36	2.34	2.33	2.33	2.33	2.33	2.33	2.33	2.33	2.33
5	2.18	2.18	2.18	2.17	2.17	2.16	2.16	2.16	2.16	2.16	2.16	2.16
	1.83	1.83	1.83	1.83	1.83	1.83	1.83	1.83	1.83	1.83	1.83	1.83
1	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

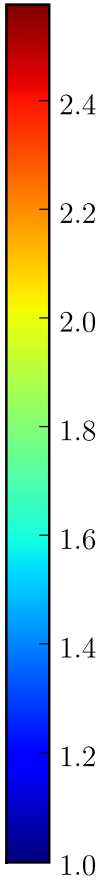

- 2 GFlops/s, mem = 4 GBytes, lat = 200 ns, bw = 5 GBytes/s.
  - Models single core, single socket of a quad-core Intel Clovertown chip
- No. of blocks p ($1 \leq p \leq p_{max}$) chosen for best perf.
- Speedups across whole range of problem sizes (at least 2.4x)

Modeled Speedup for Sequential Algorithm with Cache as Fast Memory on 27-pt Operator

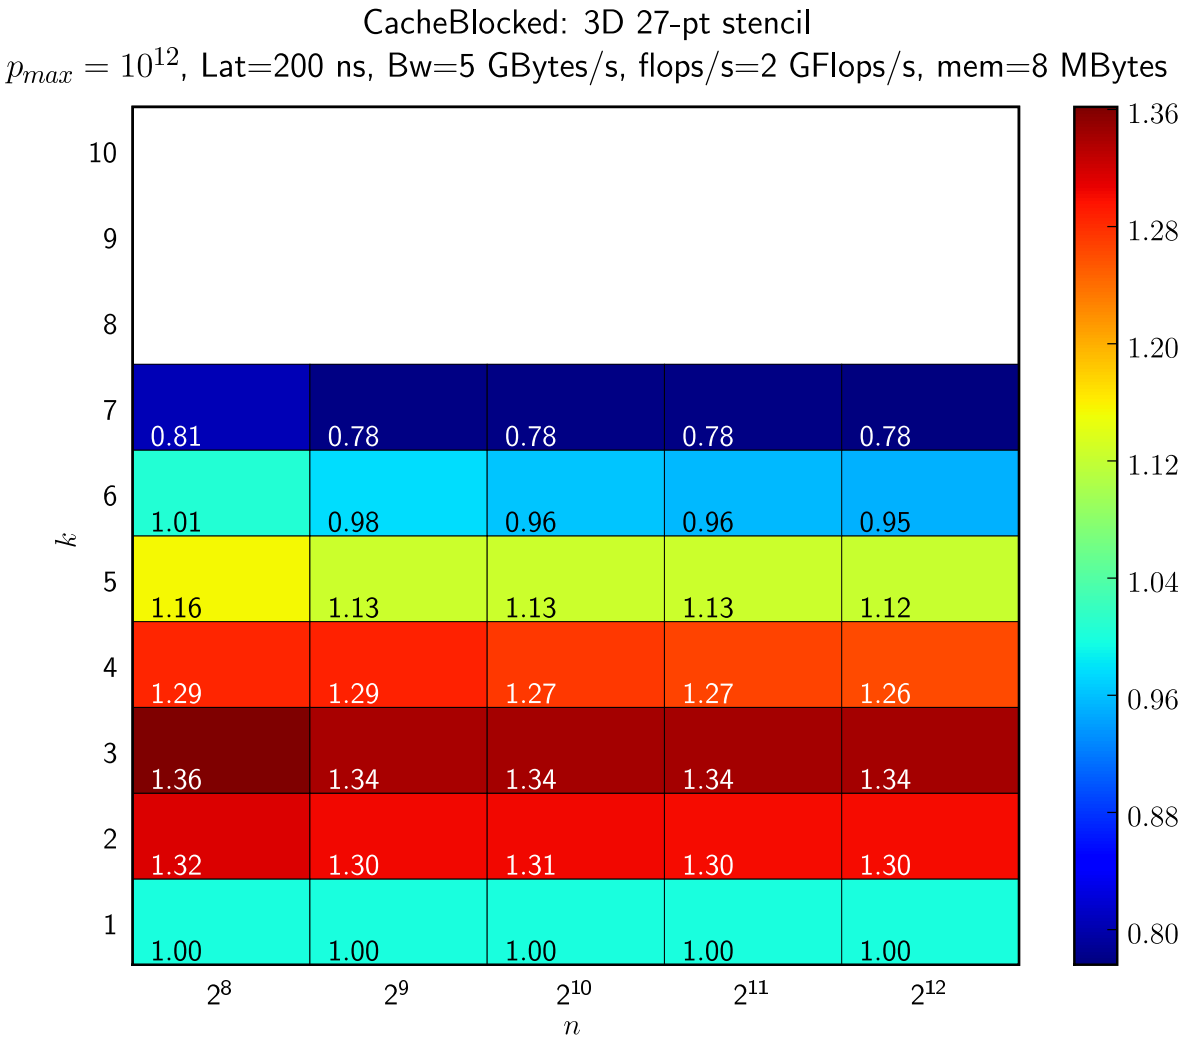

- Speedups across whole range of problem sizes (at least 1.3x)
- Surface effects limit maximum range of k