

JACOBI ITERATIVE SOLVER

Udhayasankar, Naveen - CSE 708, Fall
2022

Final Presentation – 12/06/2022



Solving a system of linear equations

$$Ax + B$$

$$\begin{aligned} a_{11}x_1 + \dots + a_{1n}x_n &= b_1 \\ a_{n1}x_1 + \dots + a_{nn}x_n &= b_n \end{aligned}$$

1. Direct or Exact methods – Result is obtained in a finite number of steps. Error is zero.
2. Iterative methods – Result is an array of approximate values, which converge to the exact result.



Jacobi Method

$$x_1 = \frac{b_1}{a_{11}} - \frac{a_{12}x_2}{a_{11}} - \dots - \frac{a_{1n}x_n}{a_{11}}$$

$$x_1 = \frac{b_1}{a_{11}} - \frac{\sum_{\substack{j=2 \\ j \neq 1}}^n a_{1j}x_j}{a_{11}}$$

$$x_2 = \frac{b_2}{a_{22}} - \frac{\sum_{\substack{j=1 \\ j \neq 2}}^n a_{2j}x_j}{a_{22}}$$

$$x_i = \frac{b_i}{a_{ii}} - \frac{\sum_{\substack{j=1 \\ j \neq i}}^n a_{ij}x_j}{a_{ii}}$$

1. Solve for an unknown using other unknowns.
2. Assume an initial value for all the unknowns, a commonly used value is 0.
3. For each iteration, calculate the unknowns – approximate solutions (vector of size n)
4. Run the iterations until an acceptable solution (a solution close to the exact value) is reached.
5. Each of the iterations produce an approximate solution for the real values of the system.

Jacobi Method – Very Large Number of Unknowns

- Matrix representation

$$\begin{array}{c}
 A * x = b \\
 \downarrow \quad \downarrow \quad \swarrow \\
 n \times n \quad \text{vector} \\
 \text{matrix} \quad \text{of size } n
 \end{array}$$

- For $n = 4$,

$$A = \begin{bmatrix}
 a_{00} & a_{01} & a_{02} & a_{03} \\
 a_{10} & a_{11} & a_{12} & a_{13} \\
 a_{20} & a_{21} & a_{22} & a_{23} \\
 a_{30} & a_{31} & a_{32} & a_{33}
 \end{bmatrix}$$

$$B = \begin{bmatrix}
 b_0 \\
 b_1 \\
 b_2 \\
 b_3
 \end{bmatrix}
 \quad
 x = \begin{bmatrix}
 x_0 \\
 x_1 \\
 x_2 \\
 x_3
 \end{bmatrix}$$

Jacobi Method – Very Large Number of Unknowns

- Matrix A is the sum of Lower Triangle, Diagonal and Upper Triangle

$$A = \begin{pmatrix}
 a_{00} & a_{01} & a_{02} \\
 a_{10} & a_{11} & a_{12} \\
 a_{20} & a_{21} & a_{22}
 \end{pmatrix}$$

$$A = L + D + U$$

$$Ax = b$$

$$(L + D + U)x = b$$

$$Lx + Dx + Ux = b$$

$$Dx = b - Lx - Ux$$

$$Dx = b + Dx - Lx - Ux - Dx$$

$$Dx = b + Dx - Ax$$

$$x = x + D^{-1}(b - Ax)$$

$$x^{k+1} = x^k + \Delta x, \text{ where } \Delta x = D^{-1}(b - Ax)$$

Jacobi Method – Algorithm

Given A , b , n , x_{old} and tolerance,
for max_iterations:

$$\Delta x = D^{-1}(b - (A * x_{old}))$$

$$x_{old} = x_{old} + \Delta x$$

if $\text{abs}(\Delta x) < \text{tolerance}$:

break

end for

- Run the algorithm for maximum number of iterations or until the algorithm converges.
- Calculate Δx
- Check for convergence i.e., if Δx is less than the given tolerance, the algorithm has converged or produced an acceptable result.
- Update the approximate solution x_{old} to be used in the next iteration.

Jacobi Method – Why?

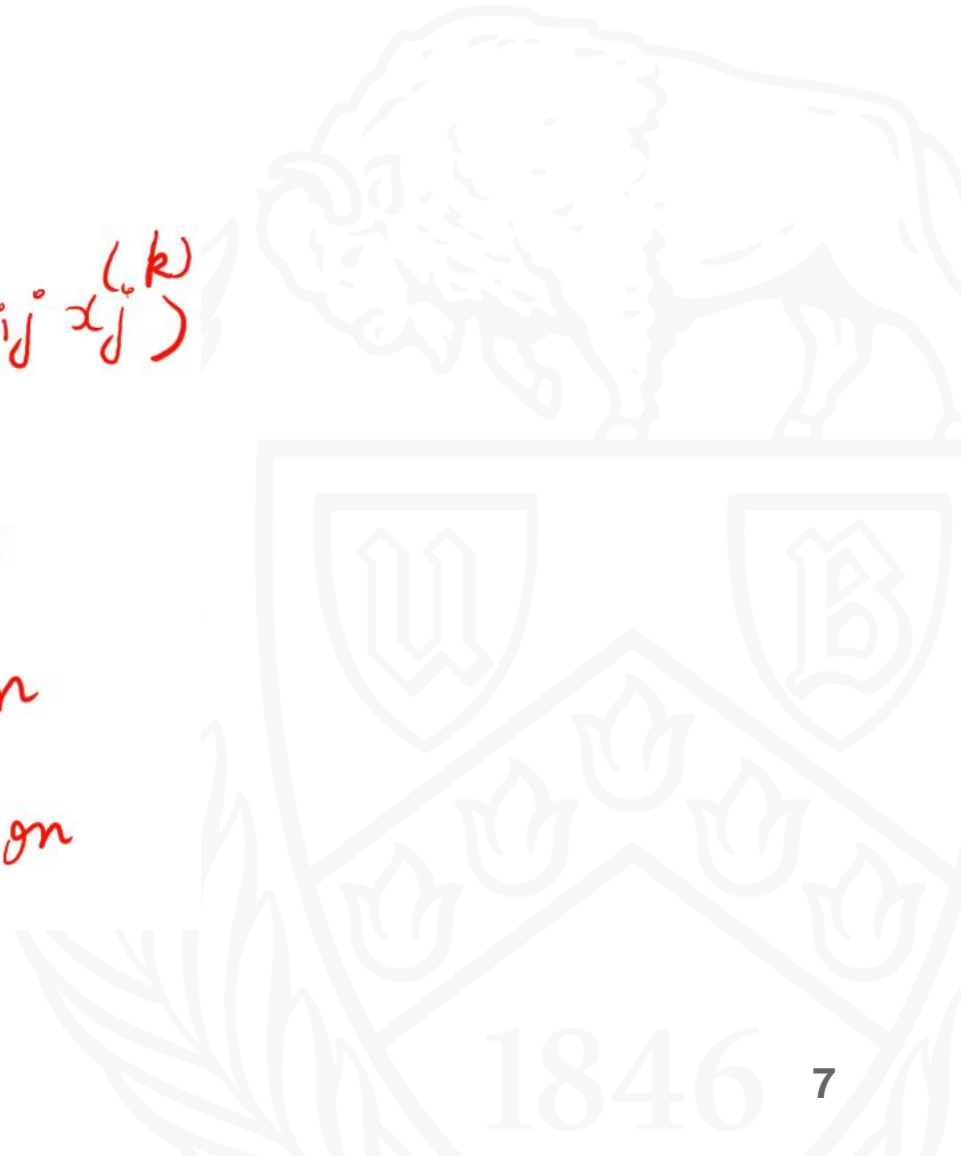
Why should the Jacobi method be parallelized?

$$x_i^{(k+1)} = 1/a_{ii} (b_i - \sum_{j \neq i} a_{ij} x_j^{(k)})$$

$i = 1, 2, \dots, n$

$k+1 \rightarrow$ current iteration

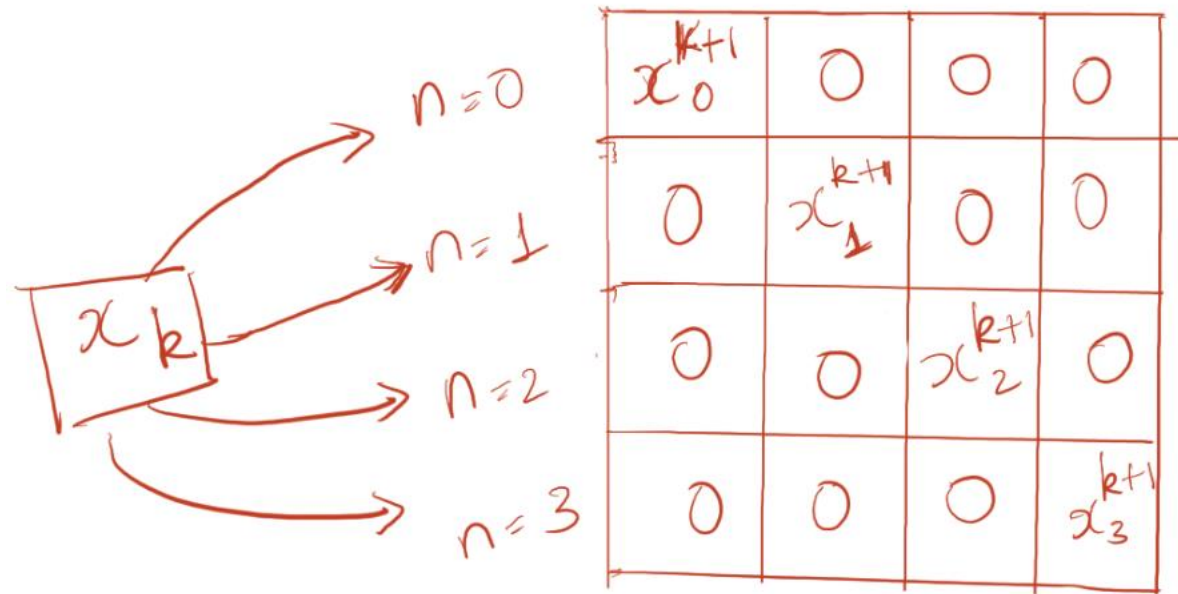
$k \rightarrow$ previous iteration



Jacobi Method – Why?

The i^{th} processor can calculate the i^{th} value to be used in the next iteration, parallelly

$n=0$	x_0^{k+1}	0	0	0
$n=1$	0	x_1^{k+1}	0	0
$n=2$	0	0	x_2^{k+1}	0
$n=3$	0	0	0	x_3^{k+1}



Jacobi Method – How?

The i^{th} processor calculates the i^{th} value to be used in the next iteration, parallelly, HOW?

The value of x from previous iteration is known to all processors.

num - proc = 4

$$\begin{aligned}
 n=0 \quad x_0^{k+1} &= [b_0 - a_{00}x_0^k - a_{01}x_1^k - a_{02}x_2^k - a_{03}x_3^k] / a_{00} \\
 n=1 \quad x_1^{k+1} &= [b_1 - a_{10}x_0^k - a_{11}x_1^k - a_{12}x_2^k - a_{13}x_3^k] / a_{11} \\
 n=2 \quad x_2^{k+1} &= [b_2 - a_{20}x_0^k - a_{21}x_1^k - a_{22}x_2^k - a_{23}x_3^k] / a_{22} \\
 n=3 \quad x_3^{k+1} &= [b_3 - a_{30}x_0^k - a_{31}x_1^k - a_{32}x_2^k - a_{33}x_3^k] / a_{33}
 \end{aligned}$$

Jacobi Method – How?

The value of x from previous iteration is known to all processors, HOW?

MPI_Allgather [3]:

MPI Allgather allows us to “gather” to all processes information from all processes with the communicator. The action of “all gather” is as if we were to gather to one process using MPI Gather, and then send from that process to all other processes the assembled information.

MPI_Allgather:

Function Call Syntax

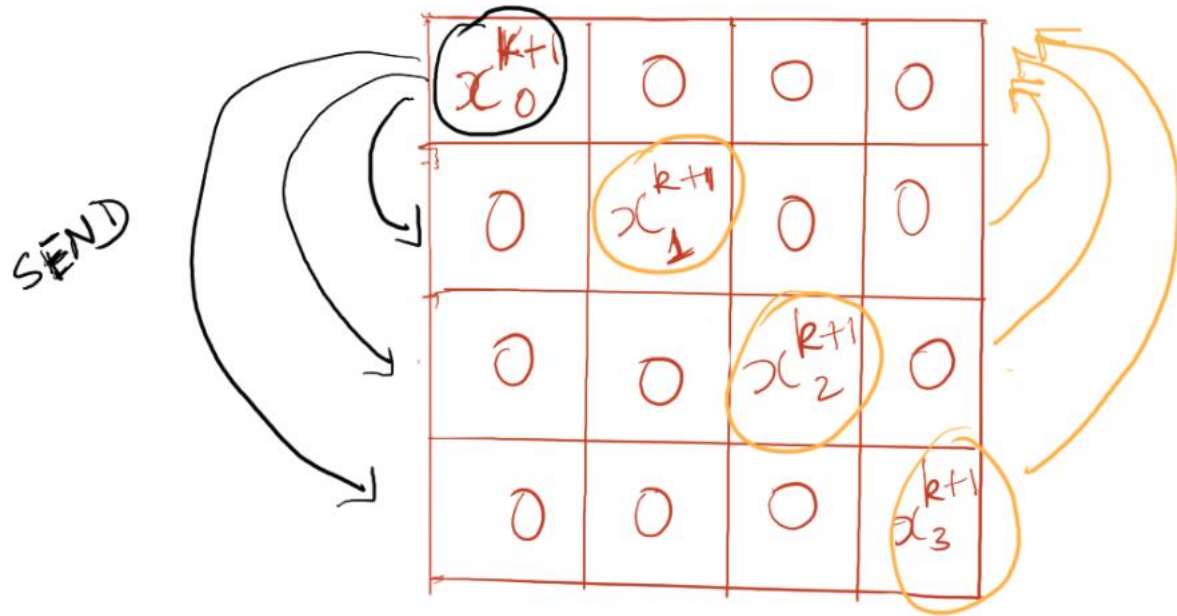
```

int MPI_Allgather(
    void*          sendbuf  /* in  */,
    int           sendcount /* in  */,
    MPI_Datatype  sendtype /* in  */,
    void*          recvbuf  /* out */,
    int           recvcnt  /* in  */,
    MPI_Datatype  recvtype /* in  */,
    MPI_Comm      comm     /* in  */)
    
```

Understanding the Argument List

- *sendbuf* - starting address of the send buffer.
- *sendcount* - number of elements in the send buffer.
- *sendtype* - data type of the elements in the send buffer.
- *recvbuf* - starting address of the receive buffer.
- *recvcnt* - number of elements for any *single* receive.
- *recvtype* - data type of the elements in the receive buffer.
- *comm* - communicator.

Jacobi Method – How?



RECEIVE

$n=0$

$n=1$

$n=2$

$n=3$

x_0^{k+1}	x_1^{k+1}	x_2^{k+1}	x_3^{k+1}
x_0^{k+1}	x_1^{k+1}	x_2^{k+1}	x_3^{k+1}
x_0^{k+1}	x_1^{k+1}	x_2^{k+1}	x_3^{k+1}
x_0^{k+1}	x_1^{k+1}	x_2^{k+1}	x_3^{k+1}

Jacobi Method – When?

When does the Jacobi method converge?

- The Jacobi method converges for strictly row-wise or column-wise diagonally dominant matrices.
- The diagonal elements of the matrix must not be zero.
- For strictly row-wise or column-wise matrices, the Jacobi relaxation will converge “to a good solution” in $\log n$ steps. In these cases, the parallel time complexity is potentially $O(\log n)$ with n processors.

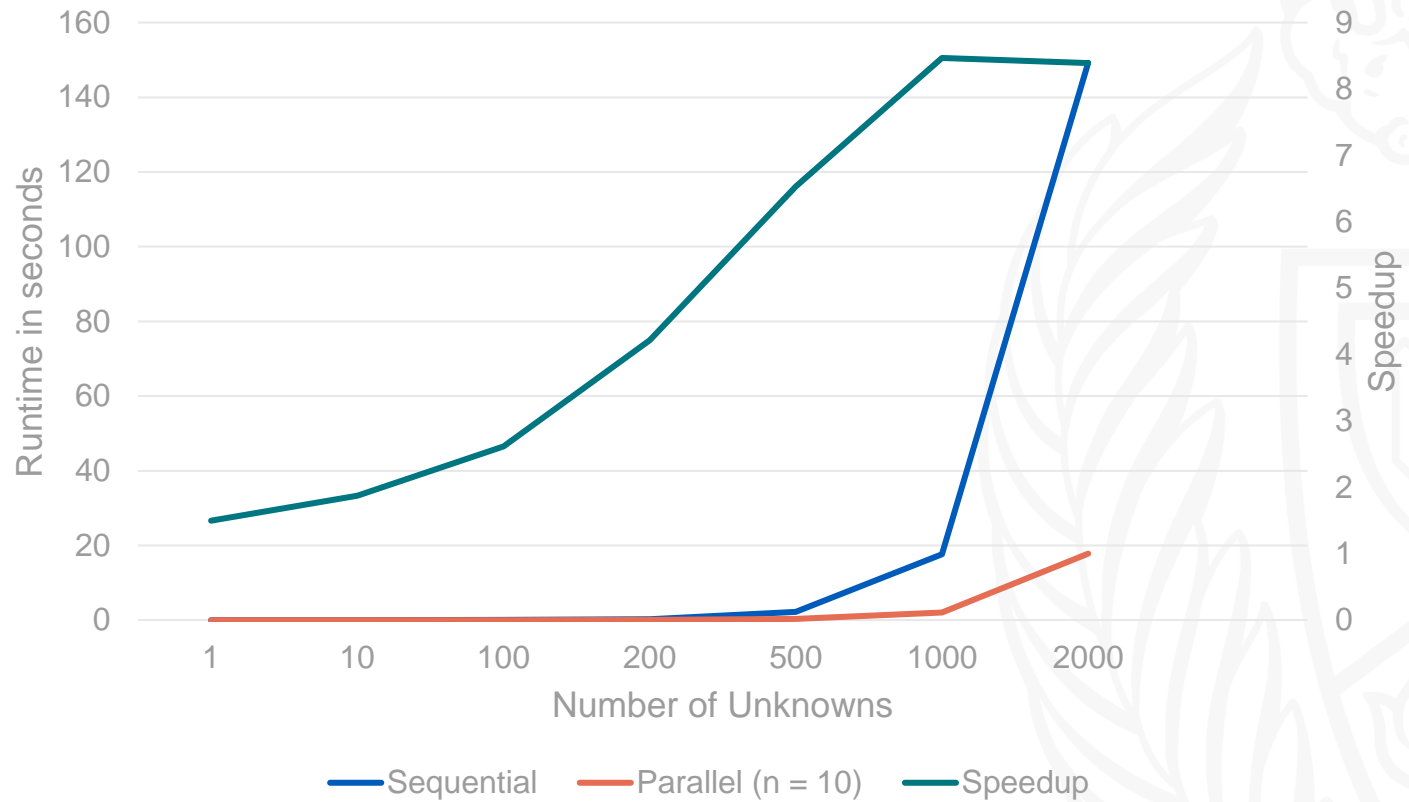
$$a_{i,i} > \sum_{j \neq i} |a_{i,j}|$$

(or)

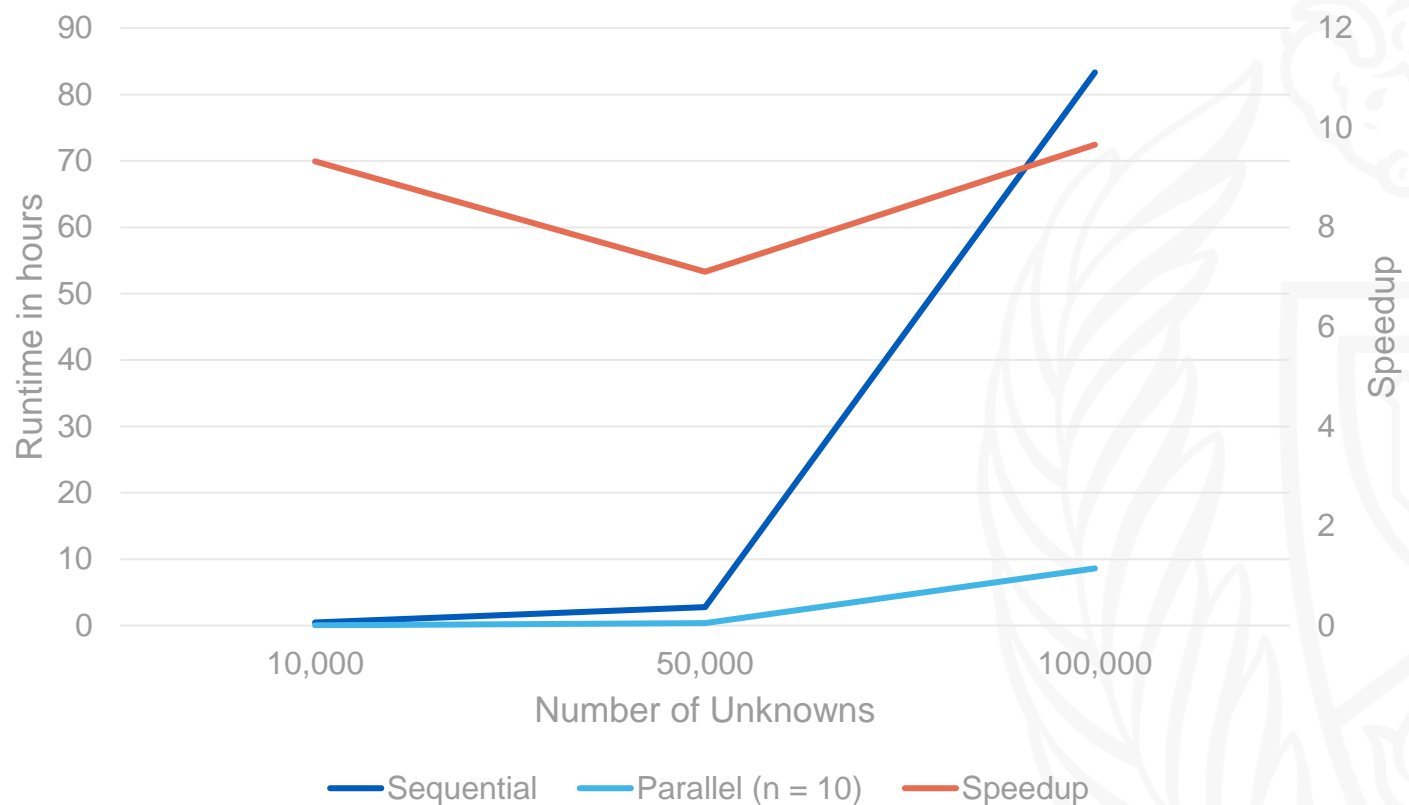
$$a_{i,i} > \sum_{j \neq i} |a_{j,i}|$$

$$i = 1, 2, \dots, n$$

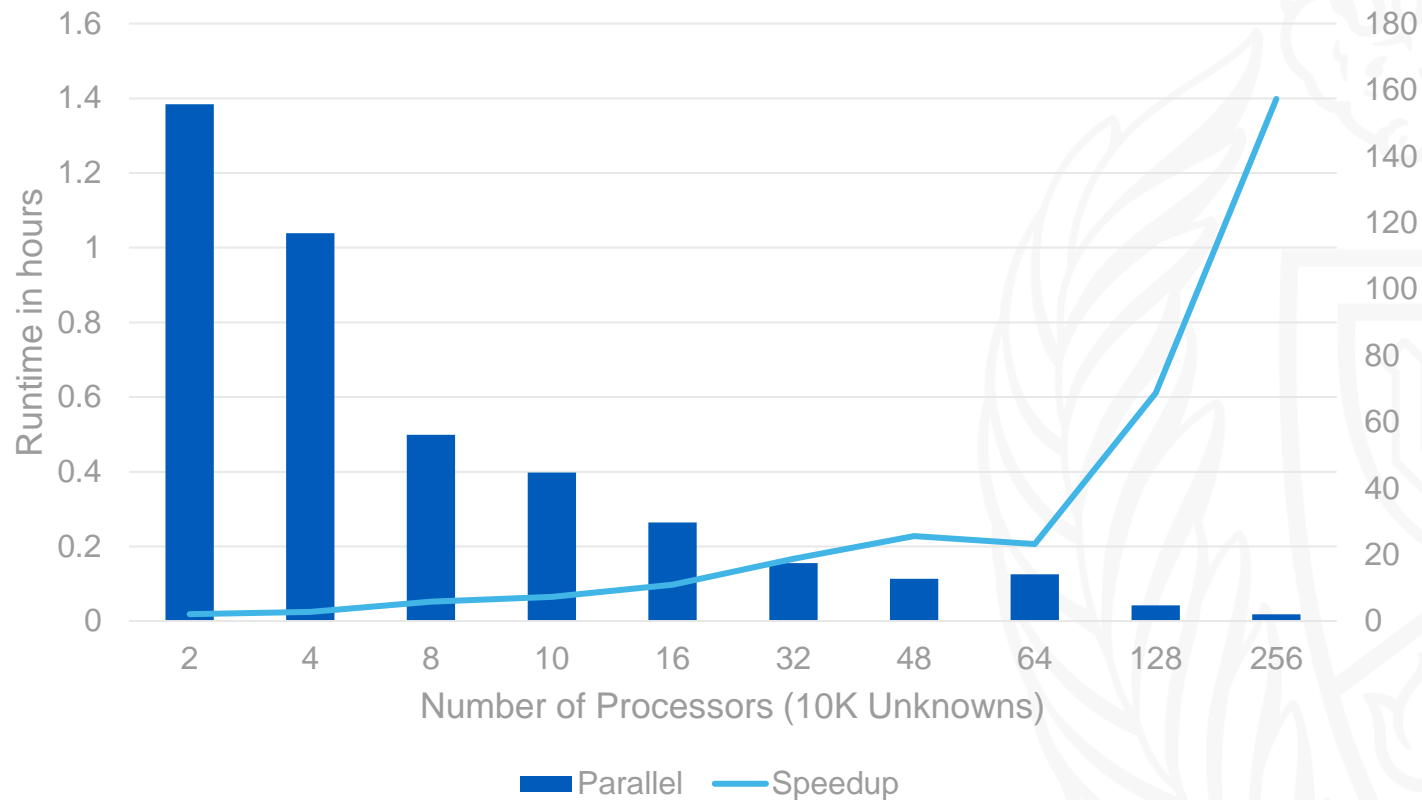
Jacobi Method – Runtime Graphs



Jacobi Method – Runtime Graphs (Larger Data)



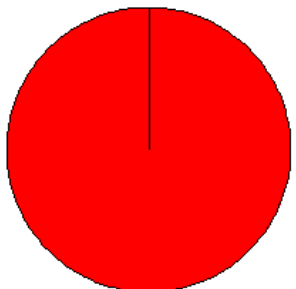
Jacobi Method – Runtime Graphs (Scaling out)



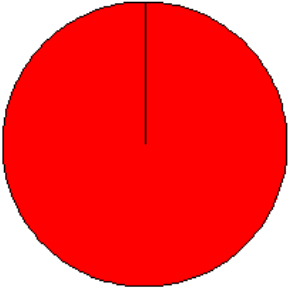
Jacobi Method – Profiling Results - % of communication – 1 Processor

unknown <ul style="list-style-type: none"> • Load Balance • Communication Balance • Message Buffer Sizes • Communication Topology • Switch Traffic • Memory Usage • Executable Info • Host List • Environment • Developer Info <p>powered by IPM</p>			command: ./mpi_jacobi 1000				
			codename: unknown	state: unknown			
			username:		group:		
			host:	mpi_tasks: 1 on 1 hosts			
			start:	wallclock: 3.01195e+01 sec			
			stop:	%comm: 0.0402785570809608			
			total memory:	0.0546227 gbytes	total gflop/sec:	0	
			switch(send):	0 gbytes	switch(rcv):	0 gbytes	
			Computation			Communication	
			Event	Count	Pop	% of MPI Time	
			<ul style="list-style-type: none"> ■ MPI_Allgather ■ MPI_Comm_rank 				

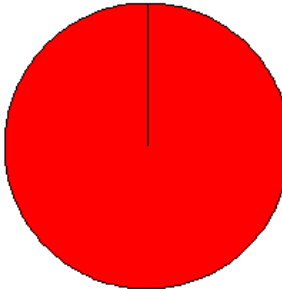
Jacobi Method – Profiling Results - % of communication – 2 Processors

unknown <ul style="list-style-type: none"> Load Balance Communication Balance Message Buffer Sizes Communication Topology Switch Traffic Memory Usage Executable Info Host List Environment Developer Info 			command: ./mpi_jacobi 1000					
			codename:	unknown	state:	unknown		
			username:		group:			
			host:		mpi_tasks:	2 on 1 hosts		
			start:		wallclock:	1.21119e+03 sec		
			stop:		%comm:	96.7734211808222		
			total memory:	0.1115722 gbytes	total gflop/sec:	0		
			switch(send):	0 gbytes	switch(recv):	0 gbytes		
			Computation			Communication		
			Event	Count	Pop	% of MPI Time		
			 <ul style="list-style-type: none"> MPI_Allgather MPI_Comm_rank MPI_Comm_size 					

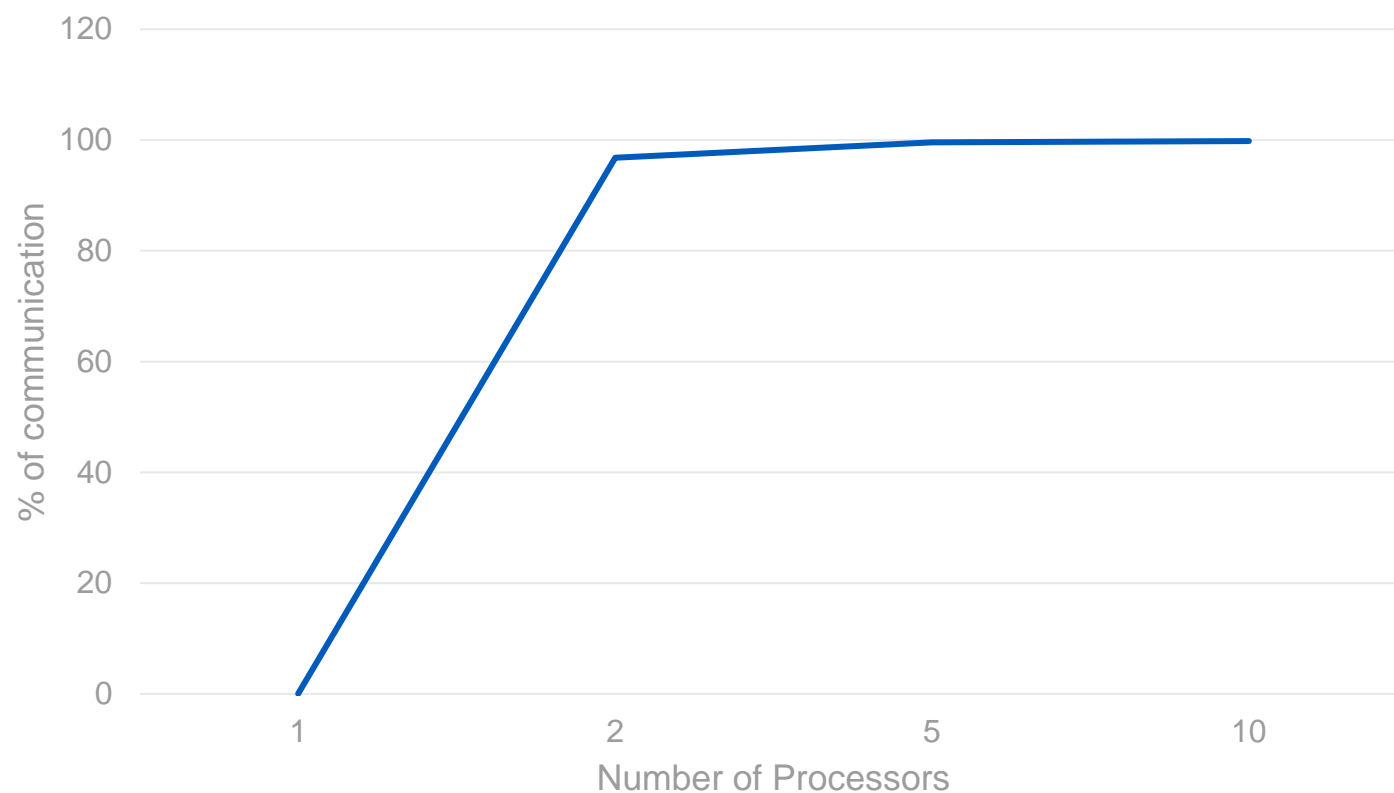
Jacobi Method – Profiling Results - % of communication – 5 Processors

unknown <ul style="list-style-type: none"> Load Balance Communication Balance Message Buffer Sizes Communication Topology Switch Traffic Memory Usage Executable Info Host List Environment Developer Info <p>powered by IPM</p>			command: ./mpi_jacobi 1000					
			codename:	unknown	state:	unknown		
			username:		group:			
			host:		mpi_tasks:	5 on 1 hosts		
			start:		wallclock:	1.69974e+03 sec		
			stop:		%comm:	99.5641686375563		
			total memory:	0.448189 gbytes	total gflop/sec:	0		
			switch(send):	0 gbytes	switch(recv):	0 gbytes		
			Computation			Communication		
			Event	Count	Pop	% of MPI Time		
			 <ul style="list-style-type: none"> ■ MPI_Allgather ■ MPI_Comm_rank ■ MPI_Comm_size 					

Jacobi Method – Profiling Results - % of communication – 10 Processors

unknown <ul style="list-style-type: none"> • Load Balance • Communication Balance • Message Buffer Sizes • Communication Topology • Switch Traffic • Memory Usage • Executable Info • Host List • Environment • Developer Info <p>powered by IPM</p>			command: ./mpi_jacobi 1000					
			codename:	unknown	state:	unknown		
			username:		group:			
			host:		mpi_tasks:	10 on 1 hosts		
			start:		wallclock:	5.19544e+03 sec		
			stop:		%comm:	99.7983038972638		
			total memory:	0.908165 gbytes	total gflop/sec:	0		
			switch(send):	0 gbytes	switch(recv):	0 gbytes		
			Computation			Communication		
			Event	Count	Pop	% of MPI Time		
			 <ul style="list-style-type: none"> ■ MPI_Allgather ■ MPI_Comm_rank ■ MPI_Comm_size 					

Jacobi Method – Profiling Results - % of communication



References

1. Introduction to Running Computations on the High Performance Clusters at the Center for Computational Research – L. Shawn Matott, Center for Computational Research University at Buffalo, SUNY
2. Wilkinson, Allen - Parallel Programming Techniques and Applications using Networked Workstations and Parallel Computers, 2nd Edition, Section 11.3, 11.4.
3. Snir, Gropp – MPI, The Complete Reference, Section 2.5
4. Karniadakis, Kirby – Parallel and Scientific Computing in C++ and MPI, Section 7.2
5. Edmond Jajaga and Jolanda Klobocishta - MPI Parallel Implementation of Jacobi, ICT Innovations 2012 Web Proceedings
6. <https://www.mpi-forum.org/docs/mpi-2.2/mpi22-report/node99.htm#Node99>
7. <https://www.mcs.anl.gov/research/projects/mpi/mpi-standard/mpi-report-2.0/node145.htm>

THANK YOU!

