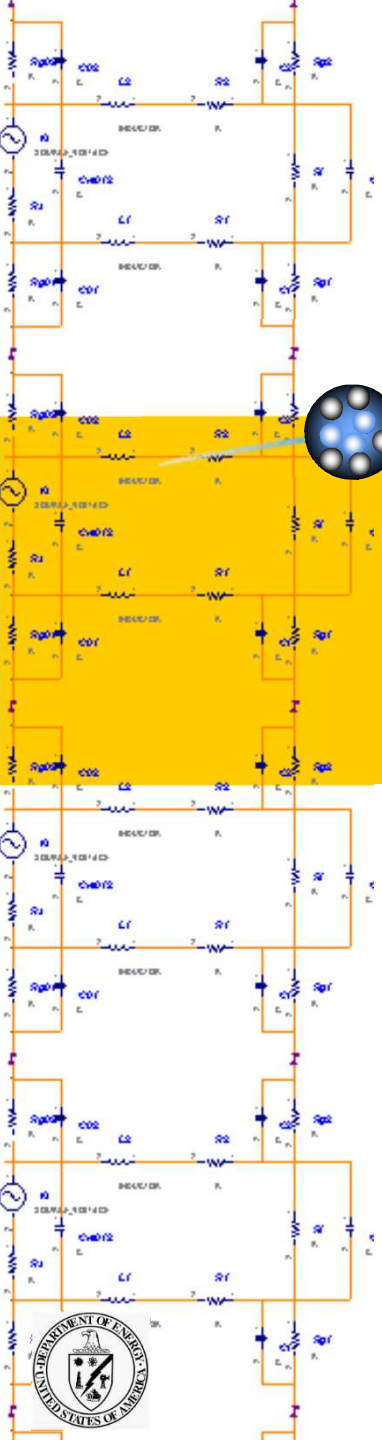

Xyce™

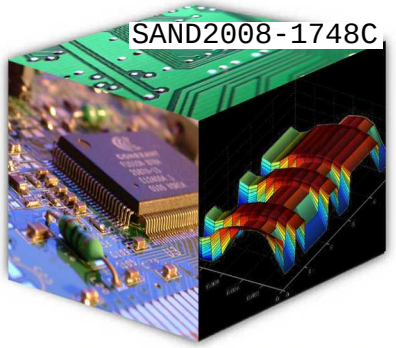

Parallel Partitioning and Preconditioning in Circuit Simulation

SIAM Parallel Processing

Feb 13, 2008

Heidi Thornquist, David Day, Erik Boman, Robert Hoekstra

Sandia National Laboratories
Albuquerque, New Mexico

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

Sandia National Laboratories

Outline

- Circuit Simulation
- Linear System Structure
- Partitioning
- Preconditioning

Electrical Simulation

- Circuit simulation is applied at many levels of detail:
  - Device (PDE) ← CHARON
  - Analog (ODE/DAE) ← Xyce™
  - Digital (VHDL)
  - Co-Simulation (Circuit + Software)
- Analog simulation models network(s) of devices coupled via Kirchoff's current and voltage laws

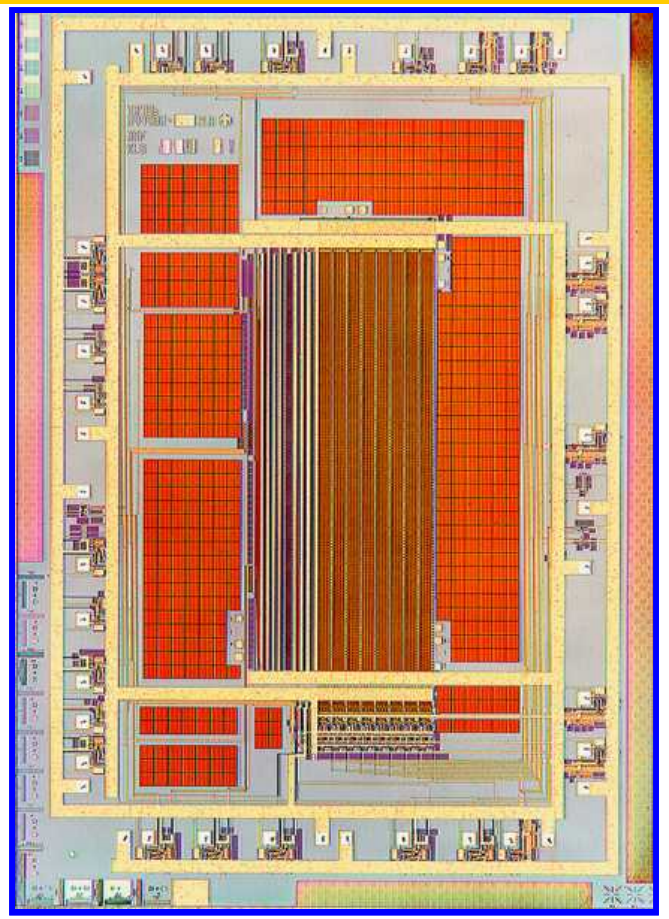

$$KCL_1 = I_1 - I_2 - I_3 = 0$$

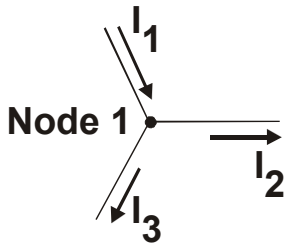

SNL Applications

- Need for high fidelity for very large scale circuits
  - Addition of environmental effects such as radiation requiring additional physics which impact all devices
- Currently, at millions of unknowns with 10s of millions approaching quickly
- Coupling to PDE simulations of a few semiconductor devices (XYCE←→CHARON)

Connectivity

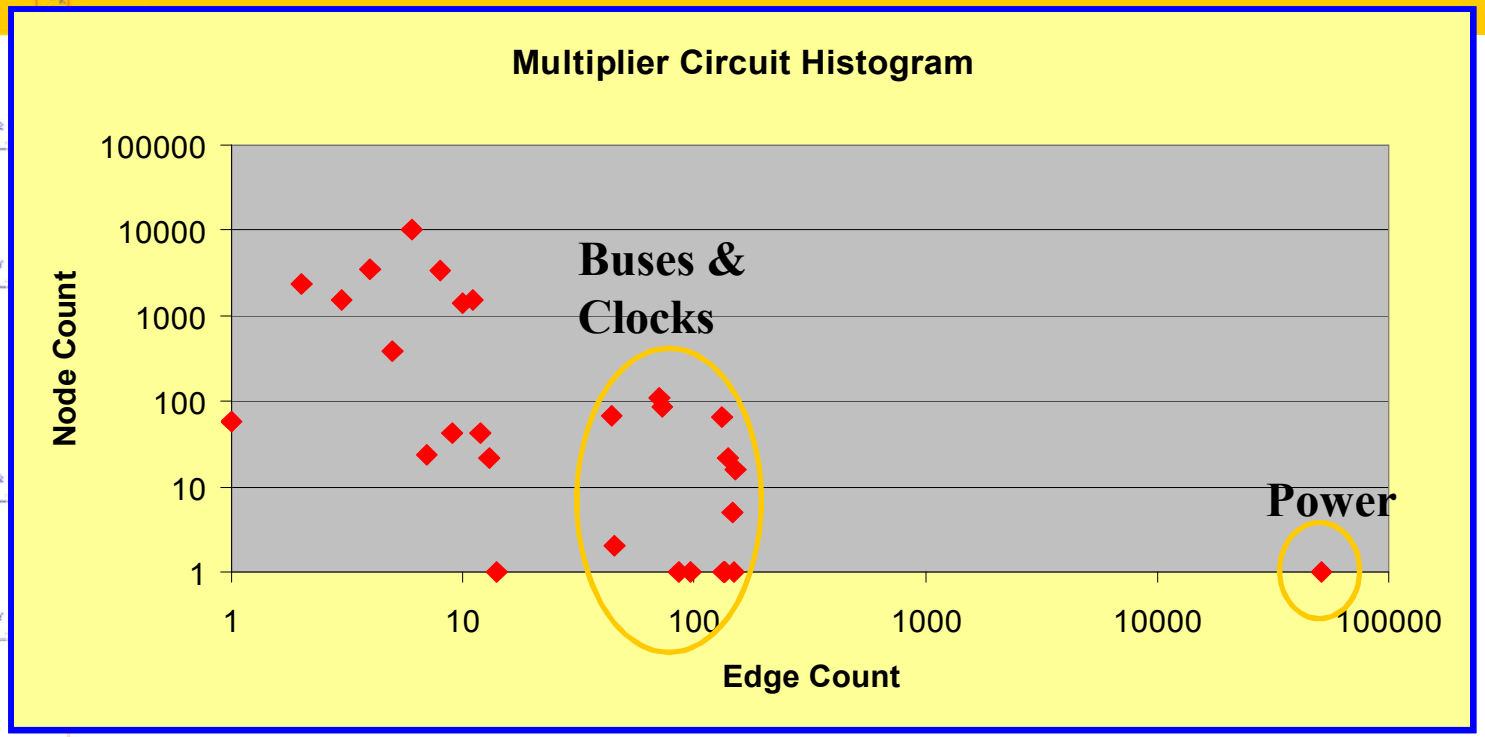

- Digital Circuits:
  - Power node generates very dense row (~0.9*N)
  - Bus lines and clock paths generate order of magnitude increases in bandwidth

Current Strategy

- Trilinos: Epetra, IfPack, AztecOO, Belos
- Strategy : GMRES
  - Overlapping Additive Schwarz
  - Singleton Filtering → DENSE ROWS
  - Zoltan/ParMETIS Partitioning
  - AMD/RCM Block Reordering
  - Row/Col Sum Scaling

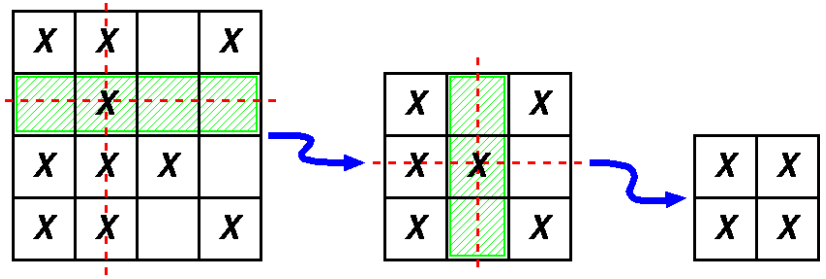

Singleton Filter → Partition LinSys → RCM → Scale

N	Total Cuts	CondEst	GMRES Iters	LinSolve Time	Newton Steps
1220	~1000	3.00E+05	500	4.7	72
1054	68	1.00E+04	127	0.43	53

Block Triangular Form

- Steady State Analog Circuit Problems are Block "REDUCIBLE"!
- BTS Algorithm: $O(n_b \cdot s_b^3 + n_b^2 \cdot s_b^2)$
- Current implementation beats our fastest sparse direct solver for n>10,000
- Ill-conditioned (>10^{16}) diagonal blocks can be better managed usin SVD with Thresholding.

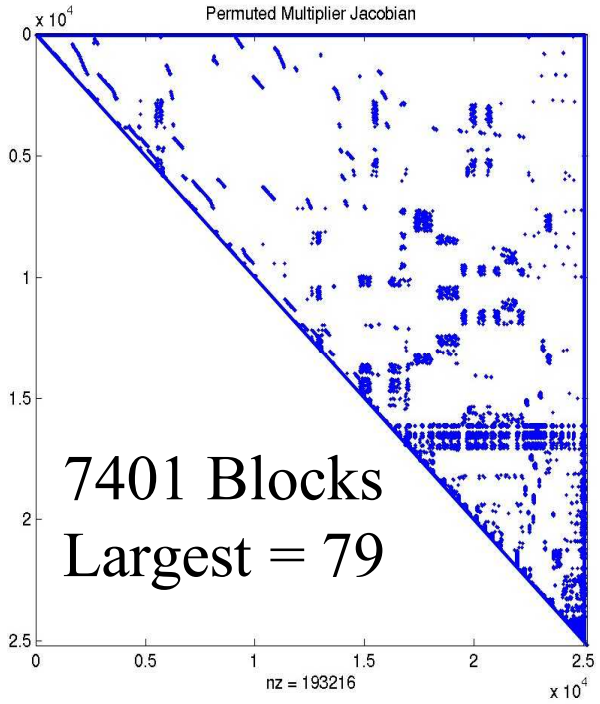

Sparse Direct Solves

- Tim Davis's KLU in AMESOS (The "Clark Kent" of Direct Solvers)
  - Gilbert/Peierel's Left-Looking Sparse LU
  - No SuperNoding since there is little or no benefit for circuit problems
  - Fastest direct solver for Xyce circuits
  - Block Triangular Factorization
  - Distributed Memory Solver
    - PARAKLETE

Method	1st Fac	Next Fac	LU nnz (10^3)	MFlops
UmfPack	.440	.257	372	2.56
KSparse	84	14	294	0.90
KLU w/ BTF	.545	.02	201	0.16

Improving Parallel Performance/Robustness

- Hypergraph Partitioning
- Global Ordering based on BTF

Zoltan Hypergraph Partitioner

- Accessed through the Trilinos Issoropia Package
- ~680k Unknown Xyce Circuit Problem (IC)

	Parmetis			PHG/Hypergraph		
Num Procs	Max Nbors	Avg Nbors	Comm Volume	Max Nbors	Avg Nbors	Comm Volume
16	15	12.5	35278	15	11.8	10121
64	53	33.0	56489	40	19.6	22801

Zoltan Hypergraph Partitioner Future

- 2D Even Better
- Minimize Number of Neighbors < sqrt(p)

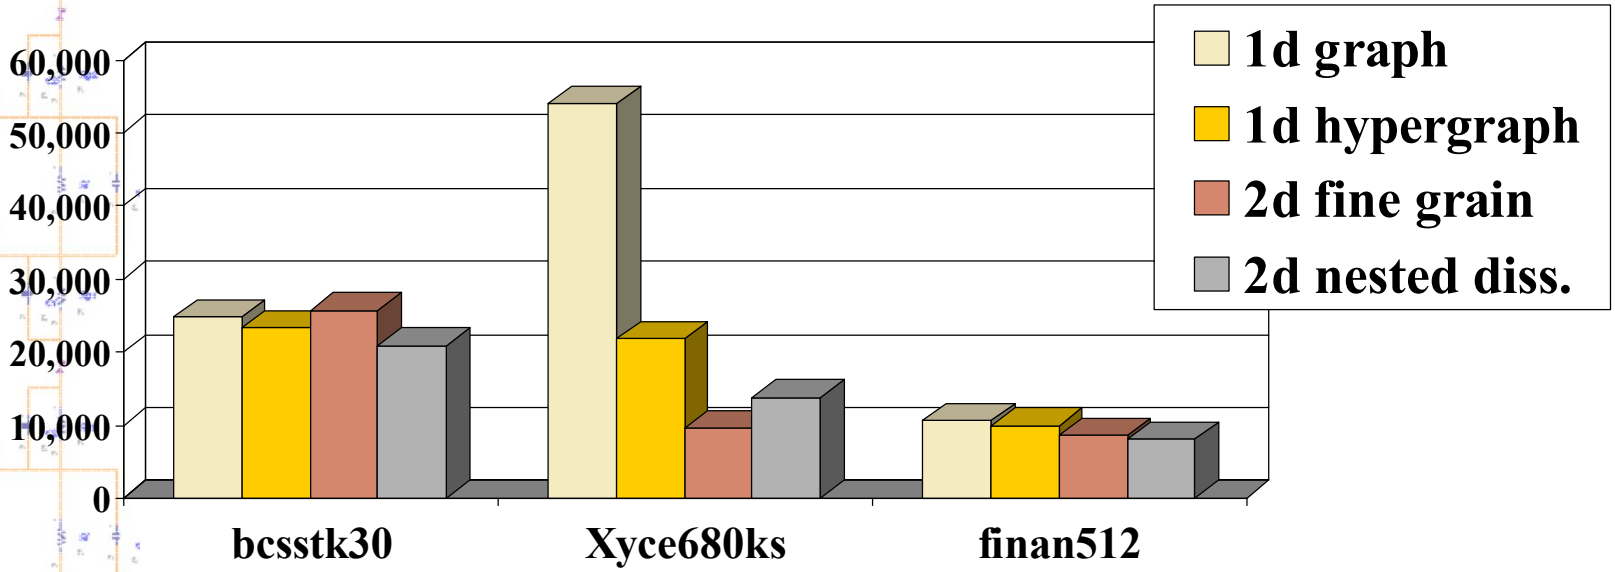

New Preconditioner

- Global Block Triangular Form Reordering
- Issoropia/Zoltan Hypergraph Partitioning

MEM_PLUS Circuit

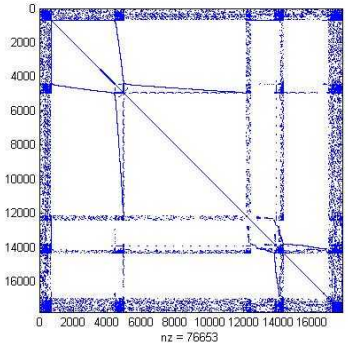

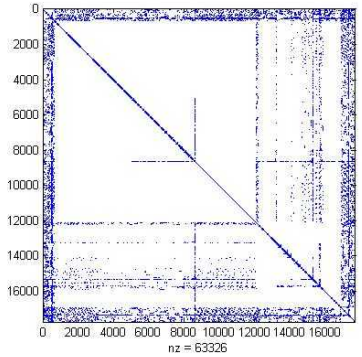

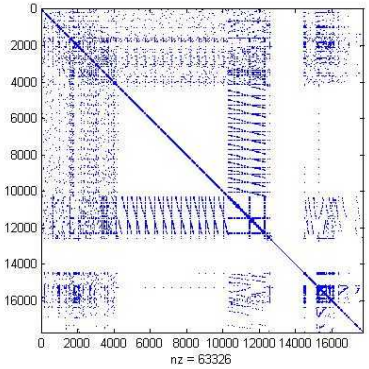

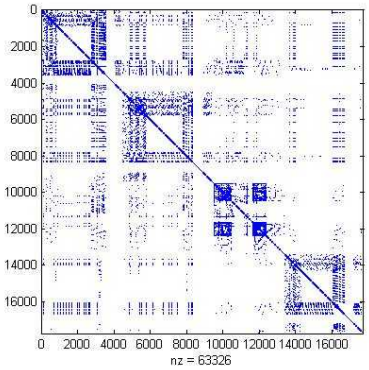

Method	Residual	GMRES Iters	Solver Time
Local AMD ILUT ParMETIS	2.245e-05 (FAIL)	500	22.36
BTF ILUT Linear	1.149e-12	500	16.79
BTF KLU Linear	5.297e-13	92	0.830
BTF KLU Hypergraph	1.451e-13	26	0.139

10K Circuit

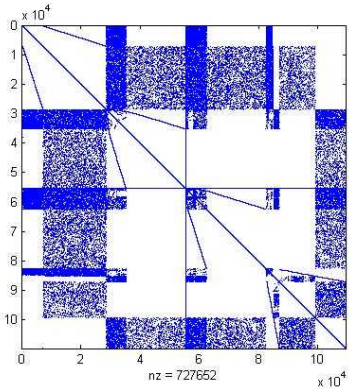

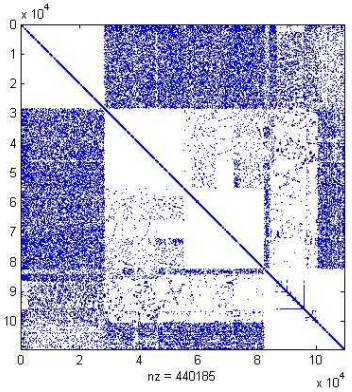

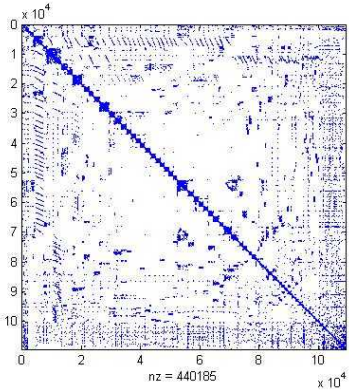

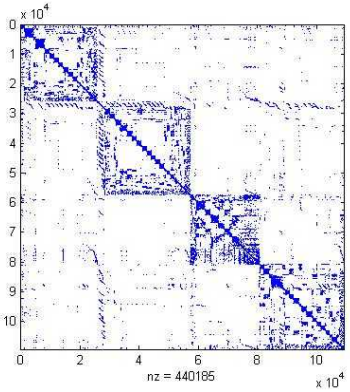

Method	Residual	GMRES Iters	Solver Time (seconds)
Local AMD ILUT ParMETIS	3.425e-01	500	302.573
BTF KLU Hypergraph	3.473e-10	3	0.139

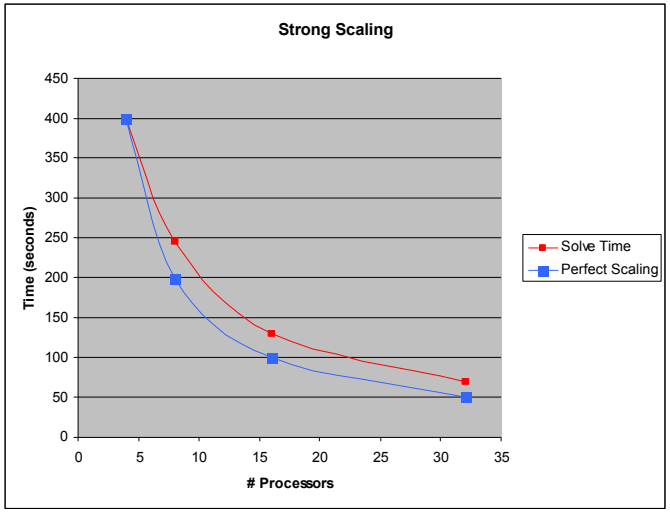

RAM2K

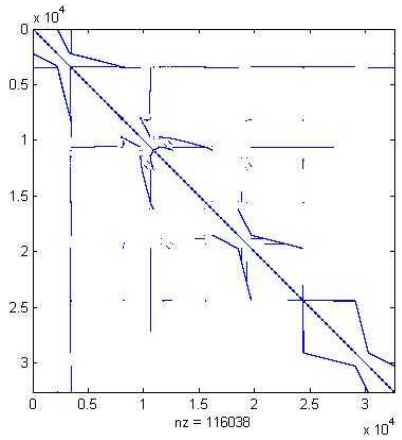

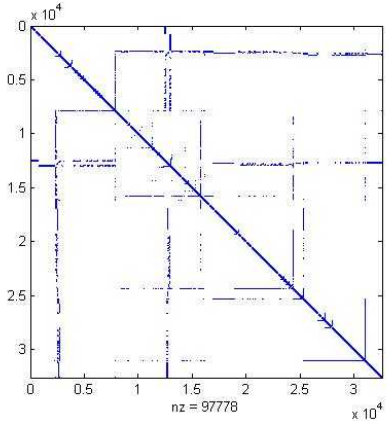

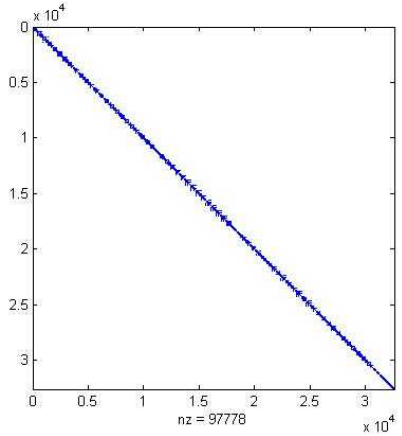

Method	Residual	GMRES Iters	Solver Time
Local AMD ILUT ParMETIS	1.730e-7	500	34.09
BTF ILUT Linear	9.044e-9	195	6.000
BTF KLU Hypergraph	8.469e-10	6	0055

CHIP2

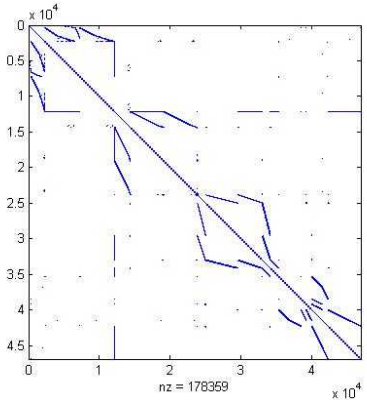

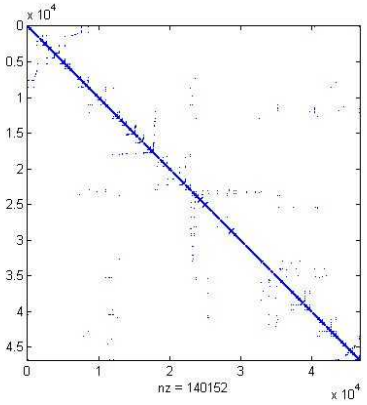

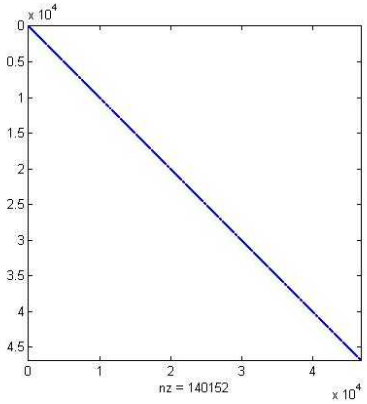

Method	Residual	GMRES Iters	Solver Time
Local AMD ILUT ParMETIS	3.313e-10	500	68.38
BTF ILUT Linear	9.666e-13	292	22.99
BTF KLU Hypergraph	6.119e-17	5	0.063

Conclusions

- Analog circuit problem sizes continue to increase
  - Parasitics
  - Environmental Effects
  - Multi-Time PDE Methods
- Robust distributed memory linear solver technology
  - Ill-Conditioning
  - Partitioning (HyperGraph → Zoltan)
  - MultiLevel Preconditioning

MultiLevel Preconditioning
Distributed Schur Complement

- Benefits
  - Scalability
  - Rapid Decrease in Coarse Mesh Size
- Demonstrated as effective for circuits by Basermann (NEC Europe) and Sosonkina (pARMS)

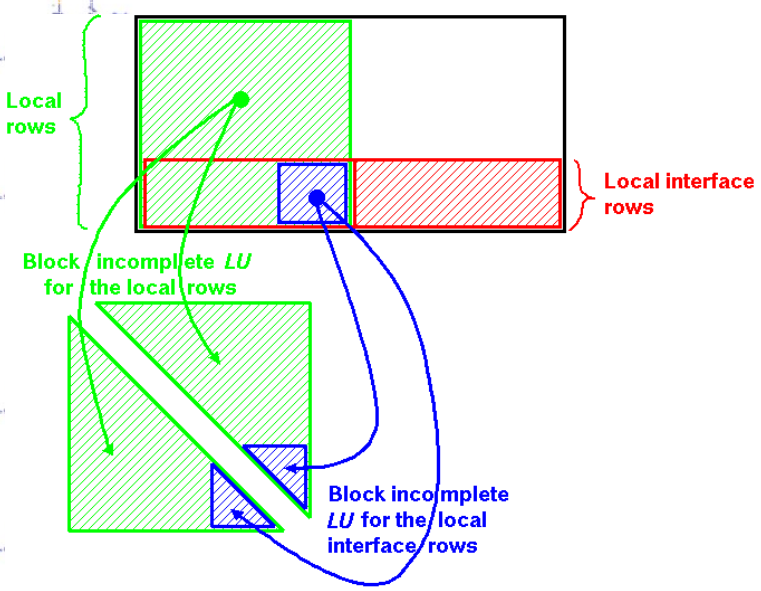

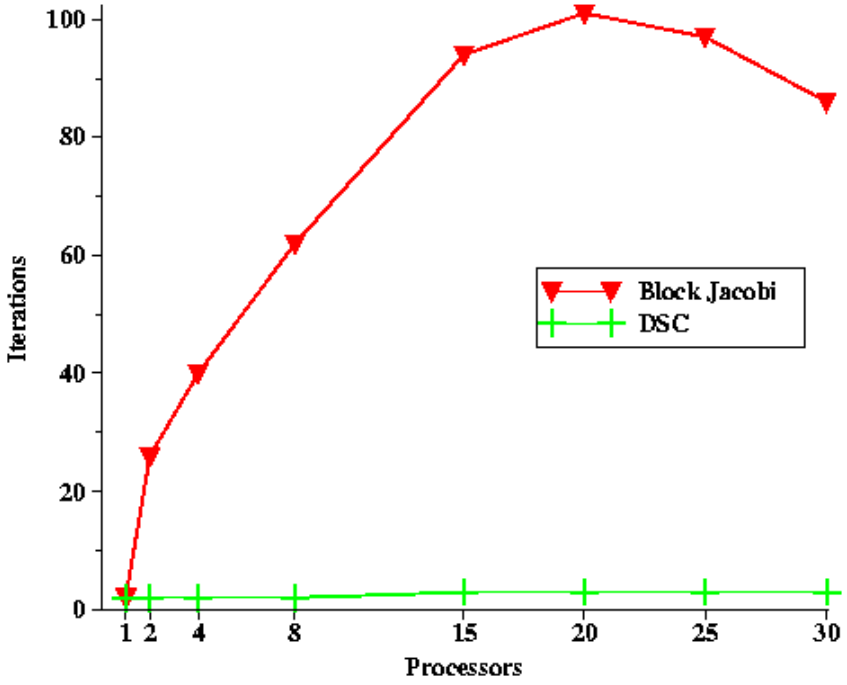

Large Memory Circuit, A. Basermann (NEC Europe)

Distributed Schur Complement Preconditioning

$$A = \begin{pmatrix} B & F \\ E & C \end{pmatrix} = \begin{pmatrix} B & 0 \\ E & S \end{pmatrix} \begin{pmatrix} I & B^{-1}F \\ 0 & I \end{pmatrix}$$

$$S = C - EB^{-1}F$$

$$LU = \begin{pmatrix} L_B & 0 \\ E U_B^{-1} & L_S \end{pmatrix} \begin{pmatrix} U_B & L_B^{-1} F \\ 0 & U_S \end{pmatrix}$$

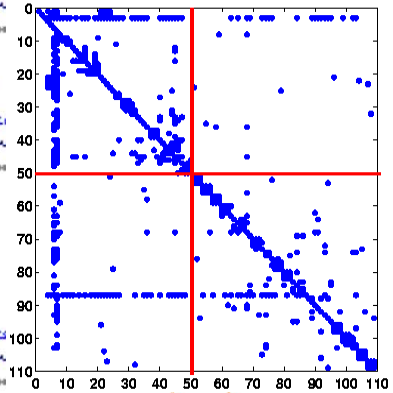

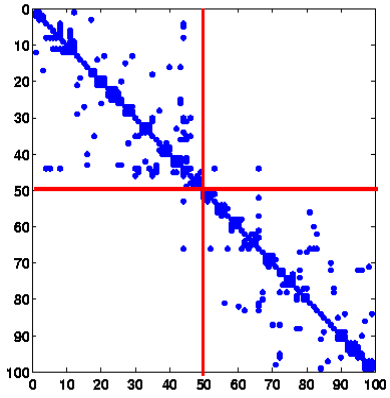

Original matrix Matrix without global nodes

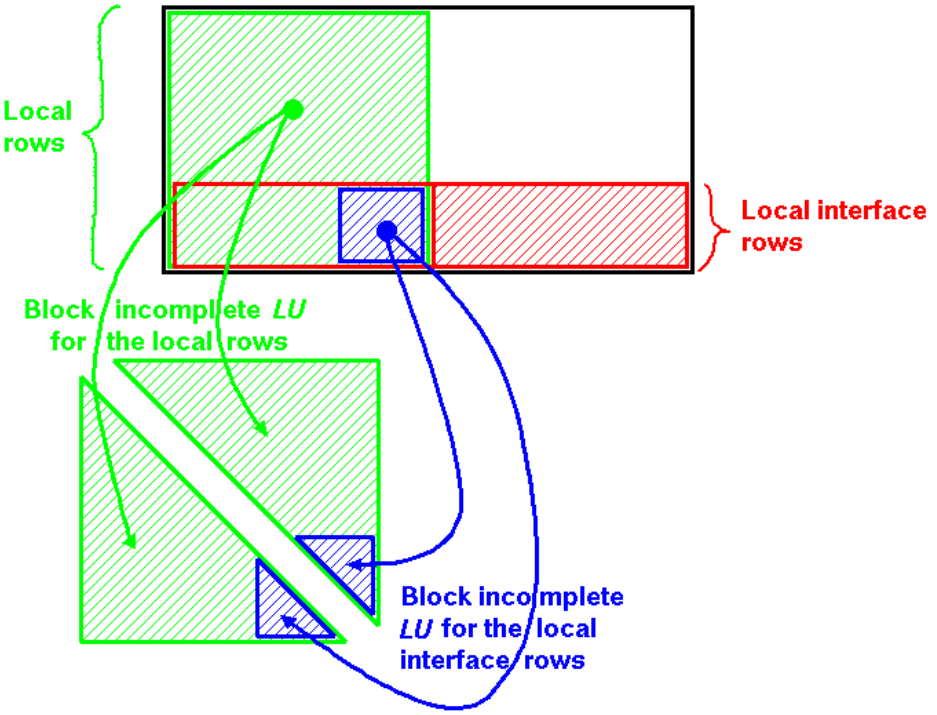

Saad and Sosonkina

Singletons

Number of…	Original Problem	After Filtering
Successful Steps Taken	925	892
Failed Steps Attempted	67	53
Jacobians Evaluated	2488	2399
Linear Solves	2488	2399
Failed Linear Solves	32	13
Linear Solver Iterations	33595	12152

- Row Singleton
  - Pre-Process

$$\begin{bmatrix} & & a_{1j} & & \\ & & a_{2j} & & \\ & & \vdots & & \\ 0 & \cdots & a_{ij} & \cdots & 0 \\ & & \vdots & & \\ & & a_{nj} & & \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_j \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_i \\ \vdots \\ b_n \end{bmatrix} \Rightarrow x_j = b_i / a_{ij}$$

- Col Singleton
  - Post-Process

$$\begin{bmatrix} & & 0 & & \\ & & 0 & & \\ & & \vdots & & \\ & & 0 & & \\ a_{i1} & \cdots & a_{ij} & \cdots & a_{in} \\ & & 0 & & \\ & & 0 & & \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_j \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_i \\ \vdots \\ b_n \end{bmatrix} \Rightarrow x_j = (b_i - \sum_{k \neq j} a_{ik} x_k) / a_{ij}$$

Voltage Source

- Independent voltage sources require current variables.
- Current through a voltage source not a function of voltage drop across the source (no I = V/R relationship!).
- Stencil: Diagonal vanishes

$$\mathbf{J}\qquad \mathbf{\Delta X} \;=\; -\mathbf{f}$$

$$\begin{bmatrix} & & 1 \\ & & -1 \\ 1 & -1 & \end{bmatrix} \begin{bmatrix} \Delta V_1 \\ \Delta V_2 \\ \Delta I_V \end{bmatrix} = \begin{bmatrix} -I_V \\ I_V \\ -V_1 + V_2 + V_d \end{bmatrix} \begin{matrix} KCL_1 \\ KCL_2 \\ Constr \end{matrix}$$

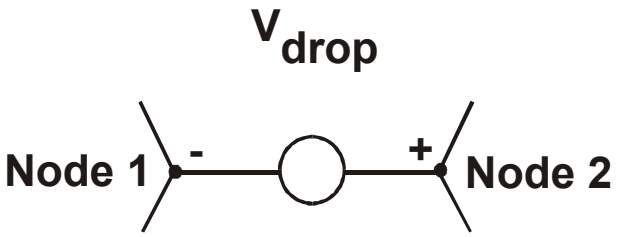

Scaled Speedup

- Analog circuit simulation on largest number of processors
  - 14 million devices
  - 1024 processors
- Details
  - Scaled microStrip Transmission Line
  - AztecOO GMRES/ILUT
    - ~10 iters/solve
  - NOX Full Newton
  - Sinusoidal Input
  - 14,000 Dev/Proc
  - 6,000 Unk/Proc

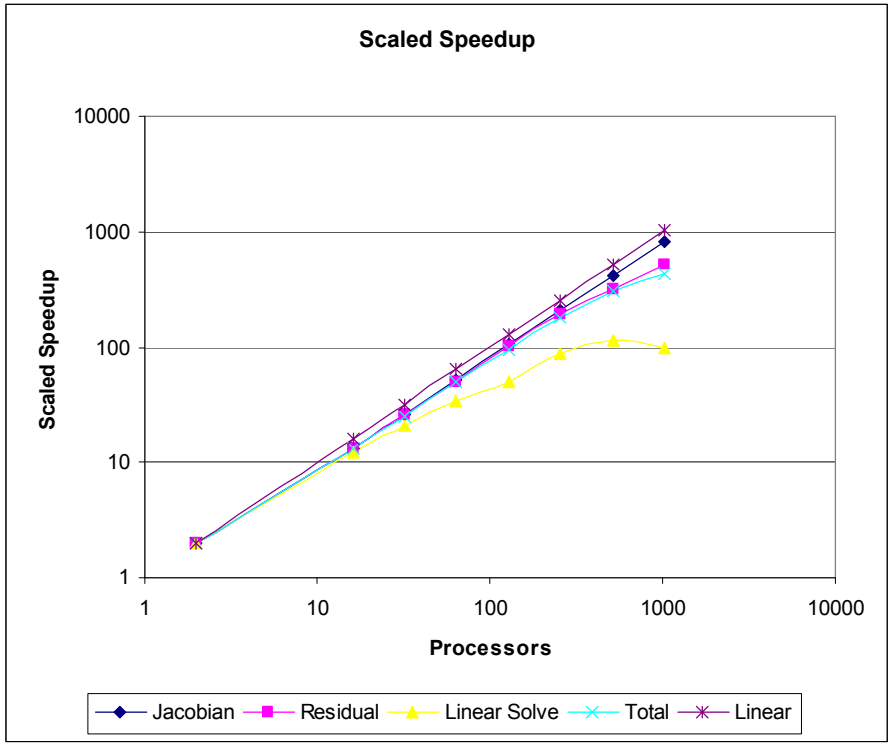

Fixed Size Speedup

- Digital Transistor Circuit
  - 0.75 Million Devices
  - >2 Million Unknowns
- Details
  - AztecOO GMRES
    - Singleton Filter
    - Zoltan/ParMETIS
    - AMD Subdomains
    - Ifpack ILU(K)
      - Fill Factor: 1
      - Overlap : 0
      - Diagonal Shift
  - NOX Full Newton

Procs	Time Total	Time Solve	Time Fill	Efficiency	Adj (Avg)	Adj (Max)	Scaled Speedup
16	8173	4132	2696	1	13.375	15	1
32	5536	2943	1353	0.738168	23.3125	31	1.48
64	3500	1900	725	0.583786	30.5	56	2.34
128	2220	1243	397	0.460191	36.78125	80	3.68