Early Attempts for FEM Miniapp Validation: Very Preliminary Study for a Semiconductor Device Simulator

Miniapplication Validation Workshop

SNL

August 24, 2010

Paul Lin, Richard Barrett, Mike Heroux

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

Response of Electronics to Radiation

- Radiation Analysis, Modeling and Simulation for Electrical Systems (RAMSES) code suite
  - Charon: semiconductor device simulator
  - Xychron: Charon coupled with Xyce (circuit modeling; next talk)
- QASPR: How well do electronics survive in radiation environments?
- Charon: Drift-diffusion model for semiconductor devices
  - QASPR: Hennigan, Hoekstra, Castro, Fixel, Pawlowski, Lin, etc.
- Charon: Resistive magnetohydrodynamics (MHD)
  - Fusion: tokamak, Z-pinch
  - DOE ASCR effort: Shadid, Pawlowski, Cyr, etc.

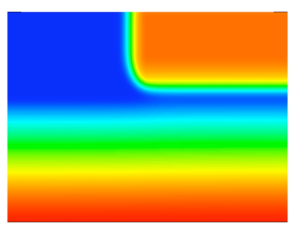

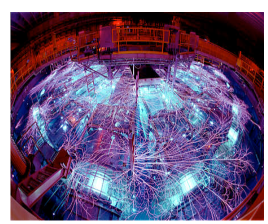

ASC™
DOE / NNSA

Sandia National Laboratories

SNL Semiconductor Device Simulations Require Substantial Resources: Reducing Time Critical

- 2D bipolar junction transistor (BJT) with full defect physics O(10^7 - 10^8) DOF; takes O(week) on O(10^3) cores
- 3D simulations? O(10^9 – 10^{10}) unknowns
- Prediction plus uncertainly required for validation requires ensemble of calculations
  - 1D simulations presently; O(10^3) simulations
  - 2D could be performed on current largest platforms (e.g. couple weeks on entire Jaguar)
  - 3D simulations?

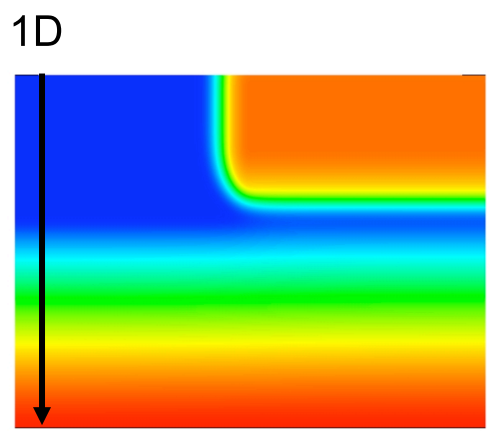

Hope: to use MiniFE to test proposed algorithmic improvements for Charon

Semiconductor Drift-Diffusion Model

Electric potential $-\nabla \cdot \epsilon \nabla \psi = q\,(p - n + C)$

$$\nabla \cdot \mathbf{J}_n - qR = q\frac{\partial n}{\partial t} \qquad \mathbf{J}_n = -qn\mu_n \nabla \psi + qD_n \nabla n$$

$$-\nabla \cdot \mathbf{J}_p - qR = q\frac{\partial p}{\partial t} \qquad \mathbf{J}_p = -qp\mu_p \nabla \psi - qD_p \nabla p$$

Each additional species adds an additional equation
(also modifies equation for electric potential)

$$-\nabla \cdot \mathbf{J}_i - q_i R_i = q_i \frac{\partial X_i}{\partial t} \qquad \mathbf{J}_i = -q_i \mu_i X_i \nabla \psi - q_i D_i \nabla X_i \qquad \mu_i = \frac{q_i D_i}{kT}$$

$$-\nabla \cdot \epsilon \nabla \psi = q\,(p - n + C) + \sum_{i=1}^{n} q_i X_i \qquad q_i \equiv Z_i q$$

(With Hennigan, Hoekstra, Castro, Fixel, Pawlowski, Phipps, Musson, T. Smith)

Sandia National Laboratories

Charon

- Stabilized FEM and FVM discretization for Drift-Difffusion
- Unstructured meshes
- Fully-implicit Newton-Krylov solver; usually GMRES
- Fully-coupled approach has advantages for complex physics, but requires efficient solution of large sparse linear systems
- Trilinos for nonlinear solver (NOX), Krylov solver (AztecOO), preconditioner (ML and Ifpack)
- Sacado for AD (for Jacobian construction/fill)
- Uses Nevada framework

- Currently MPI-only

Communication and Computation: Preconditioned Krylov Solver

- Depends on choice of Krylov solver and preconditioner
- Computation
  - Lots of dot products, mat-vec, waxpy
  - ML also has mat-mat, apply ILU factors, KLU
- Communication
  - Nearest neighbor boundary information
  - Global reductions
  - ML communication gets ugly fast
    - multiple levels
    - restriction/prolongation
    - serial coarse solve

Weak Scaling Study: 1-level and 3-level 2D BJT Steady-State Drift-Diffusion

- Charon FEM semiconductor device modeling code
- 3-level AMG preconditioner (ML library): NSA and PGSA
- “Time”: construct preconditioner and perform linear solve

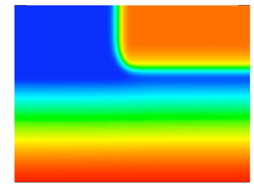

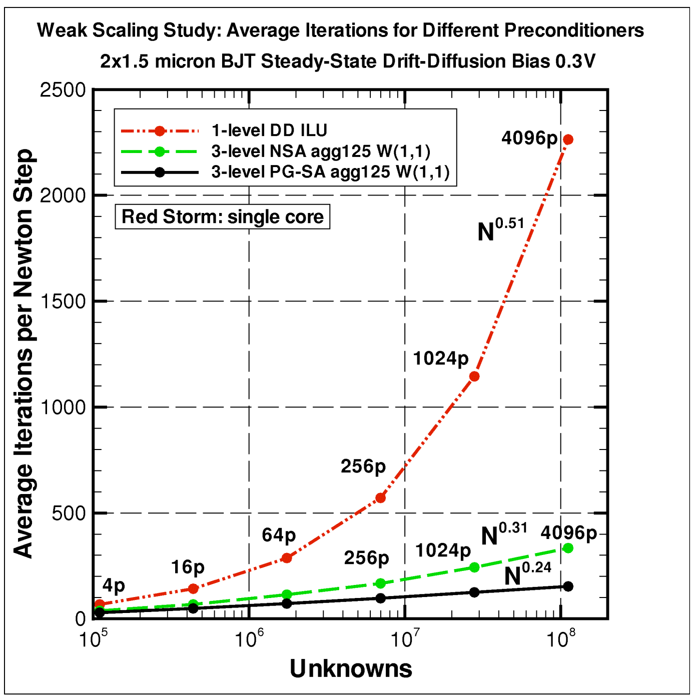

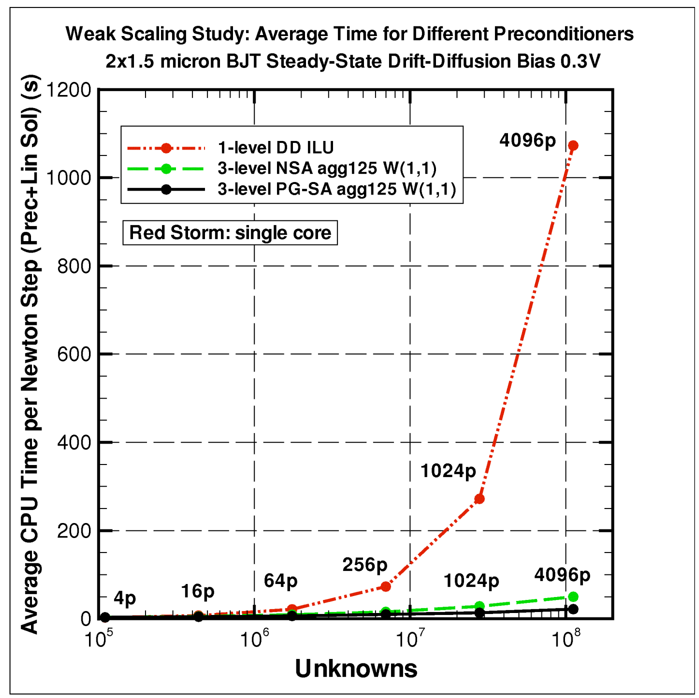

- PGSA 2.3 times faster than NSA
- PGSA ~50 times faster than 1-level

Weak Scaling to 65536 Cores on Blue Gene

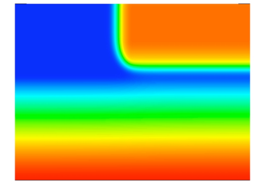

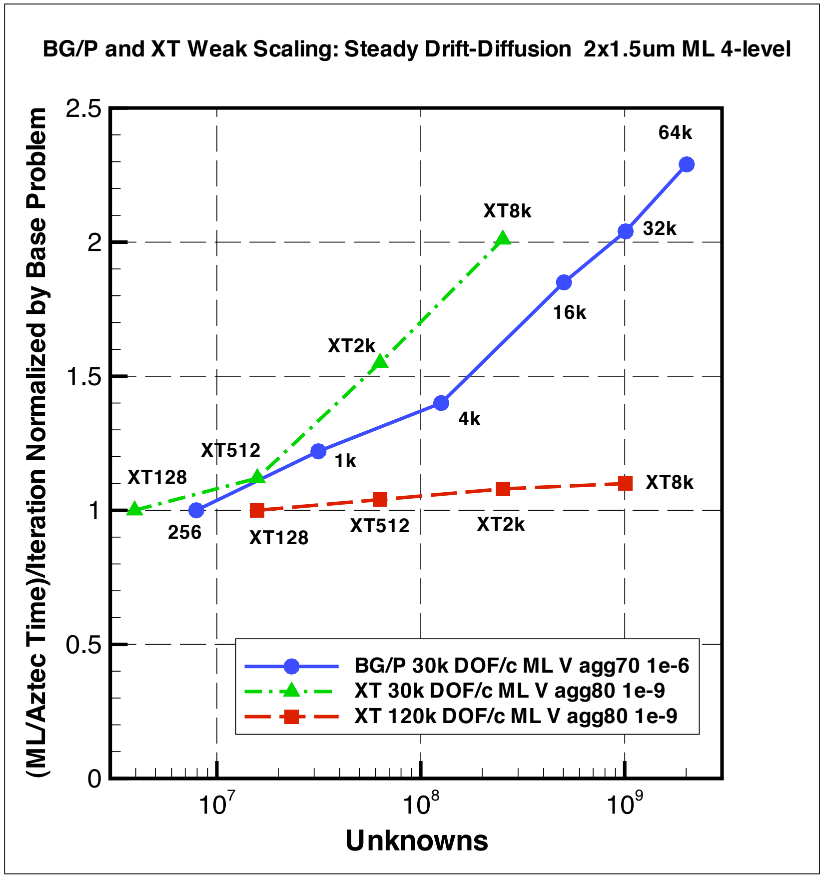

- FEM with fully implicit Newton-Krylov solver
- BJT steady-state drift-diffusion
- Problem sized increased by factor of 256 to two billion DOF on 65536 cores
- Used all four cores per BG/P node; 30k DOF/core
- TFQMR linear solver with ML PGSA 4-level
- Comparison with 30k and 120k DOF/core for Cray XT3/4: better scaling with increased work
- 2 billion DOF problem successfully run on 100k cores

Sandia National Laboratories

Preliminary Multicore Efficiency: Single Node (Quad-core CPUs)

(with J. Shadid)

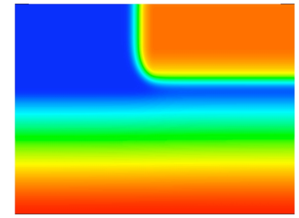

Quad socket/quad core 2.2 GHz AMD Barcelona

Weak scaling: 28K DOF/core

core	DOF	linear sys solve		Jacobian		total	
		time(s)	η	time(s)	η	time(s)	η
1	28K	9.71	Ref	3.52	Ref	14.6	Ref
4	110K	10.7	91	3.48	1.01	15.4	94
8	219K	11.6	84	3.45	1.02	16.3	89
12	329K	13.2	74	3.46	1.02	17.9	81
16	438K	15.8	61	3.13	1.12	20.1	73

- Time per Newton step
- Linear solve time (prec setup and ML/Aztec) efficiencies problematic
- Charon performance significantly affected by memory BW

Sandia National Laboratories

Preliminary Multicore Efficiency: Single Node (dual-socket/6-core CPUs)

(with J. Shadid)

Weak scaling: 28k DOF/core (2.6 GHz AMD); time per Newton step

core	DOF	linear sys solve		Jacobian		total	
		time(s)	η	time(s)	η	time(s)	η
1	28K	5.38	Ref	2.46	Ref	8.72	Ref
2	55K	5.83	92	2.46	100	9.19	95
4	110K	6.78	79	2.50	98	10.2	86
6	165K	7.65	70	2.55	96	11.1	78
8	219K	8.78	61	2.52	98	12.2	71
10	273K	9.77	55	2.52	98	13.2	66
12	329K	10.97	49	2.55	96	14.5	60

- Linear solve time (prec setup and ML/Aztec) efficiencies problematic
- Charon performance significantly affected by memory BW
- Need block methods (e.g. VBR) for more efficient memory access
- Soon will need hybrid approach: MPI/threading (Trilinos Kokkos)

Preliminary Multicore Efficiency Study: Network and Nodes

(with J. Shadid)

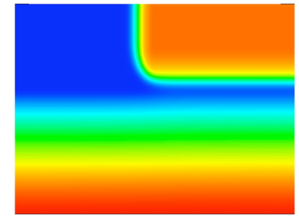

- Combines effects of network and node architecture: vary nodes and cores/node for total of 128 cores

configuration	54.5K DOF/core		218K DOF/core	
	time(s)	η	time(s)	η
128n 1ppn	25.9	Ref	147	Ref
32n 4ppn	26.0	100	152	97
16n 8ppn	27.1	96	163	90
10.5n 12ppn	30.3	86	194	76
8n 16ppn	35.5	73	229	64

Quad-socket/quad-core 2.2 GHz AMD compute nodes; InfiniBand

- Used all 16 cores per node
- 218K DOF/core case tries to maximize contention for memory BW

A miniapp that can be predictive for Charon is vital

- Charon is a large code with many TPLs
  - Charon/nevada ~700,000 lines code
  - Nevada TPLs
  - Charon TPLs (biggest TPL is Trilinos)
- Rewriting Charon to test new ideas can be extremely time consuming and painful
- Charon can be very painful to port
  - Compilers on massively parallel platforms tend to have issues with C++, especially templating (tends to trigger compiler bugs)
  - Horror stories:
    - ~6 months to port to PGI on Red Storm
    - ~4 months to port to IBM XL on Blue Gene
- Just to recompile Charon to test different compiler flags and optimizations is time consuming

Sandia National Laboratories

Miniapp: MiniFE

- Solves the steady-state 3D heat equation (Poisson equation)
- Geometry is a cube
- Finite element method with hexahedral elements
- Symmetric matrix solved by CG (no preconditioner**)**

Does MiniFE Predict Charon Behavior? Processor Ranking: 8 MPI tasks; 31k DOF/core

- Charon steady-state drift-diffusion BJT
- Nehalem (Intel 11.0.081 –O2 –xsse4.2; all cores of dual-socket quadcore)
- 12-core Magny-Cours (Intel 11.0.081 –O2; one socket, 4 MPI tasks/die)
- Barcelona (Intel 11.1.064 –O2; use two sockets out of the quad-socket)
- 2D Charon (3 DOF/node) vs. 3D MiniFE; match DOF/core and NNZ in matrix row
- Charon LS w/o or w/ ps: GMRES linear solve without/with ML precond setup time
- Try to compare MiniFE “assembling FE”+”imposing BC” time with Charon equivalent

MiniFE

	CG	FE assem+BC
1	Nehalem	Nehalem
2	MC(1.7)	MC(1.7)
3	Barc(2.7)	Barc(1.8)

Charon

	LS w/o ps	LS w/ ps	Mat+RHS
1	Nehalem	Nehalem	Nehalem
2	MC(1.7)	MC(1.8)	MC(1.46)
3	Barc(2.8)	Barc(2.5)	Barc(1.52)

Number in parenthesis is factor greater than #1 time

Sandia National Laboratories

Does MiniFE Predict Charon Behavior? Processor Ranking: 8 MPI tasks; 124k DOF/core

- Charon steady-state drift-diffusion BJT
- Nehalem (Intel 11.0.081 –O2 –xsse4.2; all cores of dual-socket quadcore)
- 12-core Magny-Cours (Intel 11.0.081 –O2; one socket, 4 MPI tasks/die)
- Barcelona (Intel 11.1.064 –O2; use two sockets out of the quad-socket)
- 2D Charon (3 DOF/node) vs. 3D MiniFE; match DOF/core and NNZ in matrix row
- Charon LS w/o or w/ ps: GMRES linear solve without/with ML precond setup time
- Try to compare MiniFE "assembling FE"+"imposing BC" time with Charon equivalent

MiniFE

	CG	FE assem+BC
1	Nehalem	Nehalem
2	MC(1.7)	MC(1.7)
3	Barc(2.6)	Barc(1.8)

Charon

	LS w/o ps	LS w/ ps	Mat+RHS
1	Nehalem	Nehalem	Nehalem
2	MC(1.8)	MC(1.8)	MC(1.47)
3	Barc(3.3)	Barc(3.0)	Barc(1.51)

Number in parenthesis is factor greater than #1 time

MiniFE Predict Charon? Compiler Ranking on Quad-socket Quadcore: 16 tasks; 31k DOF/core

- Quad-socket quadcore Barcelona node; Charon steady-state drift-diffusion BJT
- Intel 11.1.064; PGI 9.0.4; GNU 4.3.4; -O2 for all (all Open MPI 1.4.1)
- 2D Charon (3 DOF/node) vs. 3D MiniFE; match DOF/core and NNZ in matrix row
- Charon LS w/o or w/ ps: GMRES linear solve without/with ML precond setup time
- Try to compare MiniFE “assembling FE”+”imposing BC” time with Charon equivalent

MiniFE

	CG	FE assem+BC
1	Intel	Intel
2	GNU(1.01)	GNU (1.1)
3	PGI(1.04)	PGI (1.8)

Charon

	LS w/o ps	LS w/ ps	Mat+RHS
1	Intel	Intel	Intel
2	GNU(1.02)	GNU(1.01)	GNU(2.5)
3	PGI(1.06)	PGI(1.2)	PGI(3.3)

Number in parenthesis is factor greater than #1 time

MiniFE Predict Charon? Compiler Ranking on Quad-socket Quadcore: 16 tasks; 124k DOF/core

- Quad-socket quadcore Barcelona node; Charon steady-state drift-diffusion BJT
- Intel 11.1.064; PGI 9.0.4; GNU 4.3.4; -O2 for all (all Open MPI 1.4.1)
- 2D Charon (3 DOF/node) vs. 3D MiniFE; match DOF/core and NNZ in matrix row
- Charon LS w/o or w/ ps: GMRES linear solve without/with ML precond setup time
- Try to compare MiniFE “assembling FE”+”imposing BC” time with Charon equivalent

MiniFE

	CG	FE assem+BC
1	Intel	Intel
2	GNU(1.0)	GNU (1.1)
3	PGI(1.02)	PGI (1.9)

Charon

	LS w/o ps	LS w/ ps	Mat+RHS
1	Intel	Intel	Intel
2	GNU(1.01)	GNU(1.01)	GNU(2.5)
3	PGI(1.04)	PGI(1.1)	PGI(3.3)

Number in parenthesis is factor greater than #1 time

MiniFE Predict Charon? Multicore Efficiency Dual-Socket 12-core Magny-Cours : 31k DOF/core

- Charon steady-state drift-diffusion BJT; Intel 11.0.081 –O2
- Weak scaling study with 31k DOF/core
- 2D Charon (3 DOF/node) vs. 3D MiniFE; match DOF/core and NNZ in matrix row
- Efficiency: ratio of 4-core time to n-core time (expressed as percentage)
- Charon LS w/o or w/ ps: GMRES linear solve without/with ML precond setup time
- 100 Krylov iterations for both MiniFE and Charon (100 per Newton step)

MiniFE

cores	CG eff
4	Ref
8	87
12	74
16	64
20	56
24	46

Charon

cores	LS w/o ps eff	LS w/ ps eff
4	Ref	Ref
8	88	89
12	77	80
16	68	72
20	59	64
24	52	57

MiniFE Predict Charon? Multicore Efficiency Dual-Socket 12-core Magny-Cours : 124k DOF/core

- Charon steady-state drift-diffusion BJT; Intel 11.0.081 –O2
- Weak scaling study with 124k DOF/core
- 2D Charon (3 DOF/node) vs. 3D MiniFE; match DOF/core and NNZ in matrix row
- Efficiency: ratio of 4-core time to n-core time (expressed as percentage)
- Charon LS w/o or w/ ps: GMRES linear solve without/with ML precond setup time
- 100 Krylov iterations for both MiniFE and Charon (100 per Newton step)

MiniFE

cores	CG eff
4	Ref
8	89
12	73
16	61
20	54
24	45

Charon

cores	LS w/o ps eff	LS w/ ps eff
4	Ref	Ref
8	87	89
12	74	78
16	61	66
20	49	54
24	40	45

MiniFE vs. Charon

	miniFE	Charon drift-diffusion
Dimensionality	3D	Currently 2D; 3D in the works
PDE	Linear, scalar (1 DOF/node)	Nonlinear, system (3 DOF/node)
Linear system	Symmetric	Nonsymmetric
Krylov solver	CG	GMRES
Preconditioner	None	MG (ML), DD (prec setup time adds to solve time)

Sandia National Laboratories

Future: MiniFE vs. Charon

	miniFE	Charon full defect physics
Time dependence	Steady-state	Transient; stiffness issues
DOF/node	1	39
Source terms	constant	Depends on variable
Nonlinearity	linear	Strongly nonlinear
Linear system	symmetric	Nonsymmetric; likely indefinite

Concluding Remarks

- Very preliminary study comparing trends for miniFE and Charon
- Need a lot more comparisons before can draw conclusions

Thanks For Your Attention!

Paul Lin (ptlin@sandia.gov)

For further information about Charon:

• P Lin, J Shadid, M Sala, R Tuminaro, G Hennigan, R Hoekstra, "Performance of a Parallel Algebraic Multilevel Preconditioner for Stabilized Finite Element Semiconductor Device Modeling," Journal Comp Physics Vol 228 (2009), pp. 6250–6267

• P Lin and J Shadid, "Towards Large-Scale Multi-Socket, Multicore Parallel Simulations: Performance of an MPI-only Semiconductor Device Simulator," Journal Comp Physics Vol 229 (2010), pp. 6804–6818

Sandia National Laboratories