



| April 1, 2011

TAU 2011 Power Grid Analysis Contest

Zhuo Li, Frank Liu, Raju Balasubramanian
and Sani Nassif

Contests in the EDA Area

- **Live Programming Contests**

- Code and find the answer in a short period of time (hours or days)
- Cadathlon

- **ISPD Contests**

- 2005 to 2011. Released industrial benchmarks. Each team given fixed time interval to prepare the final code.
- Helped drive the new age of Physical Design. Academia started to work on placement/global routing/clock synthesis problems with real industry data.
- Provided a public platform to compare different algorithms.
- These are optimization problems, and there is no “golden” solution.

- **Power Grid Analysis Contest**

- ISPD style, research oriented contest.
- First contest in the simulation/analysis area.
- Can we find new problems? Are all old problems solved?
- Are there known optimal solution? (not really).

History and Overview

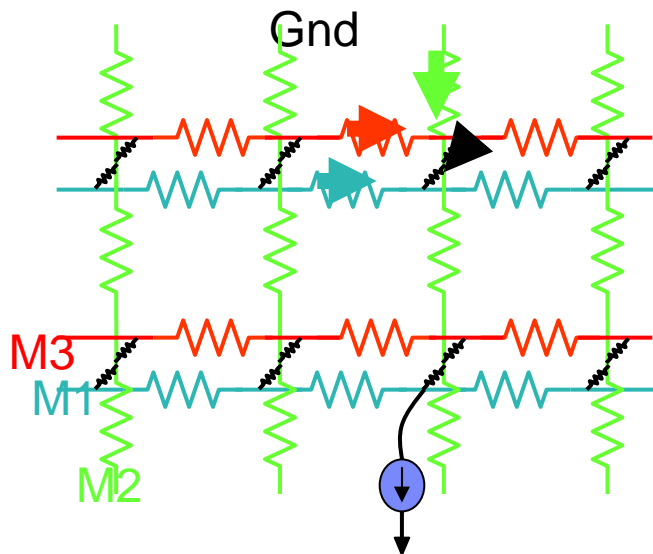
- **Nov. 24, 2010. Signed up to organize the contest. Created taucontest2011@gmail.com.**
- **Dec. 2, 2010. Organization team formed.**
- **Dec. 7, 2010. Initial contest information released. We need participants!**
- **Jan. 11, 2011. 12 teams registered!**
 - 6 teams from U.S., 6 teams from overseas
- **Jan. 12, 2011. New website.**
- **Feb. 11, 2011. 11 alpha executables received.**
- **Mar. 11, 2011. 10 final entries (5 U.S. teams, 5 oversea teams)**
- **Total 11 benchmarks.**
- **QOR metrics: Accuracy, runtime and memory.**
- **Special thanks:**
 - Cliff Sze, Gi-Joon Nam and Natarajan Viswanathan for their experience on ISPD, and webpage templates.
 - Nancy Zhou for solution pictures.

Ten Final Teams

Simulator Name	Affiliation
FIPGS	University of California, Santa Cruz
NTUPG	National Taiwan University, GIEE Department
PGSIM_TDP	Michigan Tech University
IPGS	UIUC
SEVA	Dept. of Computer Science, National Tsing Hua University; Dept. of Electrical and Computer Engineering, Missouri University of Science and Technology
PowerRush	Tsinghua University
TicTac	Texas A&M University
Resgris	INESC ID/Instituto Superior Tecnico/Technical University of Lisbon
Sato Laboratory Power Analyzer	Kyoto University
Random Walk Hybrid Solver	University of Minnesota, IBM's Thomas J. Watson Research Center (*)

Benchmark History

- Sani Nassif, “Power Grid Analysis Benchmarks”, 2008, ASPDAC.
- Released six IBM benchmarks to public. Spice Format.
- Focus on DC problems
 - Power Grid modeled as a resistive mesh.
 - Loads modeled as DC current sources.
 - Package influence modeled as resistive drops in each package connection.
 - DC analysis is the bulk of industrial power delivery analysis applications!




Name	# Nodes	# Elements	Metal Levels
IBMPG1	30638	55109	2
IBMPG2	127238	246581	5
IBMPG3	851584	1603581	5
IBMPG4	953583	1838583	6
IBMPG5	1079310	2156735	3
IBMPG6	1670494	3246725	3

New Benchmarks

- **All new benchmarks are still for DC problems**
- **Released two new benchmarks in 03/02, 1.5M nodes.**
- **Five internal power grid benchmarks from IBM's P7 processor (2M, 5M, 6M, 8M, and 22M node sizes).**
 - Three benchmarks are variations on the Nest VDD network.
 - Two benchmarks are variations on the Core VDD network.
- **Chose ibmpg3 to ibmpg6 from original set as well, to see how teams perform on released netlists (ibmpg1 and ibmpg2 are too small)**
- **Total of 11 benchmarks for evaluation**
 - Surprisingly, some teams still failed for the released benchmarks.
 - And many had some rather trivial parsing problems...
 - Some teams generated core dumps on the internal benchmarks.
 - The ranking for the top 3 teams were consistent between the external and internal benchmarks.

Some Cleanup

- **Generated by an internal solver based on real industrial designs**
- **The spice netlist generator is a little quirky since it is rarely used**
- **Sometimes it generates spice netlists that are not compatible with the standard**
 - Needs to be cleaned
- **A few scripts are written in matlab/awk/perl for evaluation, cleaning, and translation**
 - Basis for our CPU/Memory baseline



Environment and Metrics

- **Simplified case for first year (2012 Tau contest will be more challenging)**
 - No GPU/parallel algorithm
 - All codes are submitted in static-linked binary format and need to run on Linux OS
- **Quality metrics:**
 - Accuracy, runtime, and memory
 - Use penalty as score. Smaller score is better.
- **Four categories.**
 - Maximum error (worst node voltage mismatch)
 - Average error (average node voltage mismatch)
 - Runtime (total user+system on a single processor)
 - Memory (peak amount of memory used)

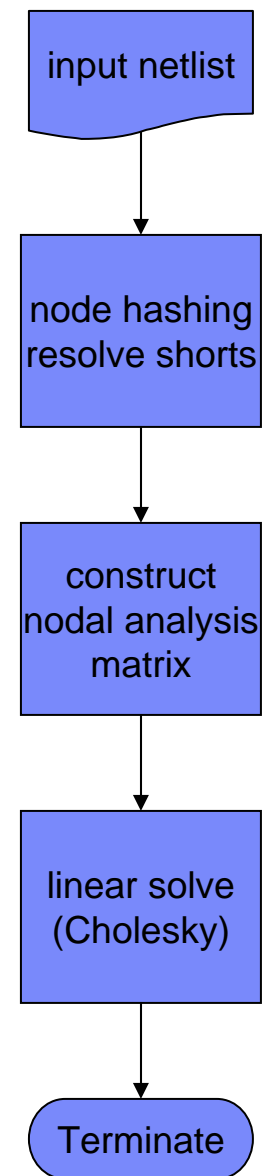
Evaluation Metrics

- **The penalty score for each category ranges from 0 to 100. The sum of scores across all categories and all benchmarks is the score for the team.**
- **Error (based on real world needs for this type of analysis):**
 - Maximum error: 100 if bigger than **10 mV**, 0 if less than **1 mV**
 - Average error: 100 if bigger than **1 mV**, 0 if less than **0.1 mV**
- **Runtime:**
 - Scale based on normalizing to the “internal solver”.
 - Score capped at 100 if runtime is longer than internal solver.
 - *Total Score set to 400 if runtime is more than 10x slower.*
- **Memory:**
 - Scale based on normalizing to the “internal solver”.
 - Score capped at 100 if memory is larger than internal solver.
 - *Total Score set to 400 if memory exceeds contest machine (64G).*

A Simple Linear Solver

- **A direct solver built with 200 lines of perl code and 200 lines of matlab code**
 - To ensure each benchmark is solvable and double check the solutions
 - Somewhat surprisingly, it is actually not too slow
- **Perl routine to hash the nodes into indices and resolve cliques of shorts**
- **Construct the pure (i.e. not modified) nodal analysis matrix**
 - Technique: remove the dependencies introduced by voltage sources from the matrix formulation. See reference below for more details
 - Benefit: matrix is symmetric positive definite → can use Cholesky solver
- **Use the sparse matrix solver in Matlab**
 - Automatically takes advantage of SPD condition
- **For team “Matlab”, total penalty is $11 * 200 = 2200$**
 - $200 = 0$ for max error + 0 for avg error + 100 cpu + 100 memory!

[†]F. Liu and P. Feldmann, “MAISE: an interconnect simulation engine for timing and noise analysis”, Proc. ISQED 2008



Results

All final 10 teams are winners! We could not have had this contest without you!

Special Note: the results show “team ID” for the results, and the top 3 teams have been notified of their status. Each team knows what his ID is. But the results beyond that point are not made public to encourage the teams to do even better in upcoming contests.

Results on Six Released Benchmarks

Team ID	CPU Score
1	80 (13.3% of matlab)
2	156
3	112
4	151
5	279
6	321
7	461
8	574
9	420
10	1355

Team ID	Memory Score
1	130 (21.6% of matlab)
2	182
3	307
4	356
5	390
6	386
7	600
8	160
9	558
10	500*

Team ID	Error Score
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	400*
9	687*
10	0 *

- **ibmpg3 to ibmpg6, newpg1 and newpg2. Matlab total CPU time: 565.9 second (600 score), total Memory: 10.8G (score: 600), total score: 1200**
- **Team 8 cannot solve two testcases.**
- **Team 9 has errors for 5 out of 6 testcases.**
- **Team 10 failed ibmnewpg1, and two testcases take more than 10x longer than matlab**

Results on All 11 Eleven Benchmarks

Team ID	CPU Score
1	136 (12.4% of matlab)
2	379
3	242
4	660*
5	1972*
6	2321*
7	1261*
8	1074
9	669
10	1723

Team ID	Memory Score
1	265 (24.1% of matlab)
2	330
3	634
4	654*
5	490*
6	386*
7	600*
8	306
9	1034
10	1000*


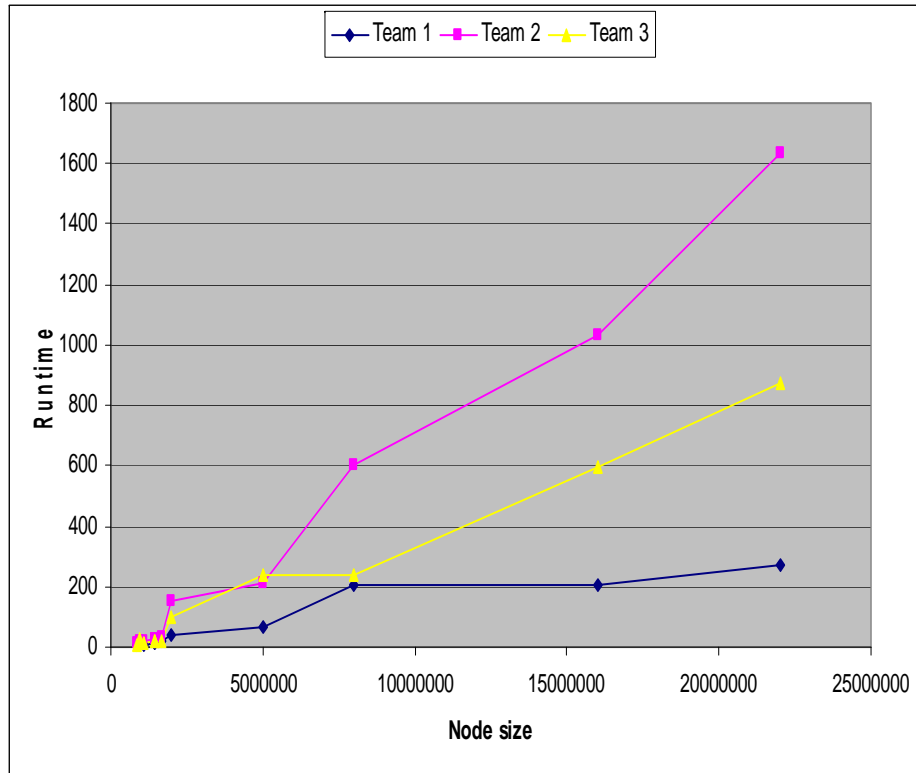
Team ID	Error Score
1	0
2	0
3	0
4	200*
5	0*
6	0*
7	0*
8	800
9	1687
10	1000 *

- Matlab total CPU time: 7986 second (1100 score), total Memory: 89.9G (score: 1100), total score: 2200
- 5 teams can run all testcases (Team 1/2/3/8/9)
- Team 4: Failed largest one, and large error on another one
- Team 5: Failed 4 of 5 testcases. No errors on good one
- Team 6: Failed on all of new testcases, can not handle only vdd cases (from error log)
- Team 7: All internals run 10x more than Matlab.
- Team 8: Large error on two internal benchmarks.
- Team 9: Large errors on all internal benchmarks
- Team 10: Large errors on all internal benchmarks

Special Notes

- **Team 2 and team 3 are very close. One is better for runtime, one is better for memory, but team 2 has better runtime for two benchmarks, and better memory for all benchmarks, so total score is better.**
- **Team 4: 3rd best runtime for six released benchmarks.**
- **Team 8: 2nd best memory for six released benchmarks and all released benchmarks. Big improvements, alpha version 10 hours for ibmpg1.**
- **Team 7: Big runtime improvement from alpha version.**
- **Failure on public benchmarks**
 - ibmpg3: team 8 and 10 (large errors/runtime)
 - ibmnewpg1: team 10
 - lbnnewpg2: team 8 and team 10
- **Most teams (other than team 9) have good accuracy.**
- **Slides with all finished data will put on website. Will notify each individual team for the score.**
- **The organizers will release the 5 internal benchmarks on the same benchmark web site (thank you TAMU for hosting).**

Top 3 Teams Runtime/Memory Chart



Third Place (Team 3)

Simulator: TicTac

Team Members: Zhiyu Zeng, Tong Xu and Peng Li

Affiliation: Texas A&M University

Award: \$150 + Plaque

Second Place (Team 2)

Simulator: SEVA

- **Team Members: Shih-Chieh Chang, Chung-Han Chou, Che-Rung Lee, Sheng-Yuan Lin, Yiyu Shi, Nien-Yu Tsai, Hao Yu**

Affiliation: Dept. of Computer Science, National Tsing Hua University; Dept. of Electrical and Computer Engineering, Missouri University of Science and Technology

Award: \$200 + Plaque

First Place (Team 1)

Simulator: PowerRush

- **Team Members: Jianlei Yang, Zuowei Li, Yici Cai, Qiang Zhou, Yuchun Ma**

Affiliation: Tsinghua University

Award: \$600 + Plaque

Next Year

- **Transient Analysis?**
- **Uncertainty Analysis?**
- **Allow Hardware Acceleration? GPU? Multicore?**

- **We will be making plans and informing the community.**
- **We will also announce the results earlier (so teams can plan on attending).**