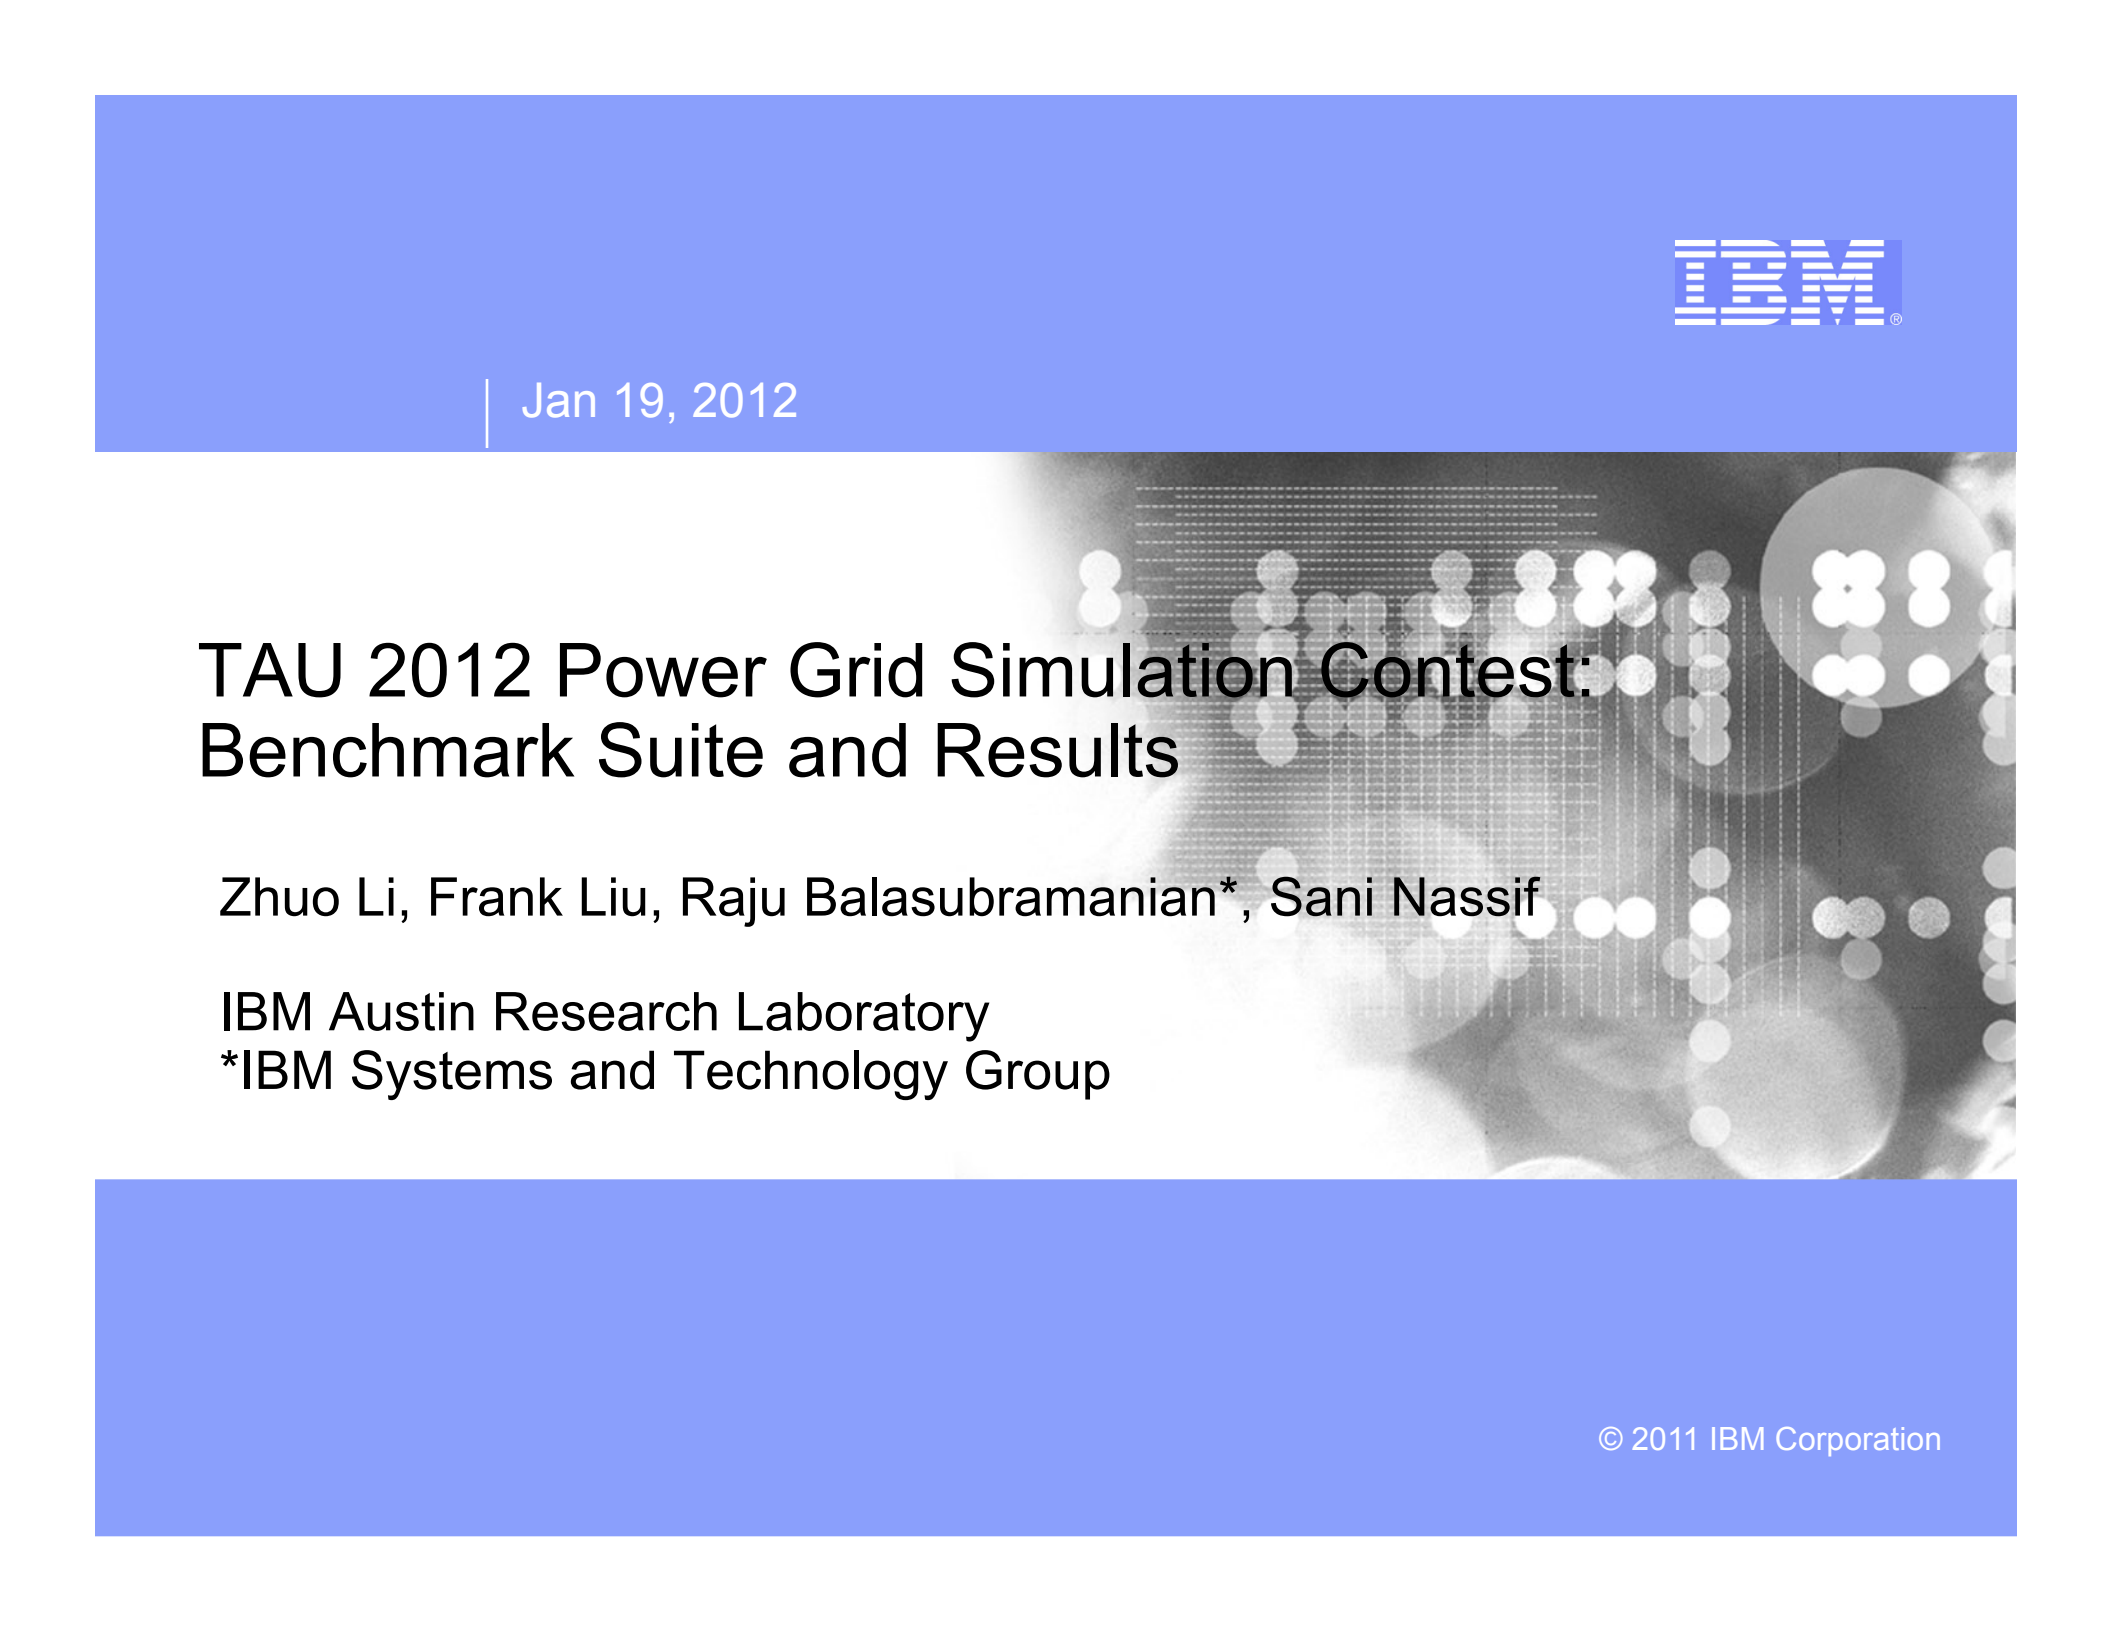

Jan 19, 2012

TAU 2012 Power Grid Simulation Contest: Benchmark Suite and Results

Zhuo Li, Frank Liu, Raju Balasubramanian*, Sani Nassif

IBM Austin Research Laboratory
*IBM Systems and Technology Group

Power Grid Contest History

- Motivate power grid simulation research
  - (one special session at ICCAD and one power grid paper in ICCAD)
- 1st Power Grid Contest 2011: DC Solution of Large Power Grids
  - 12 registered and 10 final teams, top 3:
    - PowerRush (TsingHua Univ.)
    - SEVA (National Tsing Hua Univ.; Missouri Univ. of Science & Technology)
    - TicTac (Texas A&M University)
- 2nd contest 2012: Transient Solution with Multi-Threading.
  - Build upon work done for the DC Solver.
- Special thanks to:
  - Prof. Peng Li and Albert Zeng for hosting the benchmark website.
  - Nancy Zhou for visualization help.

Timeline

- July 13, 2011. Contest announcement.
- August 30, 2011. First benchmark released.
- September 9, 2011. Eleven teams registered!
- October. 9, 2011. Eight alpha submissions.
- November 23, 2011. Second benchmark is released.
- December 4, 2011. Six final entries (5 U.S., 1 international).
  - It is a harder problem, only 5 “working” entries.
- December 30, 2011. Top 3 teams informed.
- January 19, 2012. Final results from 6 benchmarks.
  - Quality of result metrics: Accuracy, runtime and memory.

Teams

Simulator Name	Affiliation	Submitted an Entry
IITPGS	Illinois Institute of Technology	✓
PumaGazelle	Texas A&M University	✓
ETBR	University of California, Riverside	✓
pgt_solver	University of Illinois at Urban-Champaign	✓
pixel_pgsim	Michigan Tech University	✓
PowerRush	Tsinghua University	✓
Tambaguchi	Kyoto University	
Cheiron	University of Thessaly, Volos, Greece	

Benchmarks

- Six DC benchmarks were released at 2008, ASPDAC.
- Seven more industrial benchmarks ranging from 2M to 22M nodes were used in 2011 TAU Power Grid Contest (two have been released, five are pending IBM's release process).

Name	# Nodes	# Current Source	# Resistors
IBMPGNEW1	1461039	357930	2352355
IBMPGNEW2	1461039	357930	1422830
IBMPGNEW3	2256393	4516838	5425827
IBMPGNEW4	5430929	4516838	8603045
IBMPGNEW5	8969301	8411004	22091835
IBMPGNEW6	16633651	7248078	24869314
IBMPGNEW7	22214300	8411004	35355796

From DC to Transient Benchmarks

- New benchmarks were needed for this contest.
- To expedite, we enhanced the existing six ASPDAC 2008 DC benchmarks (ibmpg<n> ➔ ibmpg<n>t).
  - Using information from the original internal benchmarks, but simplified to preserve anonymity.
- We released the first two benchmarks to all teams initially.
  - We further released the third benchmark to help with some parsing problems.
- We used the five largest benchmarks to score the contest.

Basic Steps in Power Grid Analysis

- Input processing
- Electrical extraction (geometry to RLC elements)
- Load current characterization
- Load and source stitching
- Circuit simulation
- Reporting

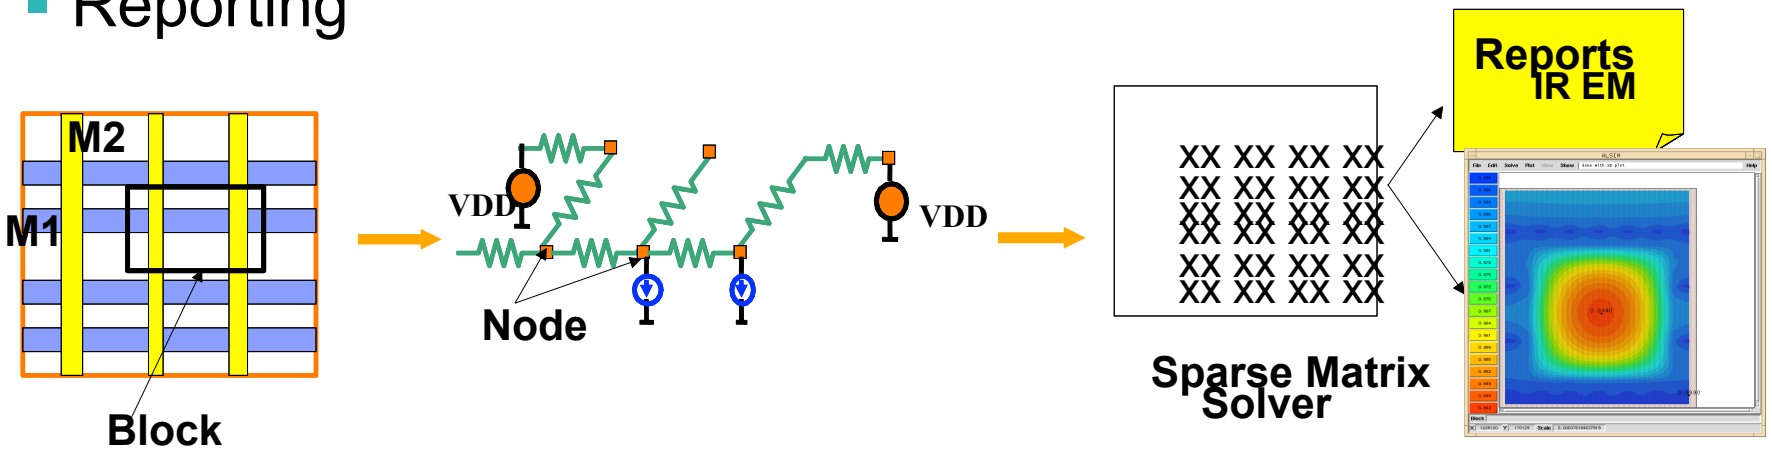

Courtesy Howard Smith, IBM

Transient Analysis Example Output

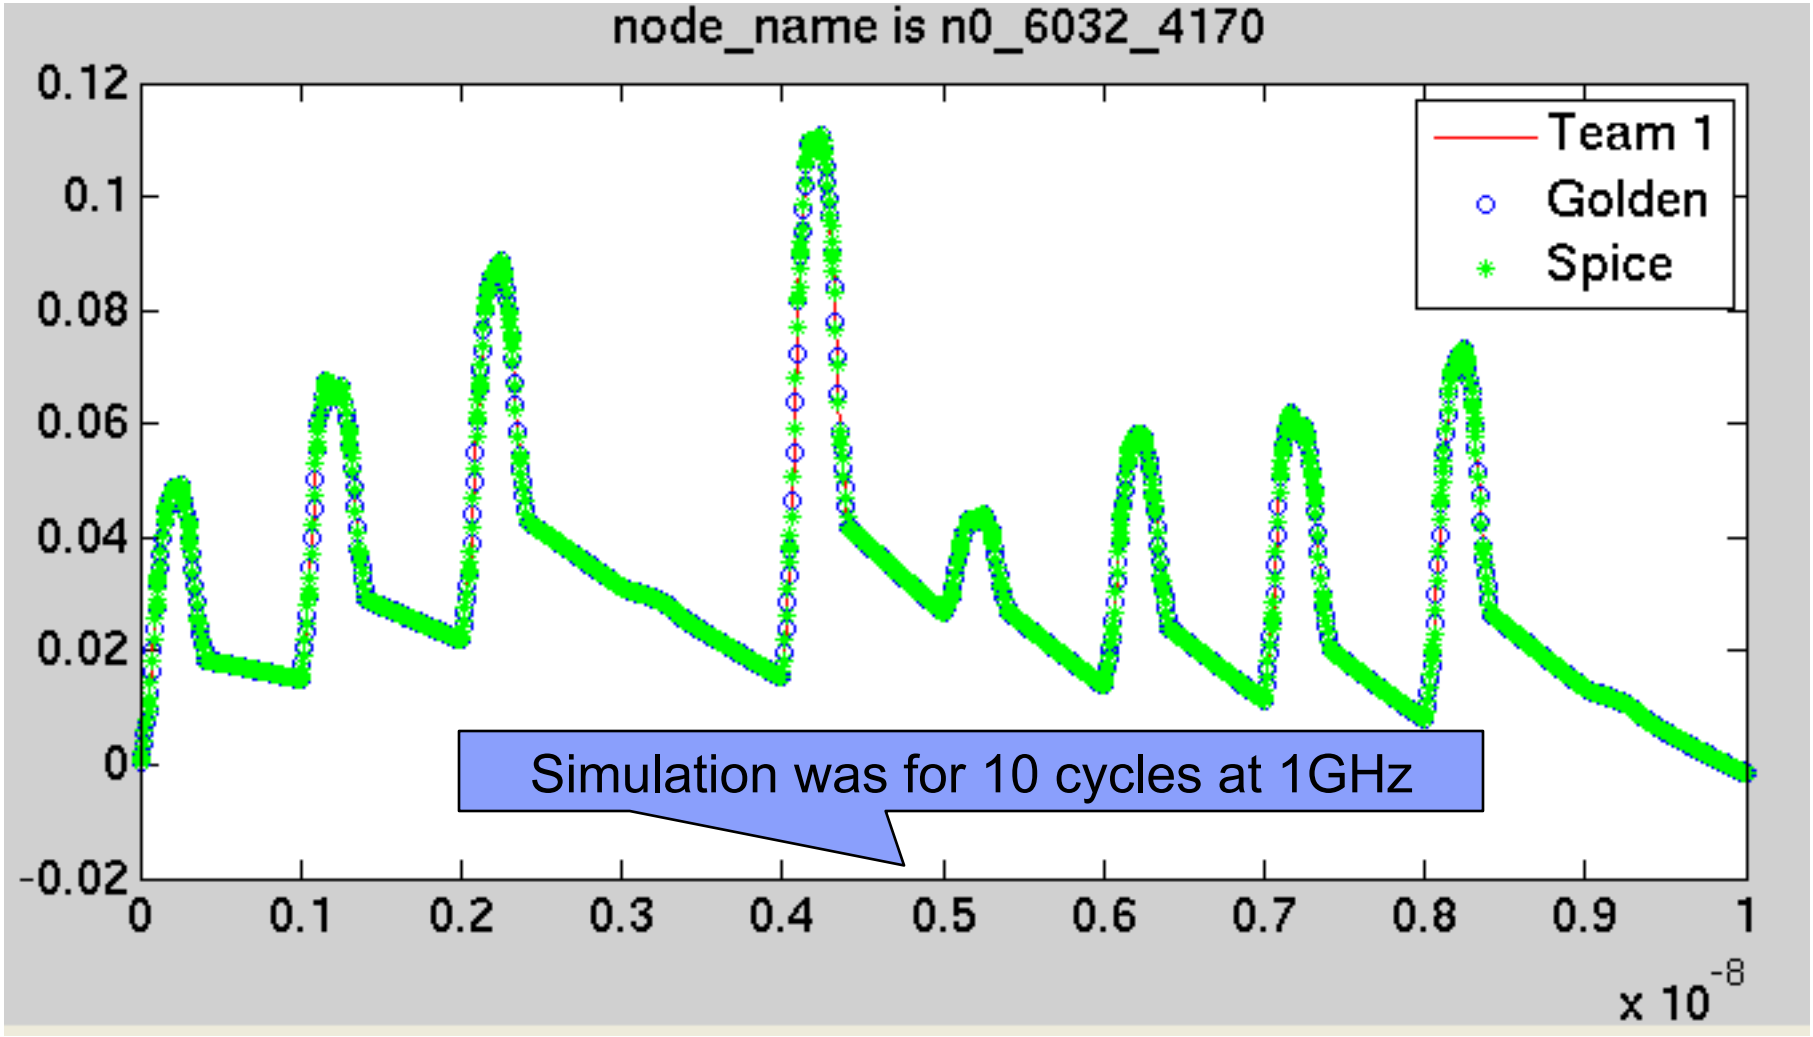

Environment and Metrics

- Allow parallel (multi-threaded) implementation.
  - The contest machine had 64 Intel CPUs, was running Linux, and all codes were submitted as statically-linked binaries.
- Quality metrics: Accuracy, runtime, and memory.
  - Score managed as a penalty, so smaller scores are better.
- Four categories:
  - Maximum error (worst node voltage mismatch).
  - Average error (average node voltage mismatch).
  - Runtime (wall clock time with constant machine).
  - Memory (peak).

Metric Details

- The penalty score for each category ranges from 0 to 100.
- The sum of the scores across all categories and all benchmarks is the final score for the team.

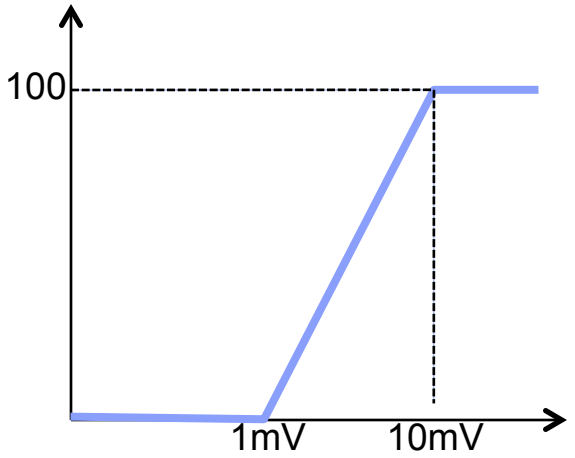

- Example: Error Metric
  - Case 1: Maximum Voltage Error:
    - 100 if bigger than 10 mV, 0 if less than 1 mV.
  - Case 2: Average Voltage error:
    - 100 if bigger than 1 mV, 0 if less than 0.1 mV.
  - Based on the real world needs for this type of analysis…

Other Evaluation Metrics

- Runtime (normalized to IBM's internal tool):
  - Score capped at 100 if runtime is longer than a reasonable scaling of internal solver.
    - No 400 penalty score this year since no contestants ran too long.
- Memory (normalized to IBM's internal tool):
  - Score capped at 100 if memory is larger than a reasonable scaling of internal solver.
    - Total Score set to 400 if memory exceeds contest machine (700G).
      - Did not happen to any of the contestants.
- IBM's "internal tool" uses a parallel solver which has been in development within IBM for many (many) years.

Results

Thanks again to all the teams for their strong participation and continuous supports!

Special Note: The names of top three teams are released, others are not made public to encourage the teams to do even better in upcoming contests.

Raw Score Results

Team ID	CPU Score	Memory Score	Error Score
1 (serial)	114	190	0
1 (parallel)	67	320	0
2 (serial)	229	100	49
2 (parallel)	169	160	49
3 (serial)	239	288	49
3 (parallel)	150	415	49
4 (serial)	240	235	250
4 (parallel)	432	281	250
5 (serial)	354	352	49
5 (parallel)	304	500	682

Some Observations on Parallelism

Perhaps we need more courses in parallel programming?

- For most of the contest entries:
  - Single threaded mode achieved the best score.
  - The ratio of wall-clock-time(single) to wall-clock-time(parallel) is less than 2X.
- CPU-time/wall-clock-time (effective number of processors) is < 32.
  - For the top 3 teams, the best ratio was 4.3!
  - One team did achieve a ratio of 26, but the wall-clock-time is not as good as the top 3 teams.
- The parallel implementation of transient power grid analysis is still an open problem with room for improvement!

Special Notes

- Team 1 achieves the fastest runtime in both single and parallel modes.
  - If I want to get one or multiple jobs done asap, as long as I have enough cpus and memory.
- Team 2 achieves the best memory in both single and parallel modes.
  - If I have limited memory resources and want to run multiple jobs.
  - It also has almost perfect memory/runtime tradeoff and get the same score.
- Team 4 ranks 3rd in memory category for both single and parallel modes.
- Team 3 ranks 2nd in runtime category in parallel mode.

Top 3 Teams Runtime Chart

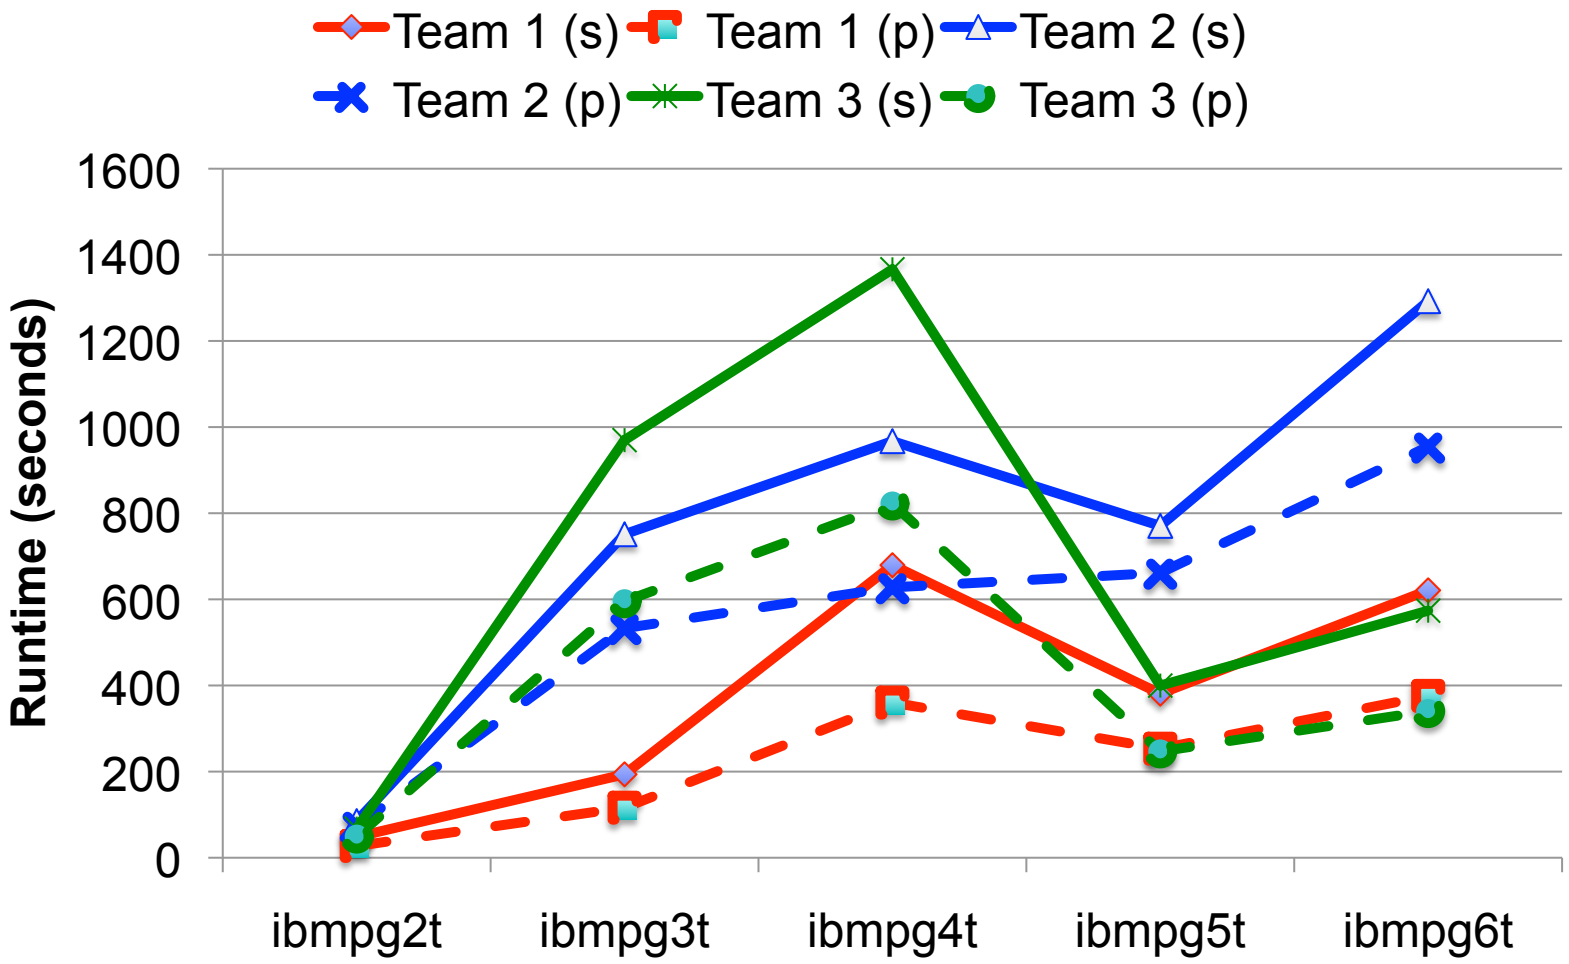

Top 3 Teams Memory Chart

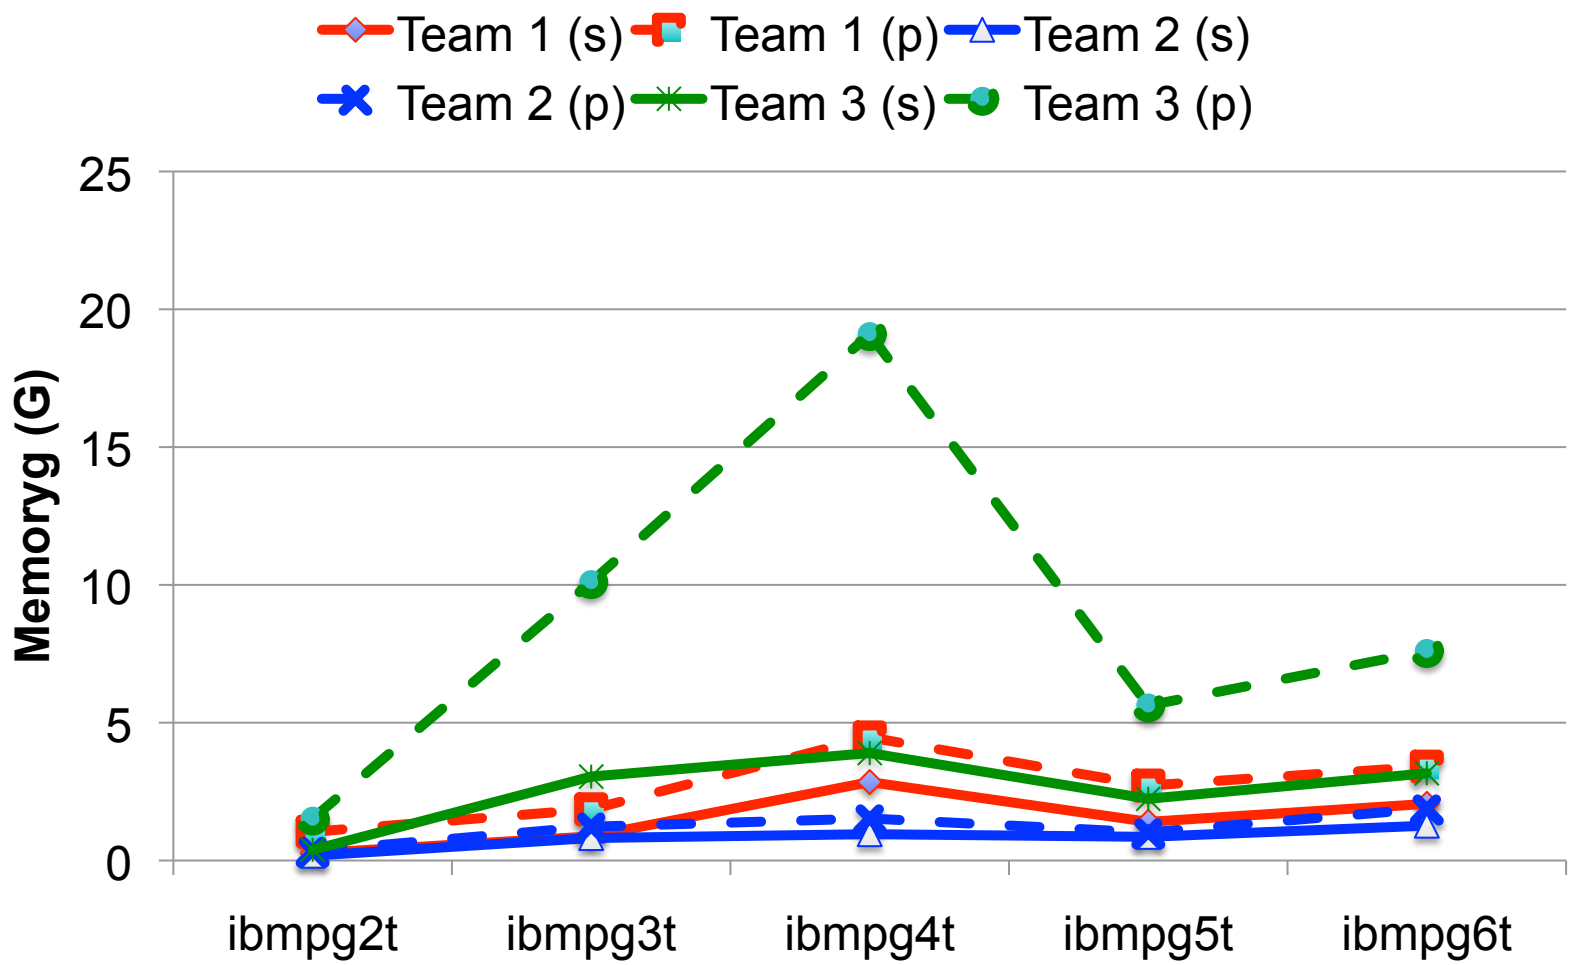

Team 1 Runtime Profile for single/parallel mode

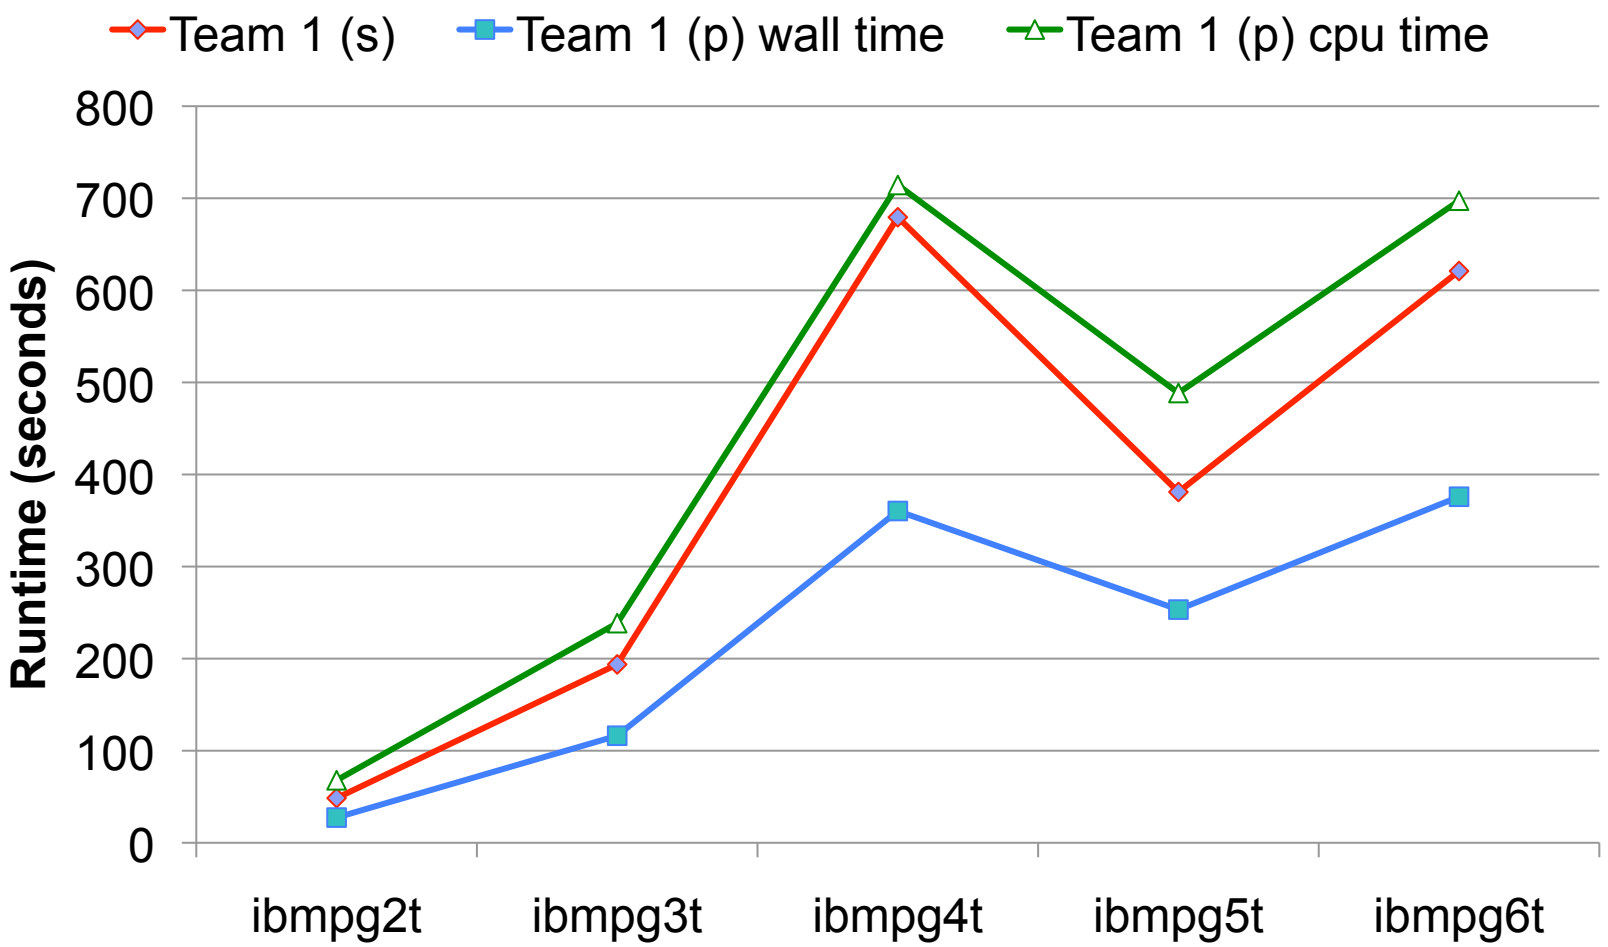

Third Place (Team 3)

Simulator: IITPGS

Team Members: Xuanxing Xiong, Jia Wang

Affiliation: Illinois Institute of Technology

Award: $300 + Plaque

Second Place (Team 2)

Simulator: PowerRush

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

Affiliation: Tsinghua University

Award: $600 + Plaque

First Place (Team 1)

Simulator: pgt_solver

- Team Members: Ting Yu, Martin D. F. Wong

Affiliation: University of Illinois at Urban-Champaign

Award: $1000 + Plaque

Next Year

- Steady State Transient Analysis?
- Uncertainty Analysis?
- Allow Hardware Acceleration? GPU?

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