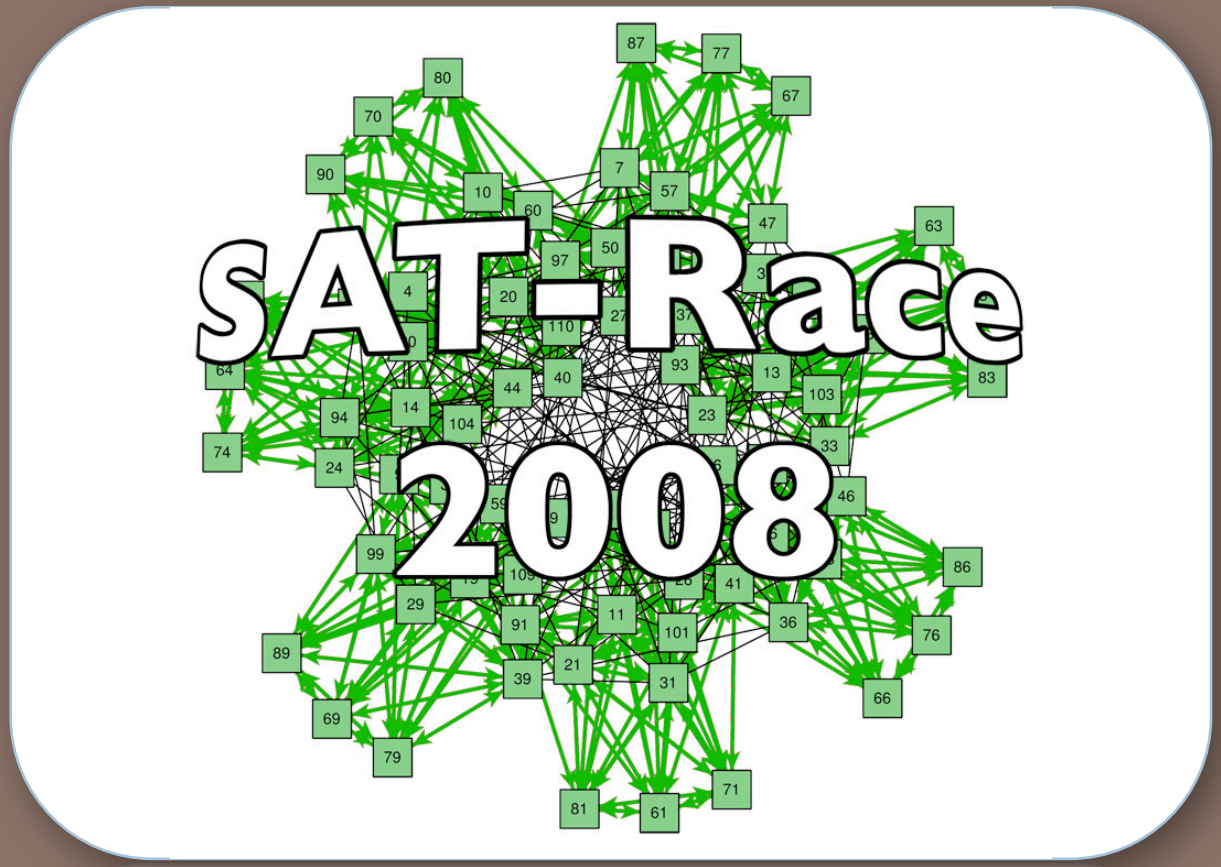

Carsten Sinz • Nina Amla • Toni Jussila
Daniel Le Berre • Pete Manolios • Lintao Zhang
Himanshu Jain • Hendrik Post

What is SAT-Race?

- „Small SAT-Competition“
  - Only industrial category benchmarks
    (no handcrafted or random)
  - Short run-times
    (15 minutes timeout per instance)
  - Mixture of satisfiable / unsatisfiable instances
    (thus not suitable for local-search solvers)
  - „Black-box“ solvers permitted
- New this year:
  - Special tracks for multi-threaded (parallel) solvers and AIG solvers

Organizers

- **Chair**
  - Carsten Sinz (Universität Karlsruhe Germany)
- **Advisory Panel**
  - Nina Amla (Cadence Design Systems, USA)
  - Toni Jussila (OneSpin Solutions, Germany)
  - Daniel Le Berre (Université d'Artois, France)
  - Panagiotis Manolios (Northeastern University, USA)
  - Lintao Zhang (Microsoft Research, USA)
- **AIG Special Track Co-Organizer**
  - Himanshu Jain (Carnegie Mellon University, USA)
- **Technical Organization**
  - Hendrik Post (Universität Karlsruhe, Germany)

Solvers

- Received 43 solvers by 36 submitters from 16 nations
  (SAT-Race 2006: 29 solvers by 23 submitters)

Australia	2
Austria	2
Canada	7
Finland	1
France	4
Germany	4
India	1
Israel	2

Netherlands	1
N. Ireland	1
P.R. China	1
Russia	1
Spain	1
Sweden	4
USA	8
UK	2

- 9 industrial solvers, 34 academic
- 27 solvers in Main Track, 8 in Parallel Track, 8 in AIG Track

Qualification

- Two qualification rounds
  - Each consisting of 50 benchmark instances
  - Increased runtime-threshold of 20 minutes
  - Successful participation in at least one round required to participate in SAT-Race
- To ascertain solver correctness and efficiency
- 1st round took place after February 10, 2nd round after March 12

Results Qualification Rounds

- Qualification Round 1:
  - 15 solvers already qualified for SAT-Race (by solving more than 40 out of 50 instances)
  - 9 in Main Track, 1 in Parallel Track, 5 in AIG Track
- Qualification Round 2:
  - 11 solvers additionally qualified (by solving more than 20 out of 50 instances)
  - 8 in Main Track, 2 in Parallel Track, 1 in AIG Track
- Overall result: 25 (out of 43) solvers qualified
  - 17 in Main Track, 3 in Parallel Track, 6 in AIG Track
  - One solver withdrawn

Qualified Solvers: Main Track

Solver	Author(s)	Affiliation
Barcelogic	Robert Nieuwenhuis et al.	Tech. Univ. Catalonia
Clasp	Torsten Schaub et al.	University of Potsdam
CMUSAT	Himanshu Jain	CMU
eSAT	Said Jabbour et al.	CRIL Lens / Microsoft
Eureka	Vadim Ryvchin, Alexander Nadel	Intel
kw	Johan Alfredsson	Oepir Consulting
LocalMinisat	Vadim Ryvchin, Ofer Strichman	Technion
MiniSat	Niklas Sörensson, Niklas Een	Independent / Cadence
MXC	David Bregman, David Mitchell	SFU
picosat	Armin Biere	Johannes Kepler University Linz
preSAT	Cédric Piette et al.	CRIL-CNRS / Microsoft
Rsat	Knot Pipatsrisawat, Adnan Darwiche	UCLA
SAT4J2.0	Daniel Le Berre	CRIL-CNRS
SATzilla	Lin Xu et al.	UBC
Spear	Domagoj Babic	UBC
Tinisat	Jinbo Huang	NICTA

Qualified Solvers: Special Tracks

Parallel Solvers:

Solver	Author(s)	Affiliation
ManySat	Youssef Hamadi	Microsoft Research
MiraXT	Tobias Schubert et al.	University of Freiburg
pMiniSat	Geoffrey Chu	University of Melbourne

AIG Solvers:

Solver	Author(s)	Affiliation
CMUSAT-AIG	Himanshu Jain	CMU
kw_aiger	Johan Alfredsson	Oepir Consulting
MiniCirc	Niklas Eén, Niklas Sörensson	Cadence Research / Independent
MiniSat++	Niklas Sörensson, Niklas Eén	Independent / Cadence Research
NFLSAT	Himanshu Jain	CMU
Picoaigersat	Armin Biere	Johannes Kelper University Linz

Benchmark Instances: Main / Parallel Track

- 20 instances from bounded model checking
  - IBM's benchmark 2002 and 2004 suites
- 20 instances from pipelined machine verification
  - 10 instances from Velev's benchmark suite
  - 10 instances from Manolios' benchmark suite
- 10 instances from cryptanalysis
  - Collision-finding attacks on reduced-round MD5 and SHA0 (Mironov & Zhang)
- 10 instances from software verification
  - C bounded model checking
- 40 instances from former SAT-Competitions (industrial category)
- Up to 11,483,525 variables, 32,697,150 clauses
- Smallest instance: 286 variables, 1742 clauses

Sizes of CNF Benchmark Instances

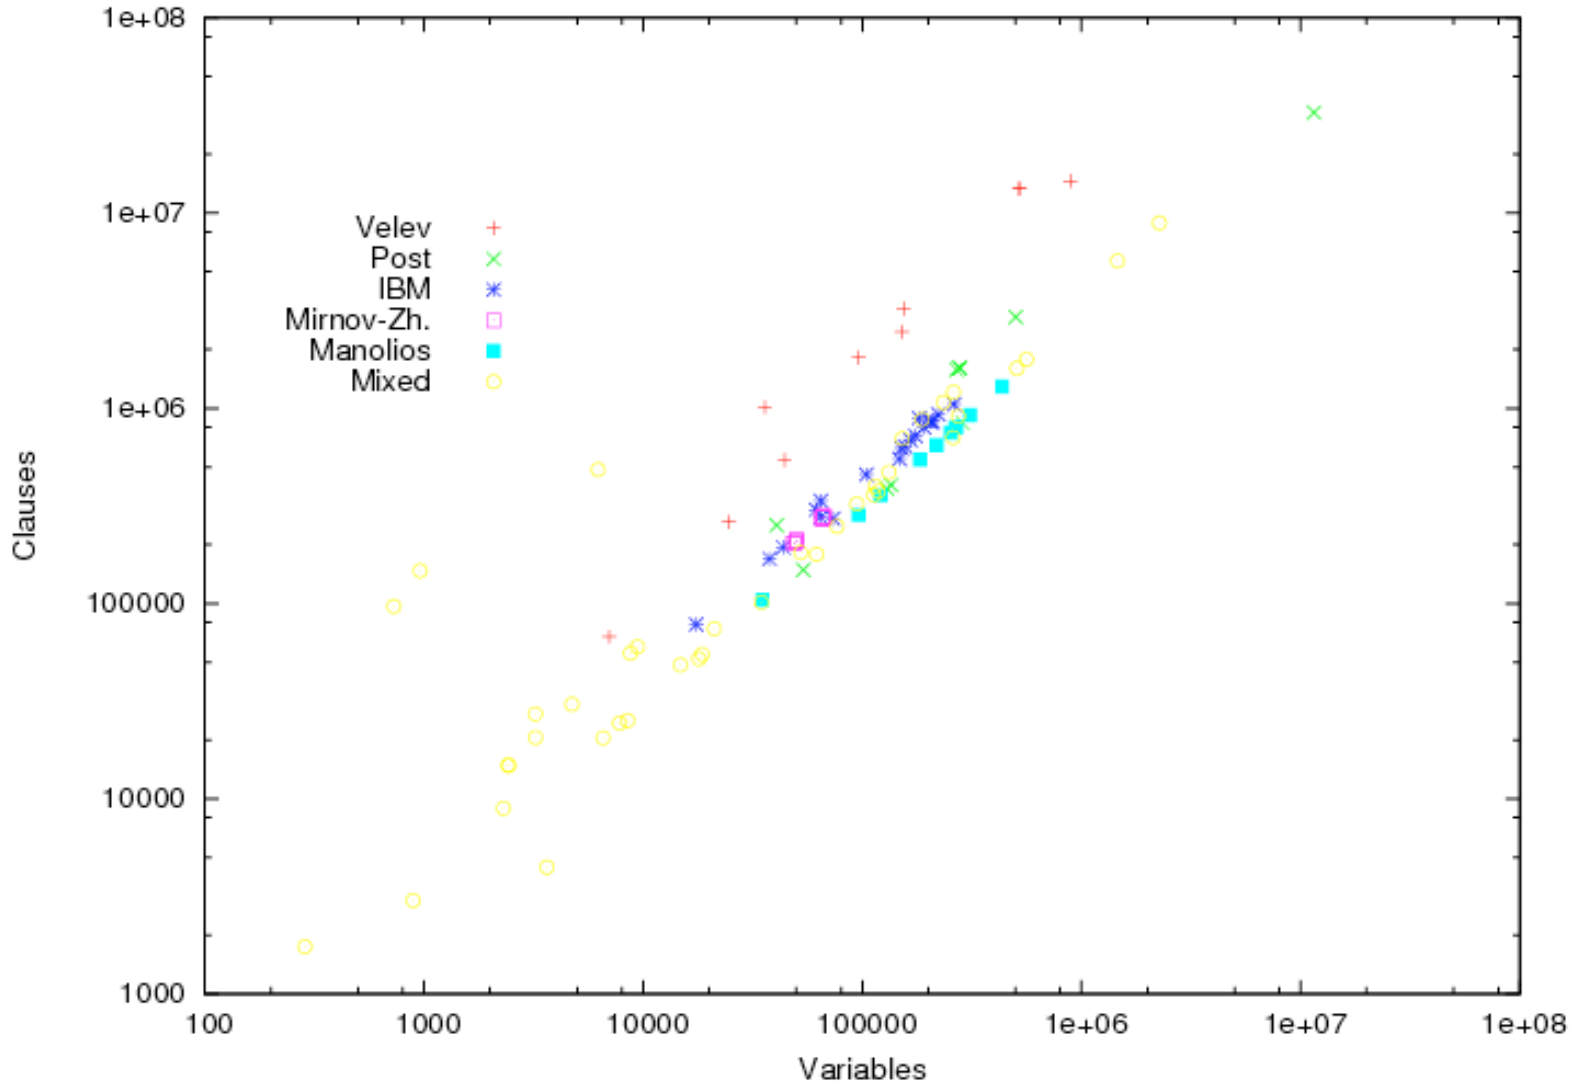

Benchmark Instances: AIG Track

- 9 Groups of Benchmark Sets:
  - Anbulagan / Babic / C32SAT / Mironov-Zhang / IBM / Intel / Manolios / Palacios / Mixed
- Instances mainly from last year's AIG Competition
- Additional instances provided by Himanshu Jain and Armin Biere

Parallel Track: Special Rules

- Run-times for multi-threaded solvers have high deviations (especially for satisfiable instances)
  - 3 runs of each solver on each instance
  - Median run-time is taken as result
  - Instance is considered as solved, if it could be solved in at least 1 out of 3 runs.

Scoring (for sequential tracks)

1. Solution points: 1 point for each instance solved in ≤900 seconds
2. Speed points:

   p_{max} = x / #successful_solvers

   $p_s = p_{max} \cdot (1 - t_s / T)$

   with x set to the maximal value s.t. $p_s \leq 1$ for all solvers and instances

Computing Environment

- Linux-Cluster at University of Tübingen
  - 16 compute nodes
  - 2 Intel Xeon 5150 Processors (Dual-Core, 2.66 GHz) per node
  - 8 MB main memory per node
  - Both 32-bit and 64-bit binaries supported
- Sequential/AIG Track: only one core per solver
- Parallel Track: 4 cores per solver

Results

Main Track (CNF Sequential)

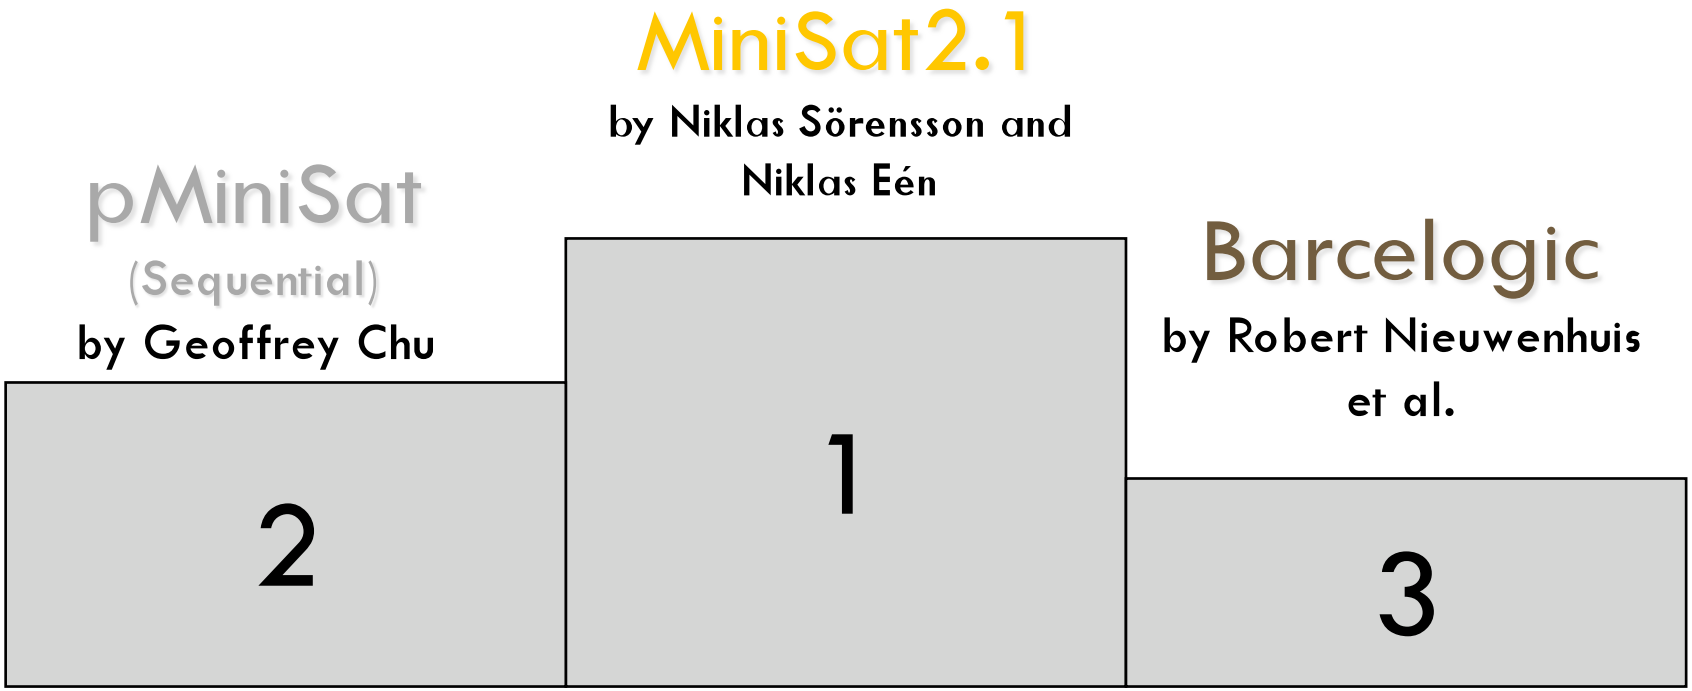

84.61 points,
79 solved instances

88.26 points,
81 solved instances

82.10 points,
77 solved instances

next best solver 81.04 points

Runtime Comparison: Main Track

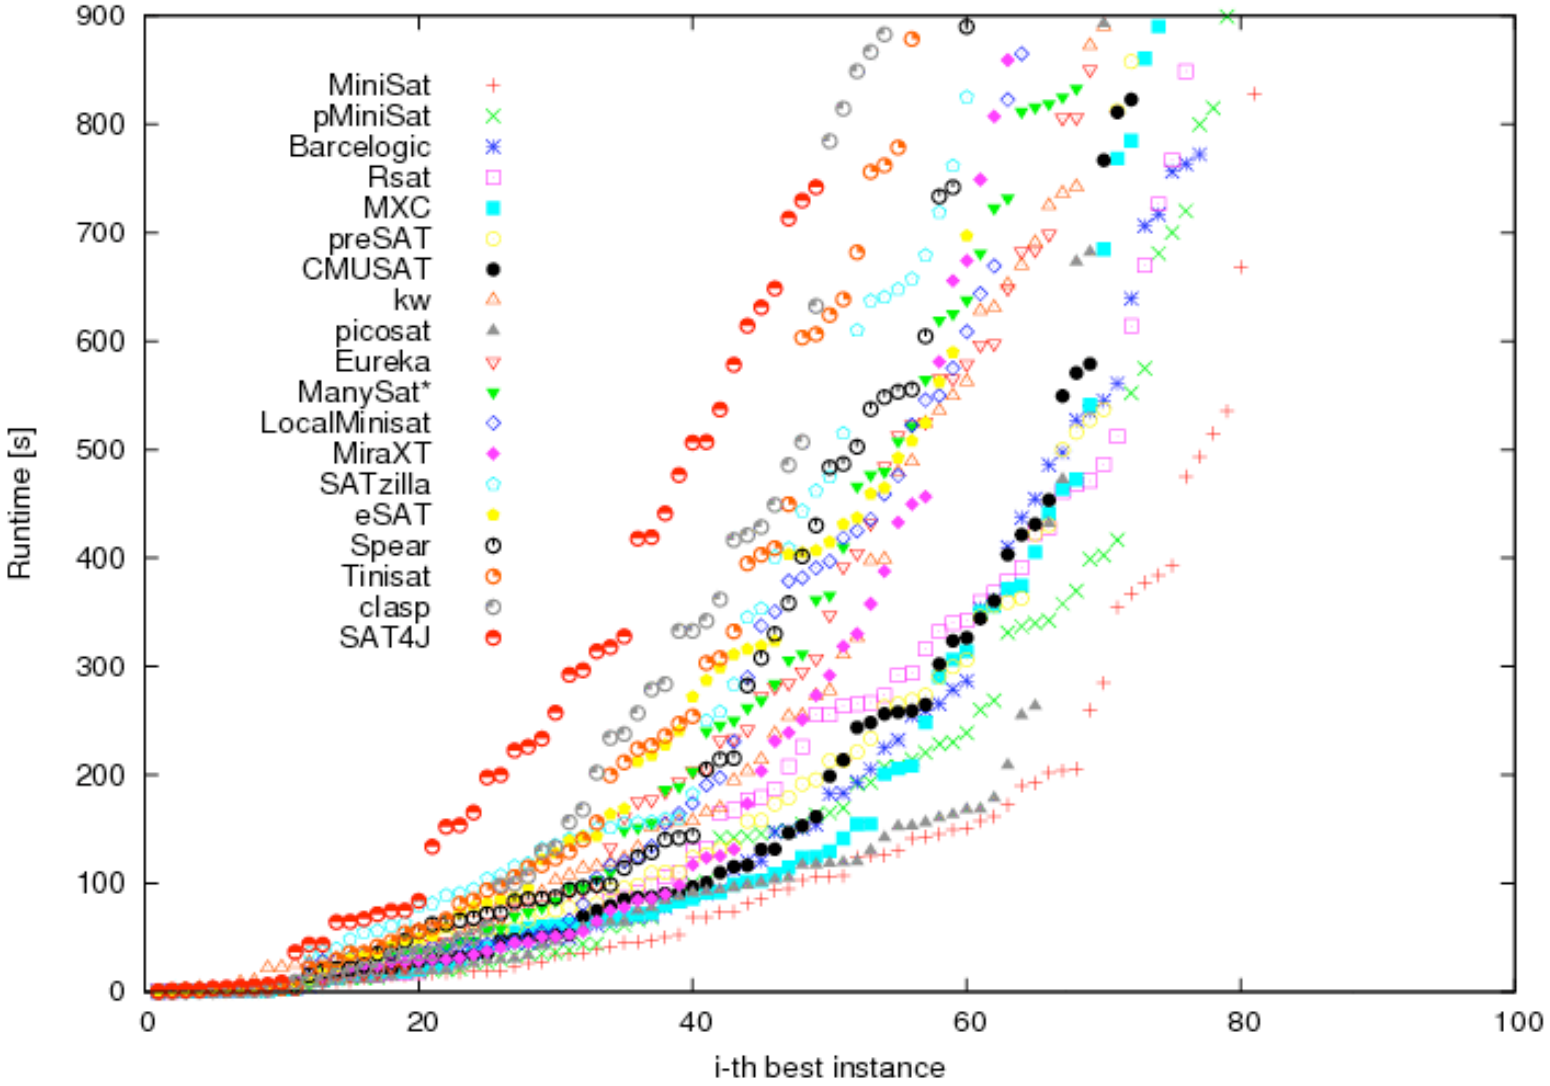

Special Track 1 (CNF Parallel)

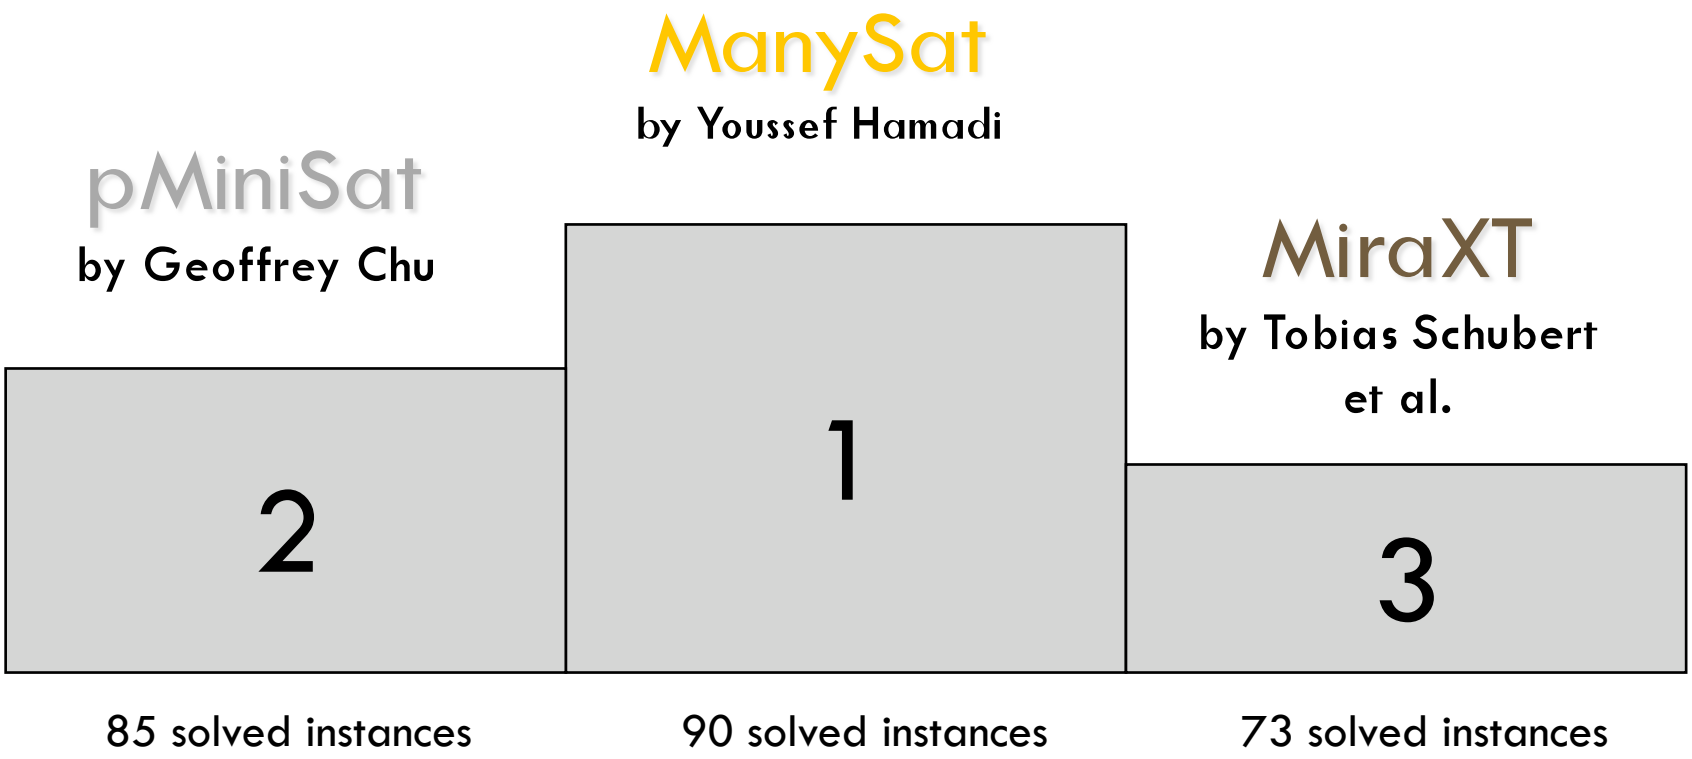

Runtime Comparison: Parallel Track

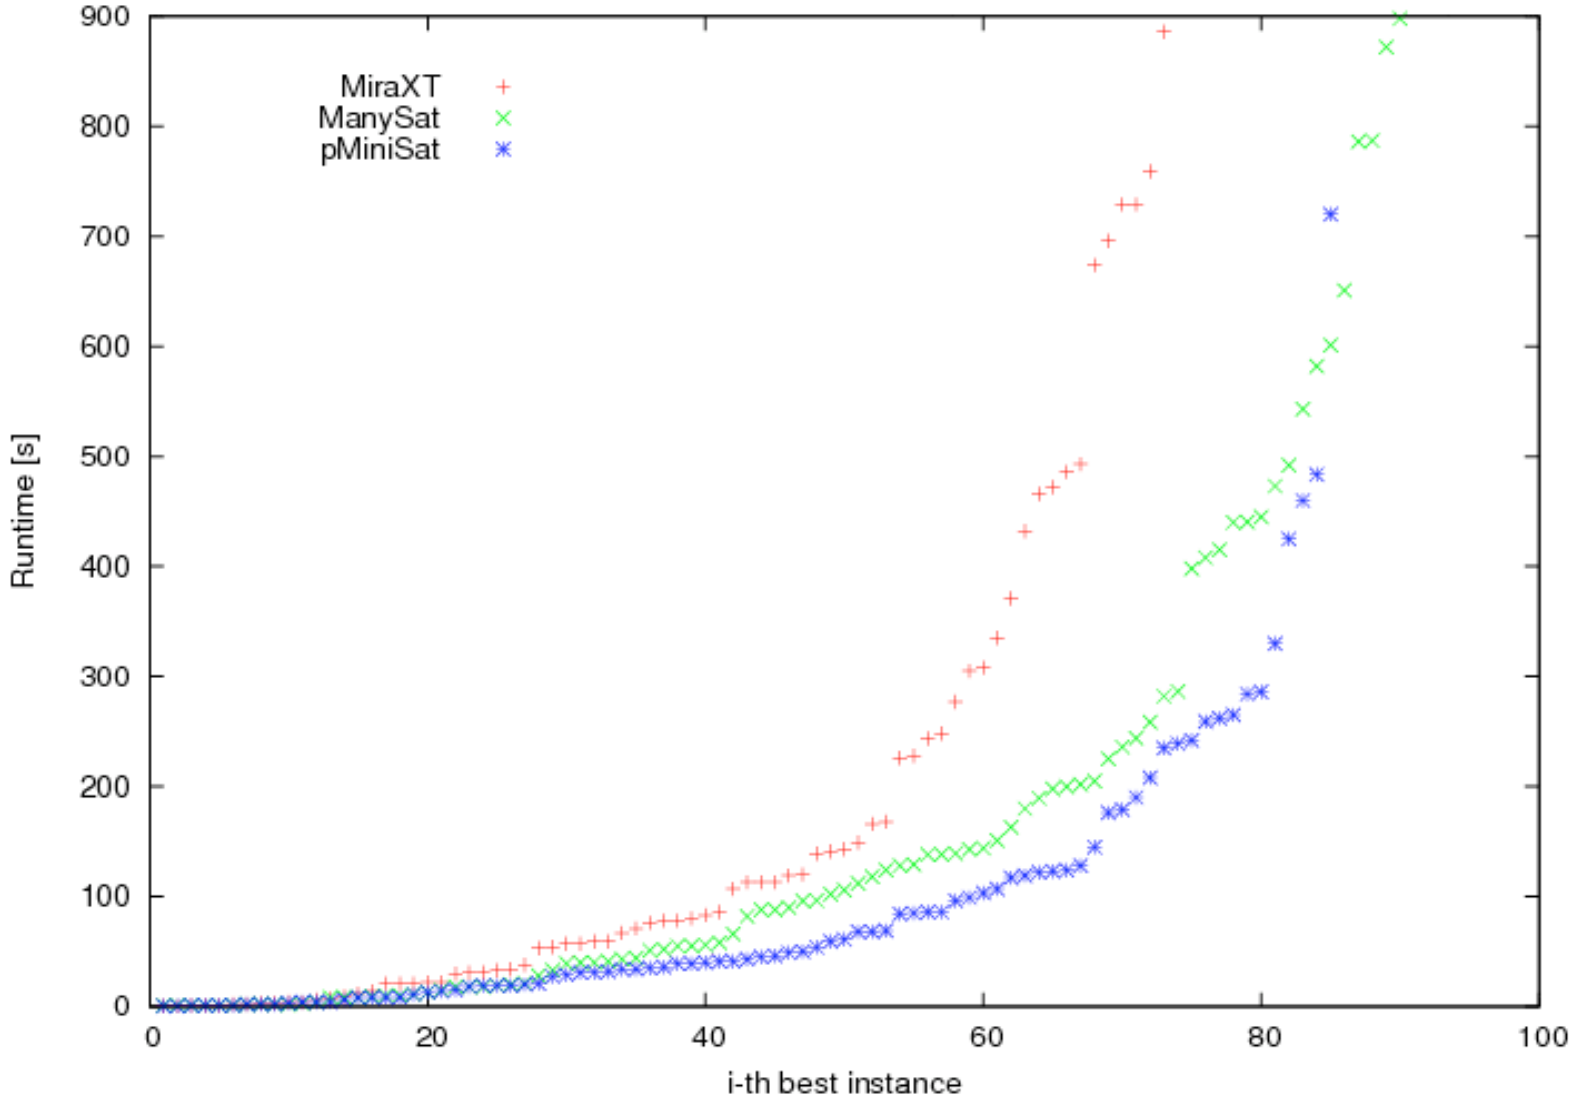

Special Track 2 (AIG Sequential)

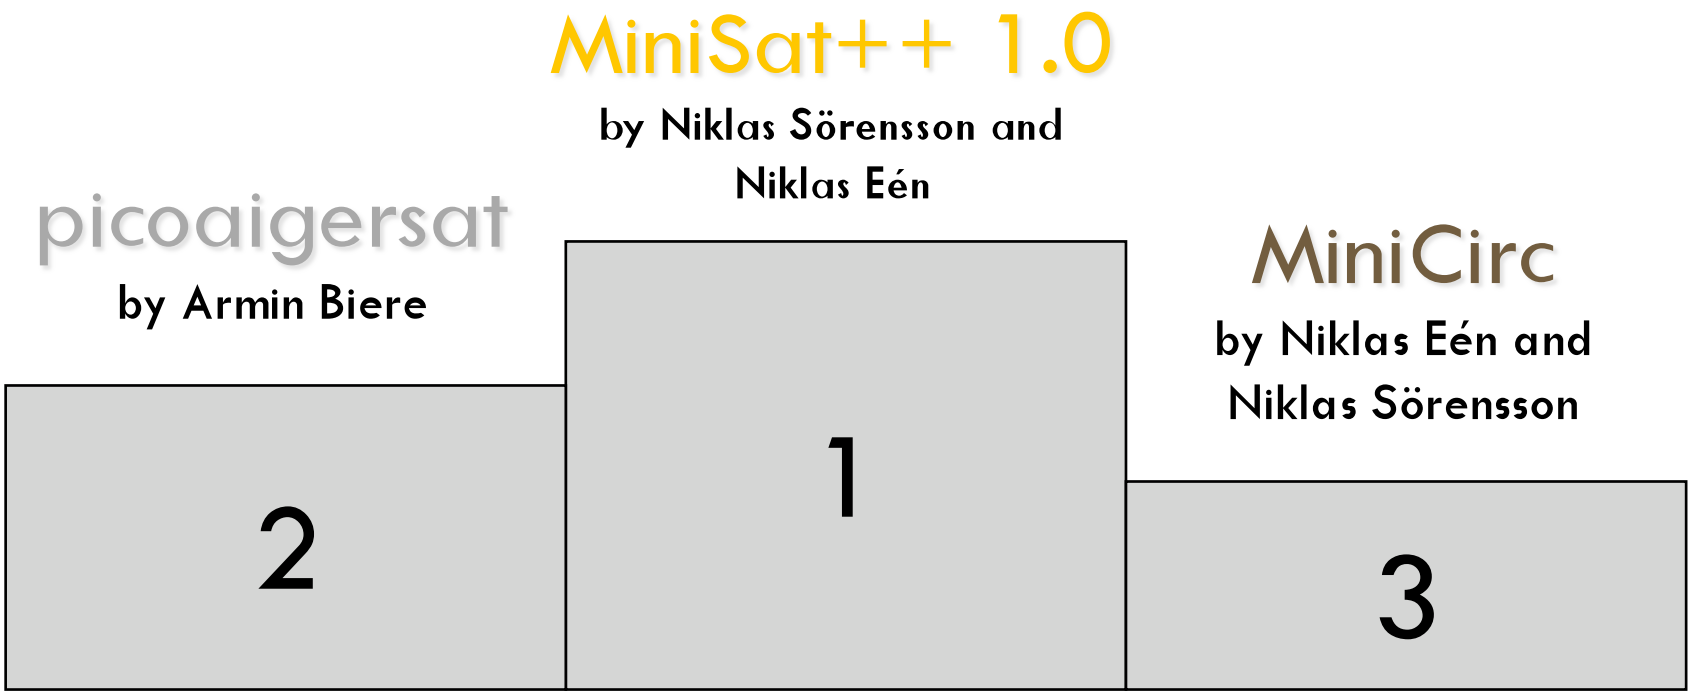

82.80 points,
69 solved instances

86.98 points,
74 solved instances

82.29 points,
70 solved instances

next best solver 81.85 points

Runtime Comparison: AIG Track

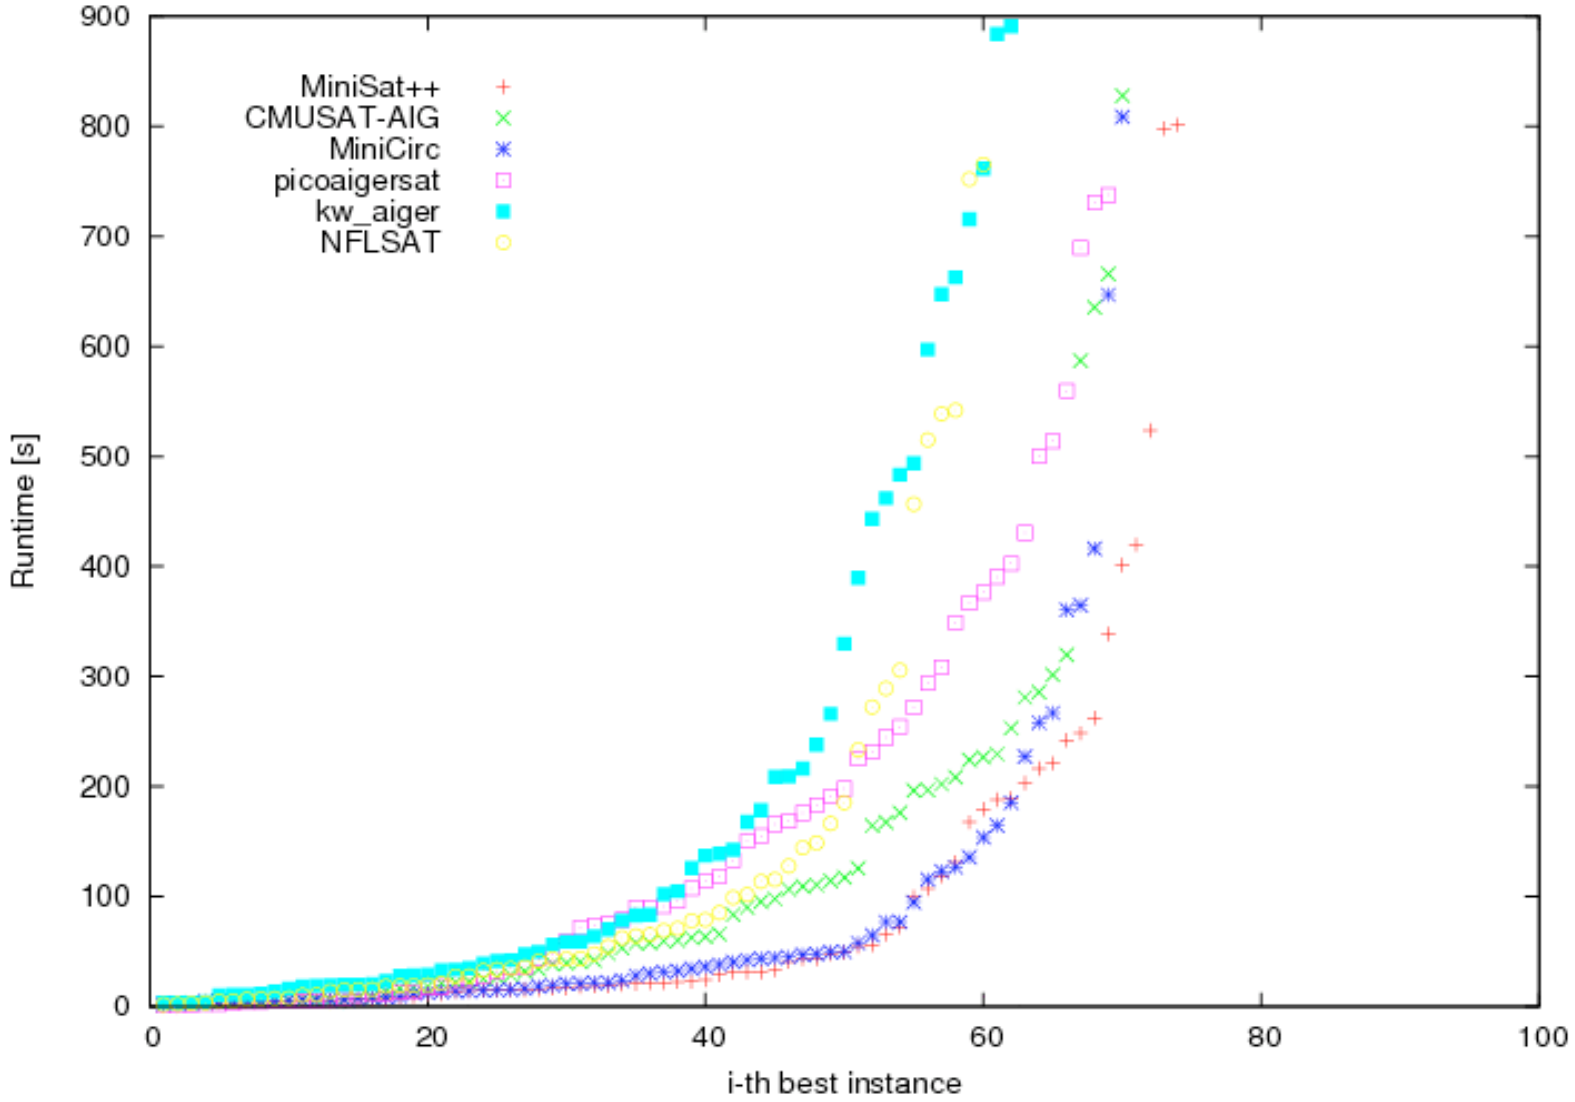

Lessons Learned

- Parallel solvers have not yet reached the quality of sequential solvers
  - 2 out of 5 solvers had to be rejected due to erroneous results
- Assessment of parallel solvers harder due to high runtime deviation
- 32-bit vs. 64-bit:
  - no clear advantage for either architecture
  - 32-bit: MiniSat; 64-bit: pMiniSat, Barcelogic
- Preprocessors are vital for large industrial instances

Conclusion

- Any Progress compared to SAT-Competition 2007?
  - SAT-Race 2008 winner can solve 6 more instances than SAT-Competition 2007 winner (SAT+UNSAT Industrial Category)
  - Four solvers out-perform SAT-Competition 2007 winner
  - Third best solver of SAT-Competition 2007 would have reached place 17 only
- New ideas for implementation, optimization
  - See solver descriptions on Poster Session this afternoon
- Many new solvers
  - but mostly slight variants of existing solvers