Benchmarking Polyhedral Operators

Duong NGUYEN QUE

duong@cri.ensmp.fr

Centre de recherche en informatique

École des mines de Paris

Numerical & Symbolic Abstract Domains 2005

21 January 2005

PARIS

Context

- PIPS (www.cri.ensmp.fr/pips)
- Dual polyhedral representations:
  - System of constraints
  - System of generators
- Polyhedral operators:
  - integer “satisfiability”
  - dual conversion
  - convex hull
  - etc ...
- Problems with time, memory space and coefficient magnitude

- Example: Satisfiability
  - Hardware: PC 2.4Ghz, RAM: 2GB
  - Program ocean.f (4373 LOC), PERFECT Club
  - Total calls of satisfiability test: 11916
  - Number of overflows: 424
  - Number of calls with execution time greater than 3 seconds: 3521
  - Largest system : 906 530 constraints

Benchmarking: Why?

- Existing implementations
  - Reuse
  - New heuristics
- Previous experimental evaluation
  - theory
  - < 200 tests
- Program analysis and transformation with real-life examples

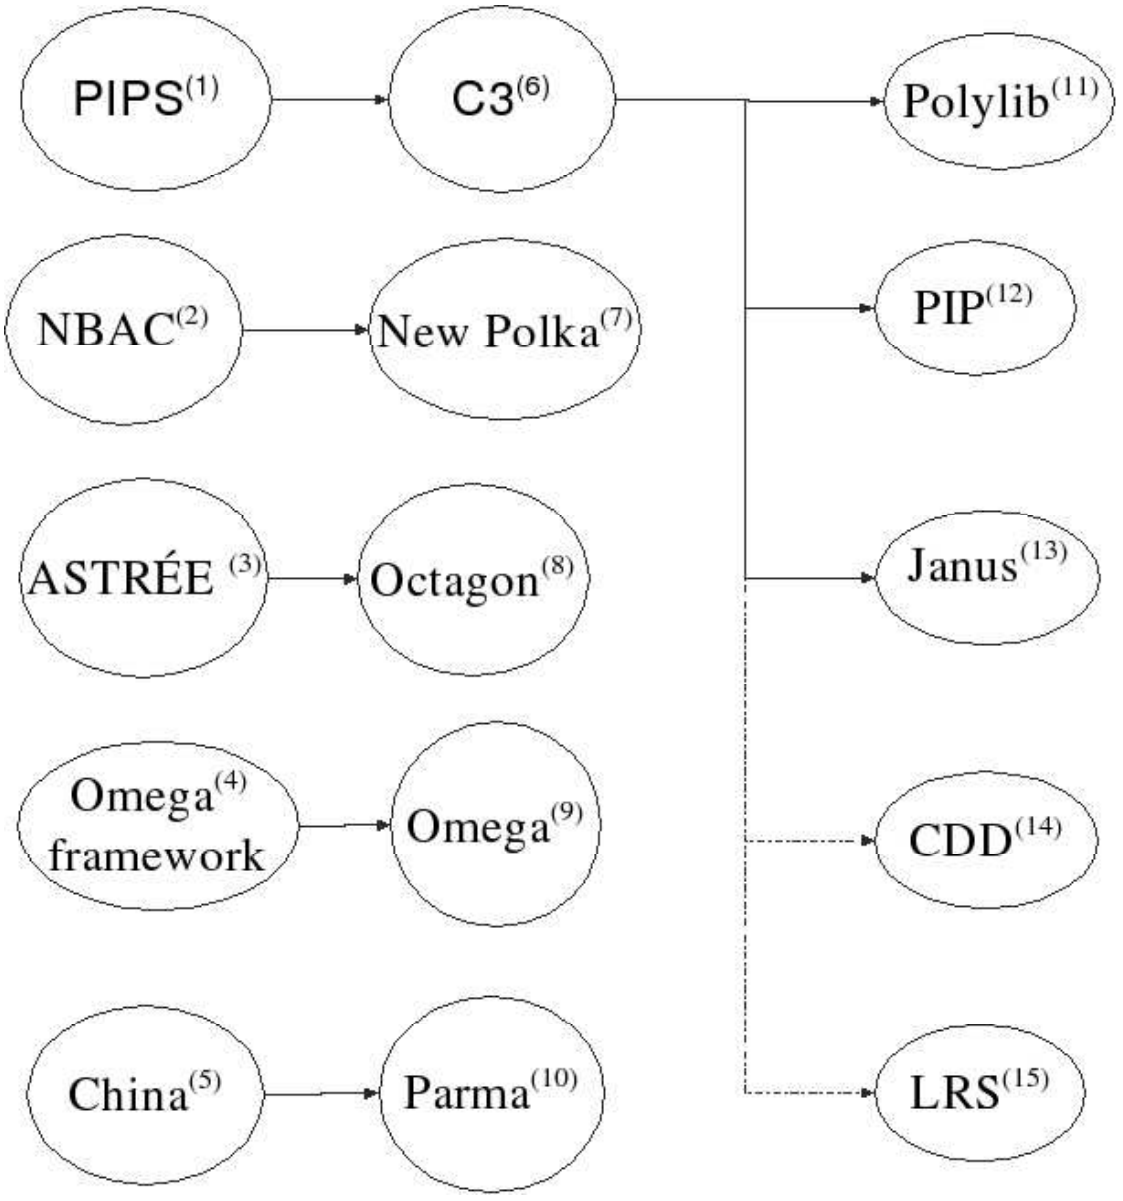

Polyhedral Benchmark Generation

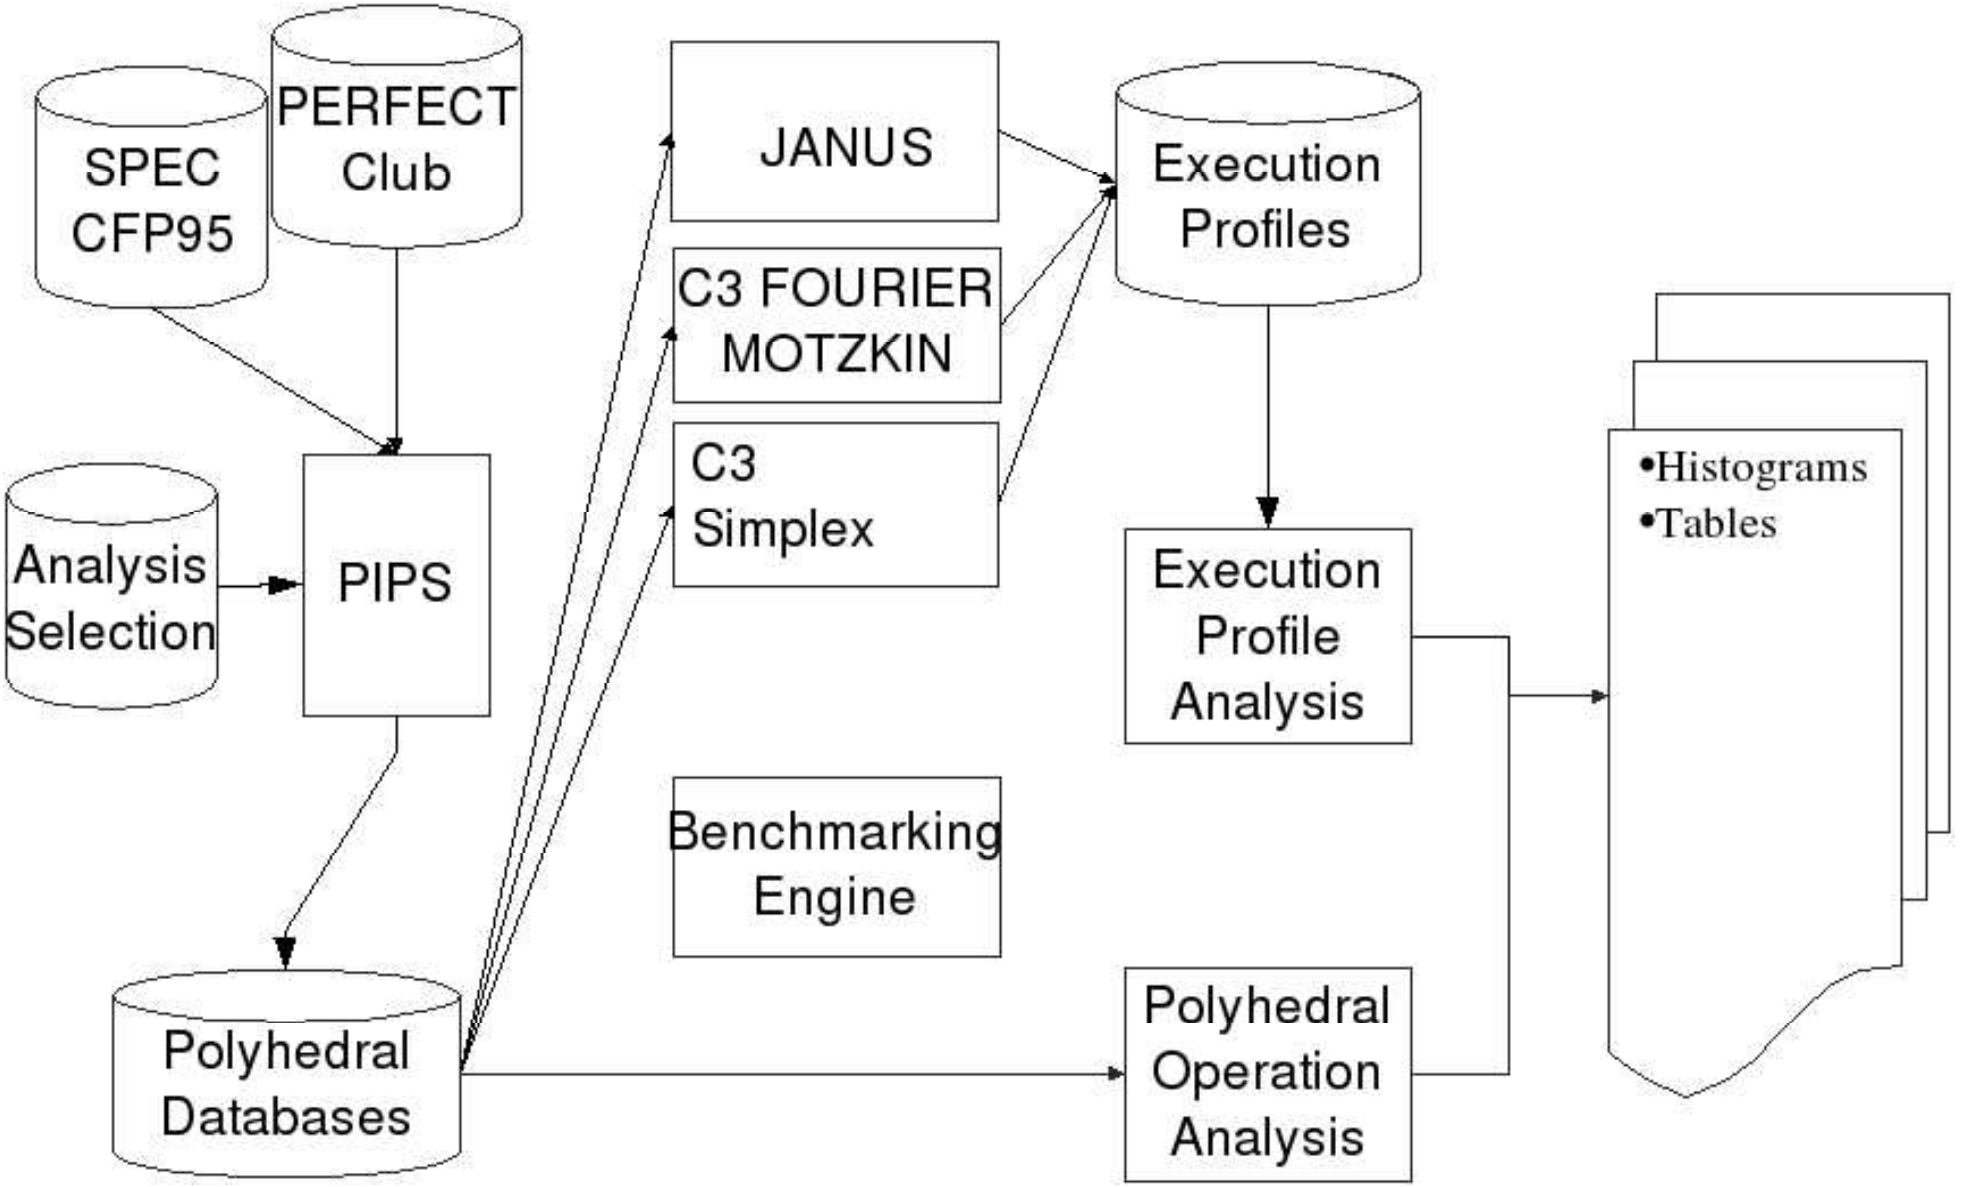

Notes on Analyzed Programs

- Input: CPU Benchmarks
  - PERFECT Club: 13 Fortran programs (61000 LOC)
  - SPEC CFP 95 : 10 Fortran programs (32000 LOC)
- Database generation:
  - Tracing calls of operators (e.g. satisfiability test, convex hull)
  - PIPS's analyses performed: Inter-procedural Precondition and Convex Array Region
  - Selection: sampling or filtering
- Output:
  - Histograms by criteria, head-to-head comparisons
  - Tables: statistics for operations, numbers of exceptions, comparison of global execution times

Examples of Database Distribution

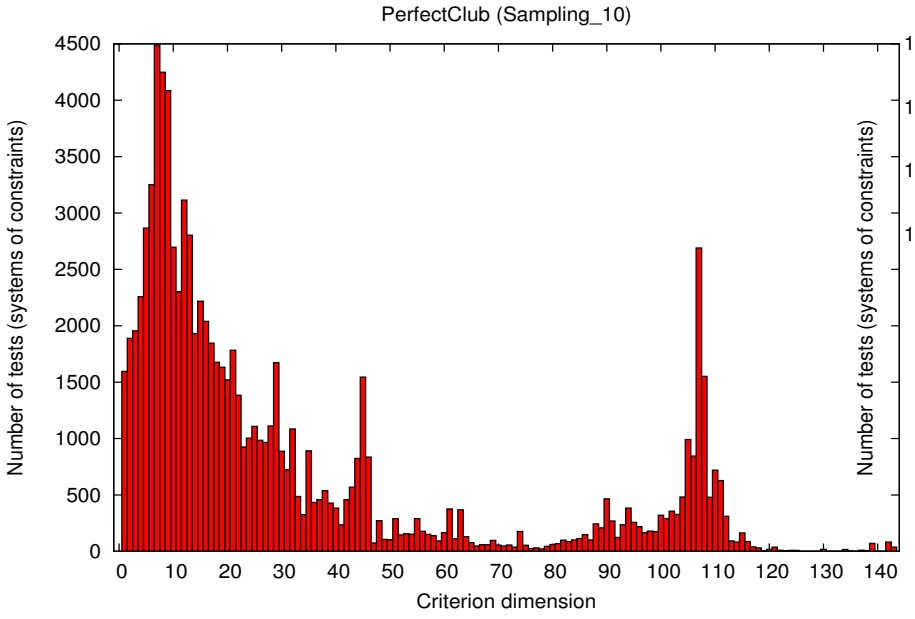

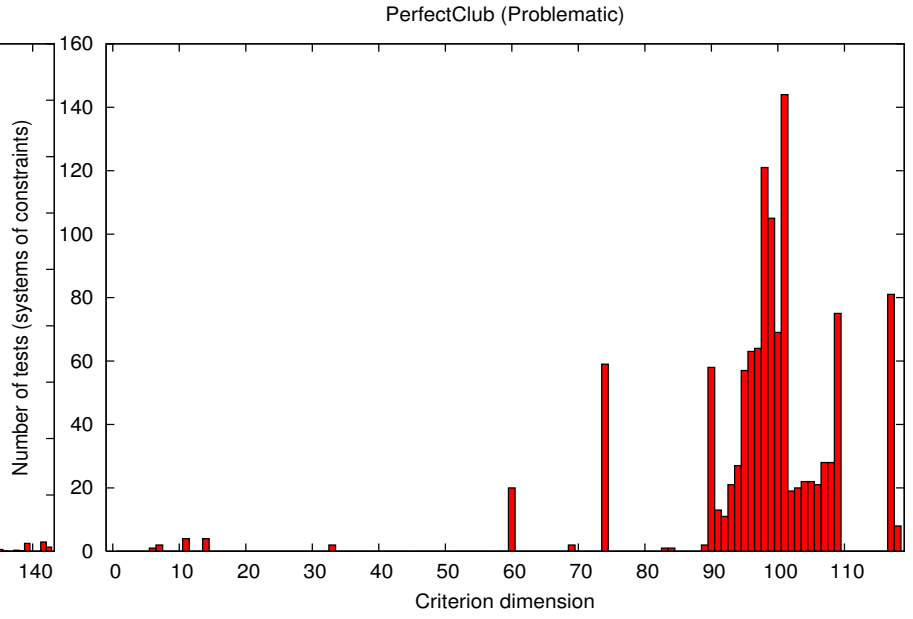

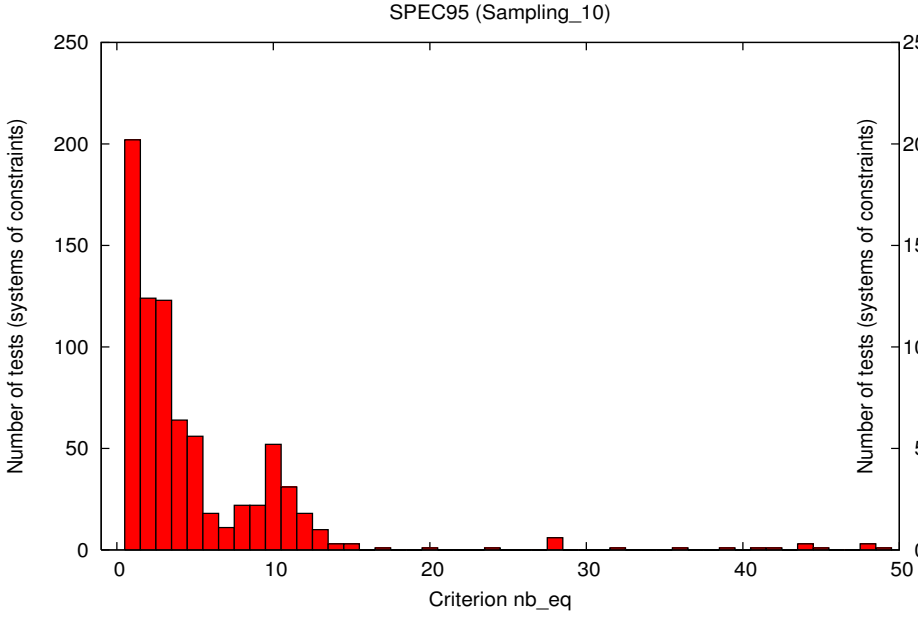

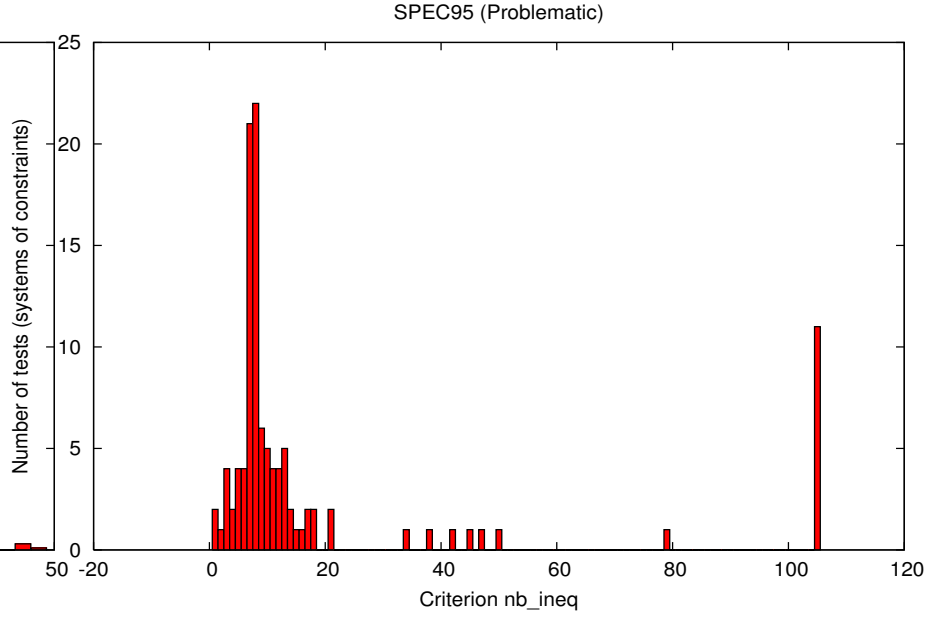

Impact of Filter: PerfectClub, Satisfiability

- Filter criterion used:
  - (space dimension > 60) & (#constraints > 70) & (#density_index > 150) & (max coefficient > 800)

	Sampling at 10 % (91825 tests)			Filtered (1556 tests)		
Parameters	Max	Min	Average	Max	Min	Average
File Size	25306	78	671.42	17879	382	5596.73
Space Dimension	100	1	18.25	67	1	46.97
#equations	85	0	12.40	33	0	17.40
#inegalities	437	0	6.43	104	1	38.89
#constraints	492	1	18.84	128	1	56.30
#sparsity_index	12	0	1.22	8	1	3.90
#vertices	12528	0	4.54	18432	0	1118.76
#rays	1017	0	2.70	876	0	69.10
#lines	29	0	1.18	23	0	6.78

Satisfiability

- Janus 64-bit - INRIA
  - 3400 LOC
  - integer solution
- C3 Simplex - CRI
  - 1488 LOC
  - rational solution
- C3 Fourier-Motzkin (FM) - CRI
  - 1000 LOC
  - integer/rational solution
- Comparison criteria:
  - Execution time
  - Number of exceptions
- Timer resolution: 10 ms

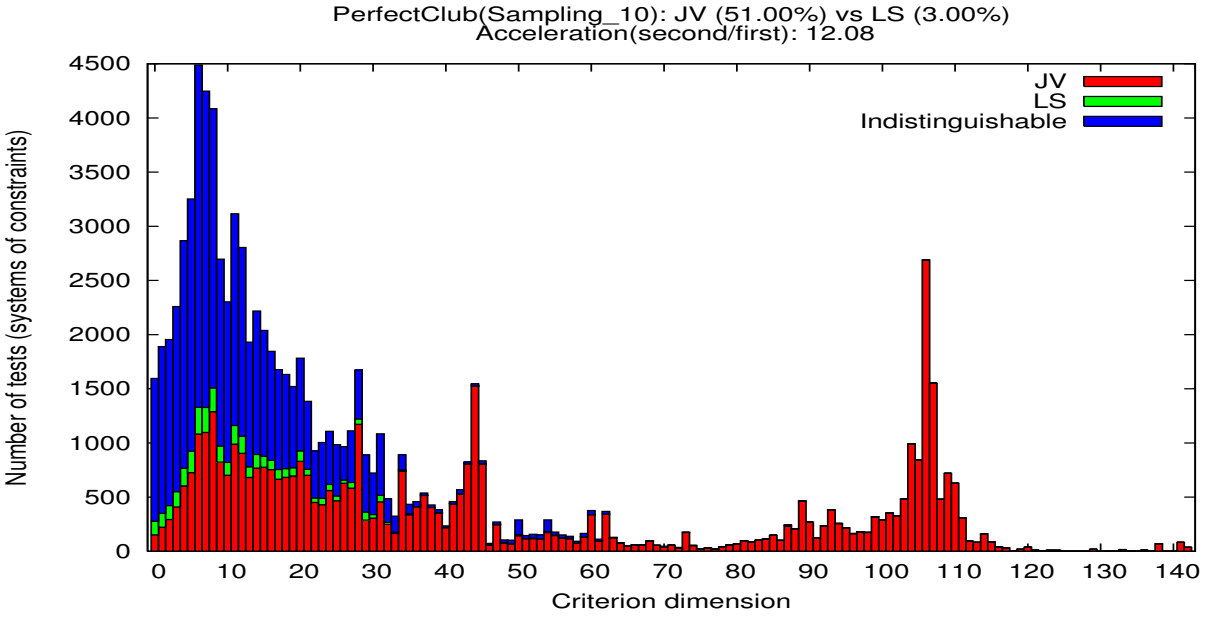

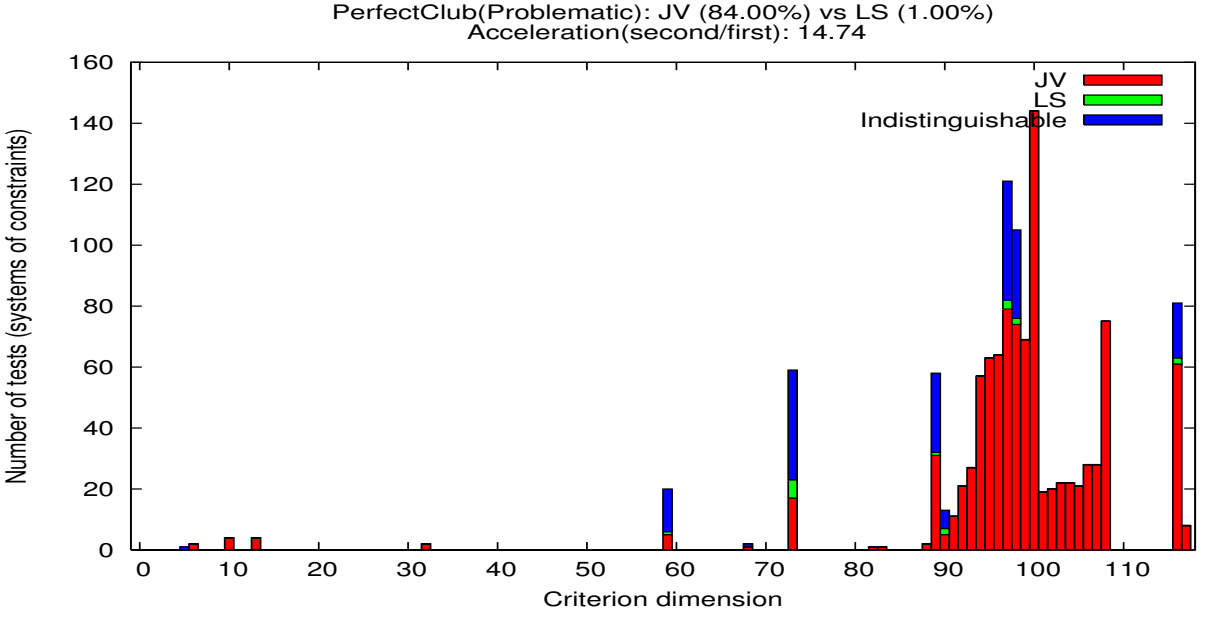

Integer/Rational Solution: Accuracy

- Janus is more accurate for integer solution

PerfectClub	Sampling at 10 % (91810 tests)	Filtered (1177 tests)
Janus <> C3 Simplex	16	0
Janus <> C3 FM	4	0

SPEC95	Sampling at 10 % (10597 tests)	Filtered (4676 tests)
Janus <> C3 Simplex	4	21
Janus <> C3 FM	2	20

- Impact of analyzed programs

Satisfiability: Execution times

- Global execution times (in seconds)

PerfectClub	Sampling at 10% (91810 tests)		Filtered (1177 tests)	
Janus/C3 Simplex	682/8241	1/12	54/801	1/15
Janus/C3 FM	750/22276	1/30	55/1555	1/28

SPEC95	Sampling at 10% (10597 tests)		Filtered (4676 tests)	
Janus/C3 Simplex	30/258	1/8	44/708	1/16
Janus/C3 FM	33/238	1/7	5/159	1/30

Satisfiability: Exceptions

- Number of exceptions: Magnitude overflow and timeout of 2 minutes.

PerfectClub	Sampling at 10 % (91810 tests)		Filtered (1177 tests)	
	overflow	timeout	overflow	timeout
Janus	48	0	0	0
C3 Simplex	587	10	2	0
C3 FM	0	19	0	0

SPEC95	Sampling at 10 % (10597 tests)		Filtered (4676 tests)	
	overflow	timeout	overflow	timeout
Janus	75	0	2987	0
C3 Simplex	65	0	3174	0
C3 FM	0	12	8	620

- Janus overall winner

Dual Conversion

- Polylib - IRISA
  - > 4000 LOC (Dual conversion: 435 LOC)
  - Format conversion 60 LOC (CRI)
- CDD - IOR, Zürich, Switzerland
  - > 4000 LOC (Dual conversion: 400 LOC)
- Comparison criteria:
  - Execution time
  - Number of exceptions
- Timer resolution: 10 ms

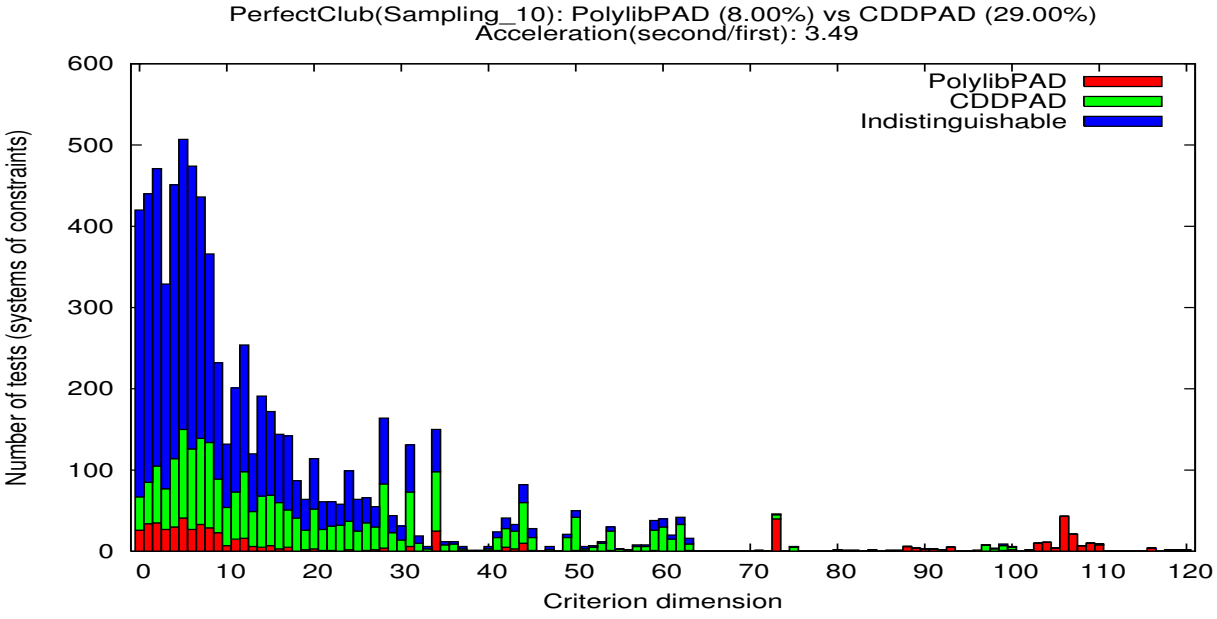

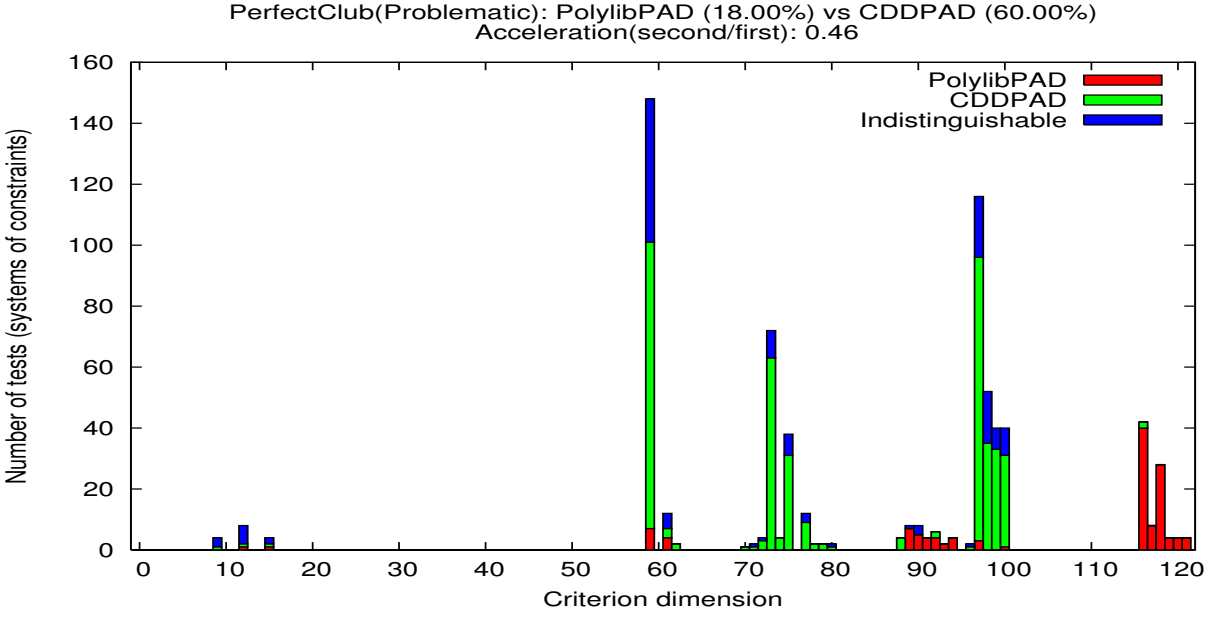

Dual Conversion: Execution times

- Global execution times (in seconds)

PerfectClub	Sampling at 10 % (7550 tests)		Filtered (693 tests)	
Polylib/CDD	41/145	1/3	457/210	1/0.5
Parallel/Polylib/CDD	15	1/3/9	139	1/3/1.5

SPEC95	Sampling at 10 % (4652 tests)		Filtered (178 tests)	
Polylib/CDD	24/11	1/0.5	4.45/3.7	1/0.8
Parallel/Polylib/CDD	7	1/3/1.5	3	1/1.5/1.2

Dual Conversion: Exceptions

- Number of exceptions (magnitude overflow and timeout of 2 minutes)

PerfectClub	Sampling at 10 % (7550 tests)		Filtered (693 tests)	
	overflow	timeout	overflow	timeout
CDD	0	3	0	0
Polylib	0	0	0	0

SPEC95	Sampling at 10 % (4652 tests)		Filtered (178 tests)	
	overflow	timeout	overflow	timeout
CDD	0	0	0	0
Polylib	2	0	0	0

- Heuristics?

Conclusion

- Results:
  - First large-scale and real-life experimental evaluation of key polyhedral operators
  - No unique winner: heuristics required
  - Unexpected sensibility to experimental data sets
  - Limited interest of “hard” cases
  - Bug detection and non-regression testing
- Future work
  - More operators
  - More implementations