Understanding the Performance of Highly Configurable Systems

Sergiy Kolesnikov, Norbert Siegmund, Christian Kästner*, Sven Apel

University of Passau, *Carnegie Mellon University

FOSD Meeting 2015, Traunkirchen

Highly Configurable Systems

192 configurations

Apache

eclipse

2560 configurations

LLVM

1152 configurations

Performance Models and Their Properties

performance model

What is the performance of Config. 100? Apache

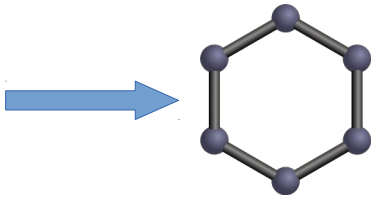

3.14159265358979323846264338

Accuracy

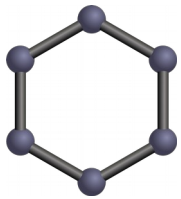

Complexity

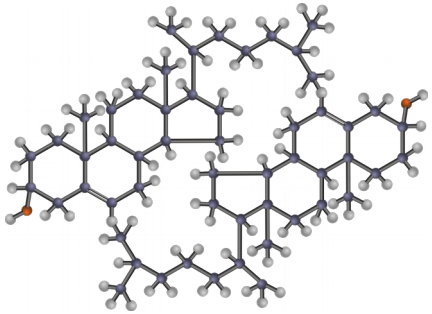

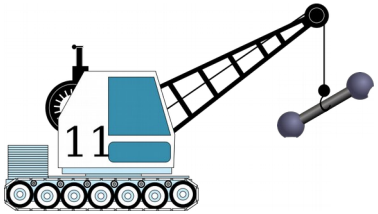

Compute time

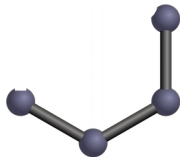

Real-World Example

- HSMGP highly scalable multi-grid solver for large datasets
  - 32 configuration options
  - 3456 configurations
  - Performance = Solving time

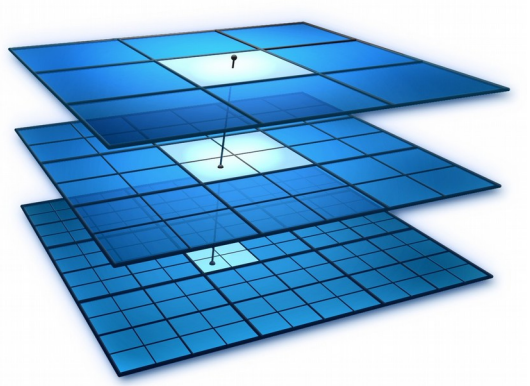

- Benchmarked all configurations
- Applied a configurable machine-learning algorithm to learn the influence of the configuration options on the system's performance

HSMGP Performance Models

- **Model 1:**

1898.97·CGS + 4579.98·GSACBE + 254.46·GSRB + -513.26·GS + -409.03·GSAC + 257.11·numCores_4096 + -166.14·CGS_IP_AMG + 1944.30·GS·numPre_0·numPost_1 + 1996.71·numPre_1·numPost_0 + -4826.99·GSACBE·numPre_1·numPost_0 + 2326.29·GSRB·numPre_0·numPost_1 + -1674.72·GSACBE·numPost_1 + 1791.76·GSACBE·numPre_6 + 1170.06·GSACBE·numPre_5 + -1416.57·GSRBAC·numPre_1·numPost_0 + 2087.99·JAC·numPre_0·numPost_1 + 85.81·numCores_1024 + -438.45·GSRBAC·numPre_1 + -2497.00·GSACBE·numPre_0·numPost_2 + 363.91·GSRBAC·numPost_6 + -2918.50·GSACBE·numPre_2·numPost_0 + -379.99·GSRBAC·numPre_2 + 478.58·numPre_2·numPost_0 + -1139.78·GSAC·numPre_1·numPost_0 + 232.59·numPre_6·numPost_6 + 689.11·GS·numPre_1·numPost_1 + 260.11·GSRBAC·numPost_5 + -233.95·GSRBAC·numPre_3 + 913.87·GSRB·numPre_1·numPost_1 + -95.52·CGS_IP_CG·numCores_64 + 698.16·GS·numPre_0·numPost_2 + -1971.52·GSACBE·numPre_2·numPost_2 + 879.54629303334·GSRB·numPre_0·numPost_2 + 811.11·GSAC·numPre_0·numPost_1 + -1865.21·GSACBE·numPre_1·numPost_2 + -1852.20·GSACBE·numPre_3·numPost_0 + 2259.25·GSACBE·numPre_0·numPost_6 + 2205.56·GSACBE·numPre_4·numPost_4 + 2204.54·GSACBE·numPre_3·numPost_5 + 2206.79·GSACBE·numPre_2·numPost_6

Error: 6%
Time: 25 days

- **Model 2:**

$$1945.1 \cdot \mathrm{CGS} + 4607.5 \cdot \mathrm{GSACBE} + 364.5 \cdot \mathrm{GSRB} - 458.3 \cdot \mathrm{GS} - 391.6 \cdot \mathrm{GSAC} + 239.1 \cdot \mathrm{Cores} - 154.2 \cdot \mathrm{CGSIP}$$

Error: 13%
Time: 10 hours

What are the Tradeoffs?

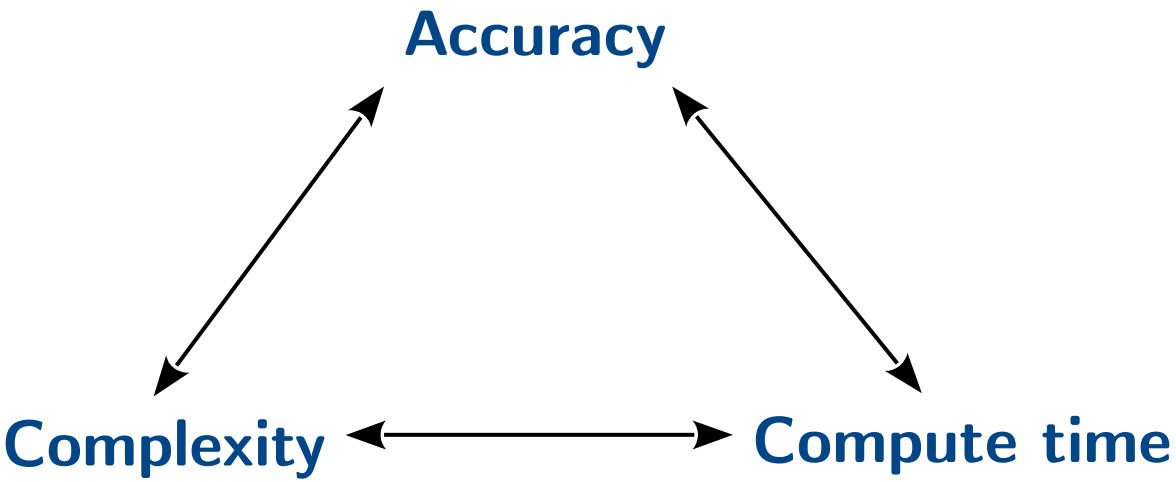

- Adapted an existing linear regression algorithm to learn performance models
- Empirical study on 10 highly-configurable real-world systems to quantify the tradeoffs

Example DBMS

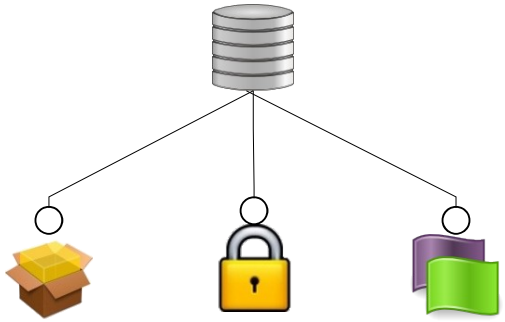

Id	Config.	Perf.
1		100
2		90
3		70
4		100
5		80
6		90
7		70
8		80

Performance – write speed in MB/s

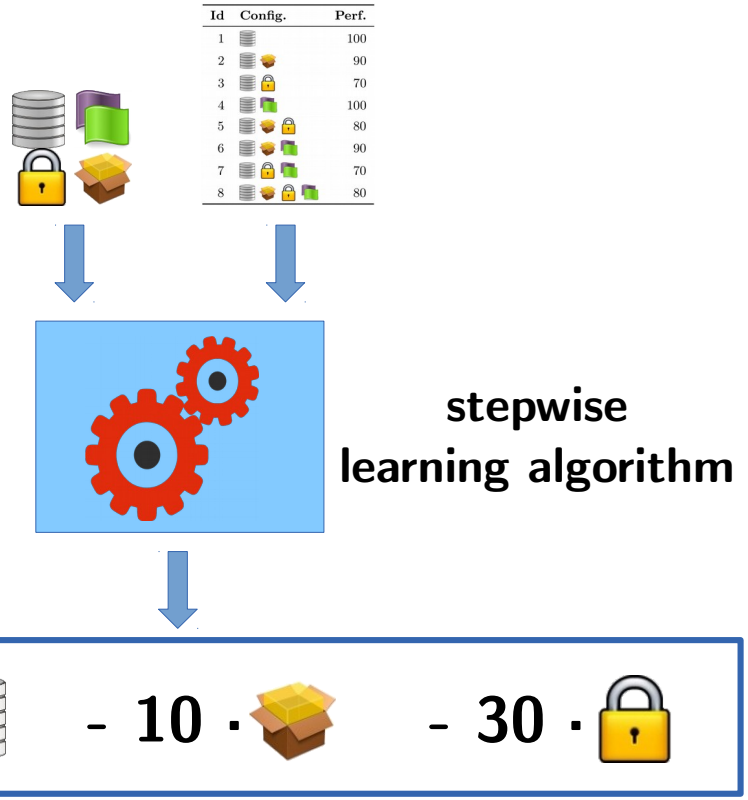

Performance of configuration 2:
$100 \cdot \mathit{1} - 10 \cdot \mathit{1} - 30 \cdot \mathit{0} = 90$

Performance of configuration 5:
$100 \cdot \mathit{1} - 10 \cdot \mathit{1} - 30 \cdot \mathit{1} = 60$

Stepwise Learning Algorithm

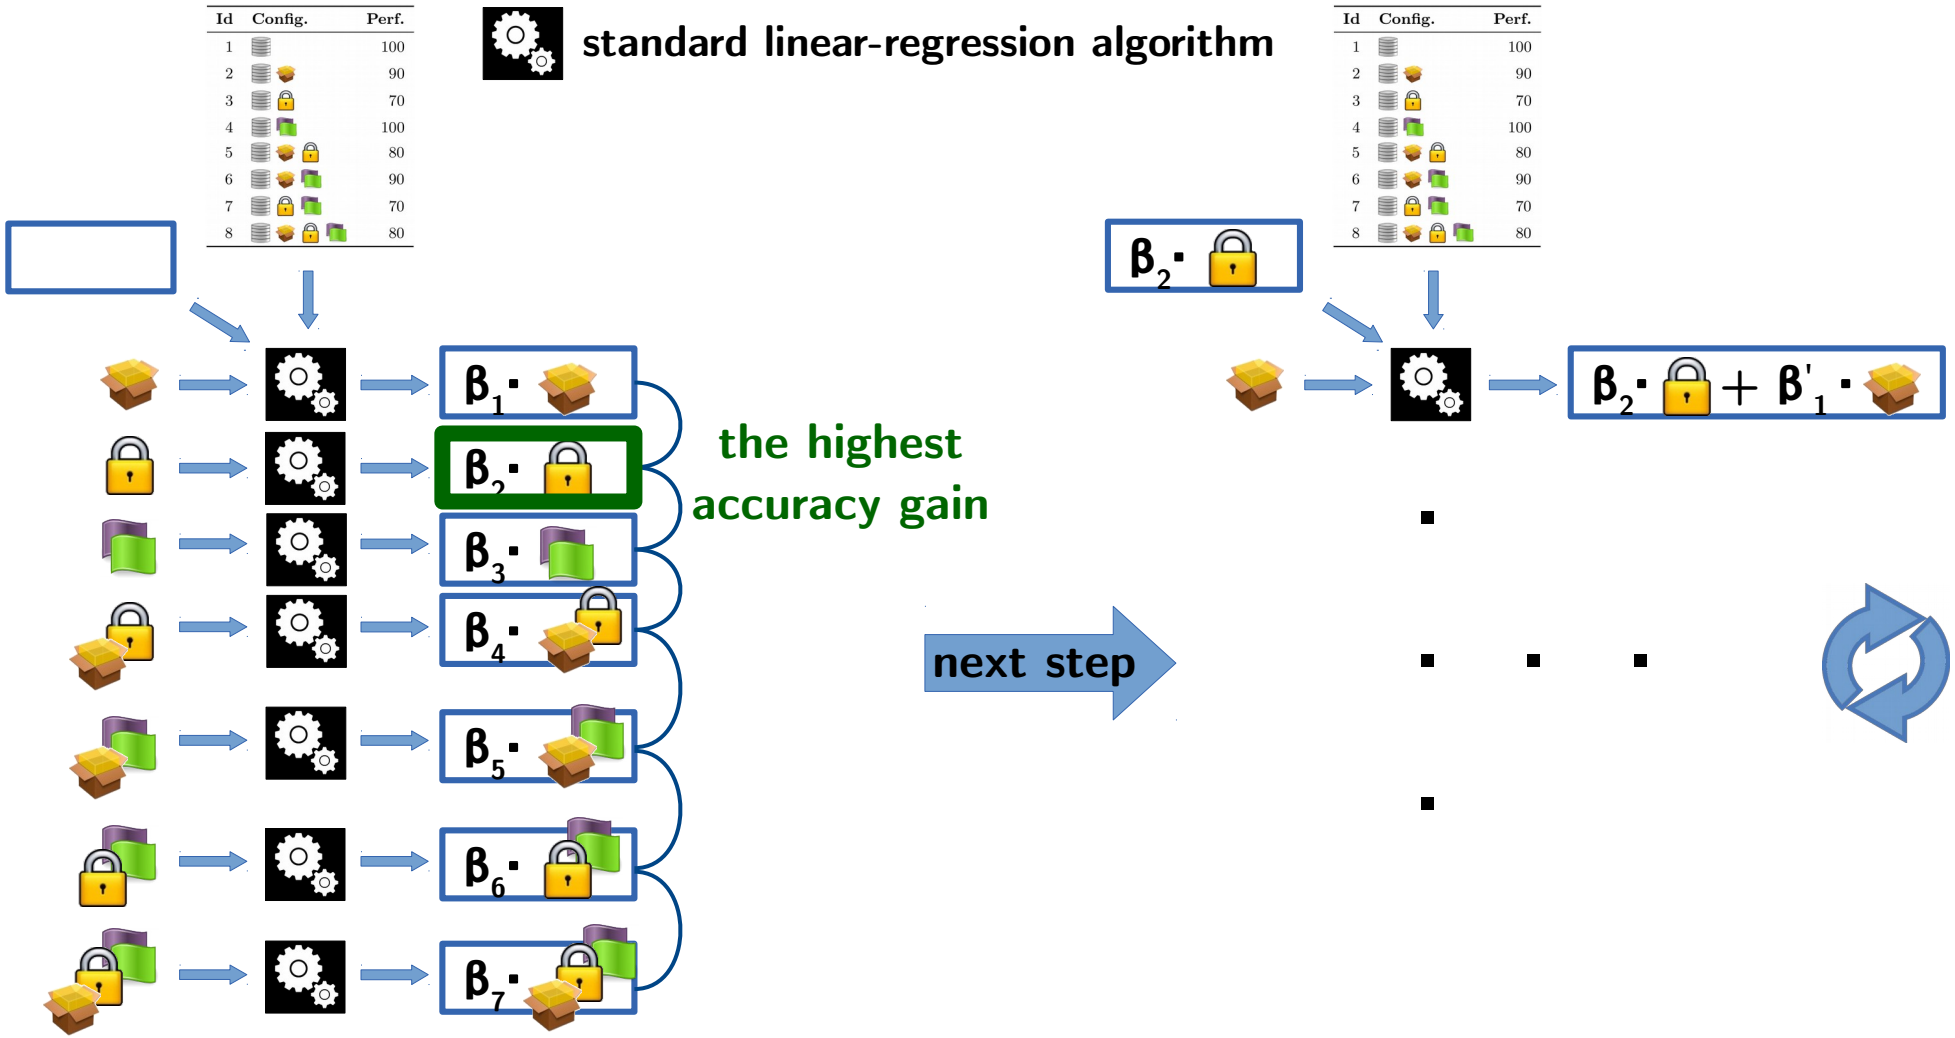

Hierarchical Stepwise Learning Algorithm

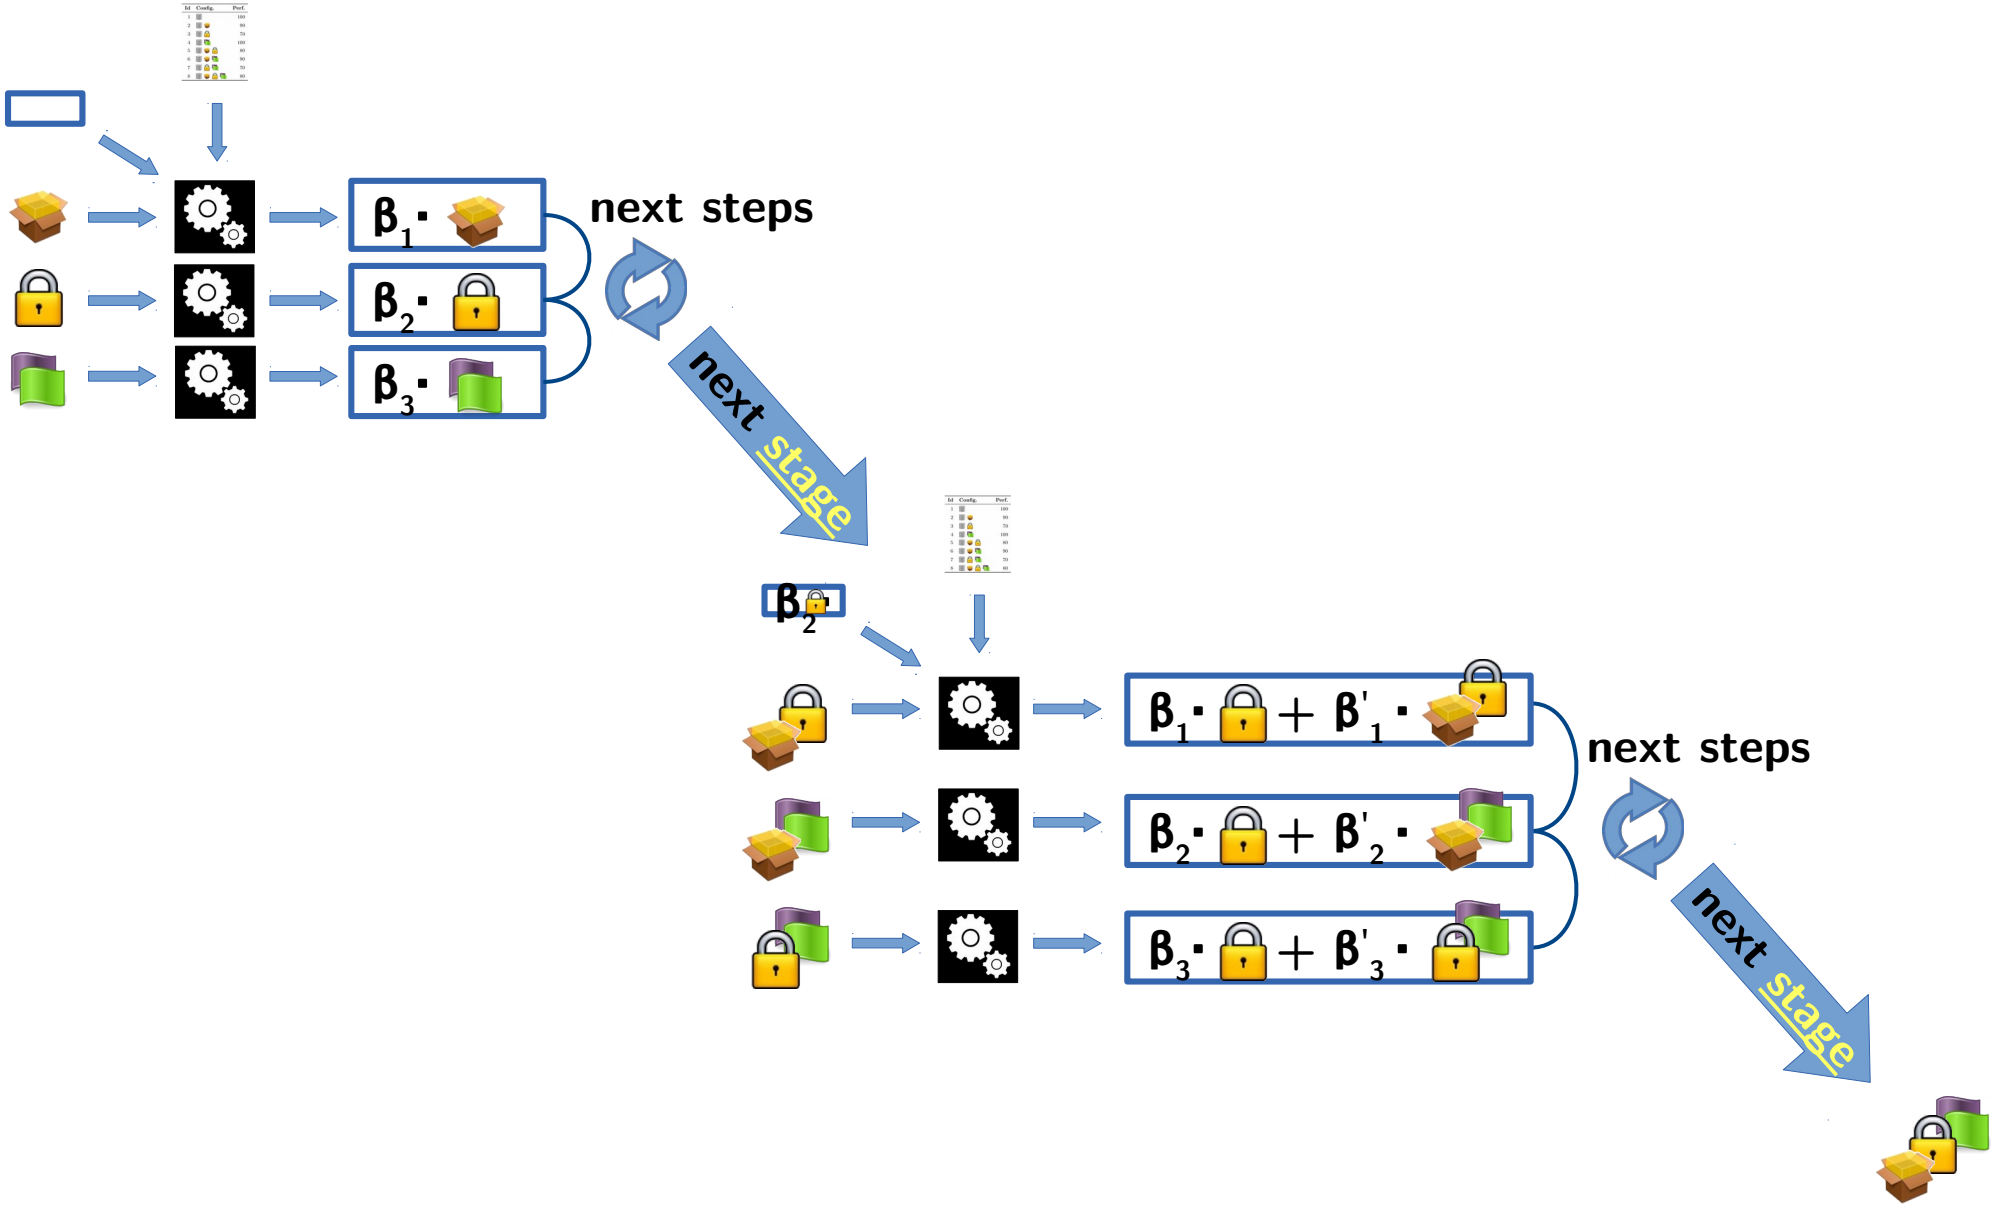

Candidate-Size Penalty

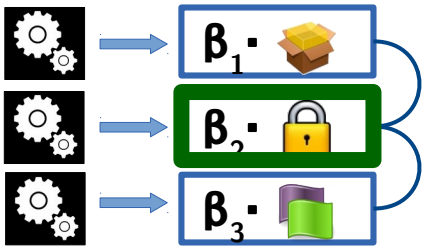

the highest accuracy gain

p() – size-penalty function

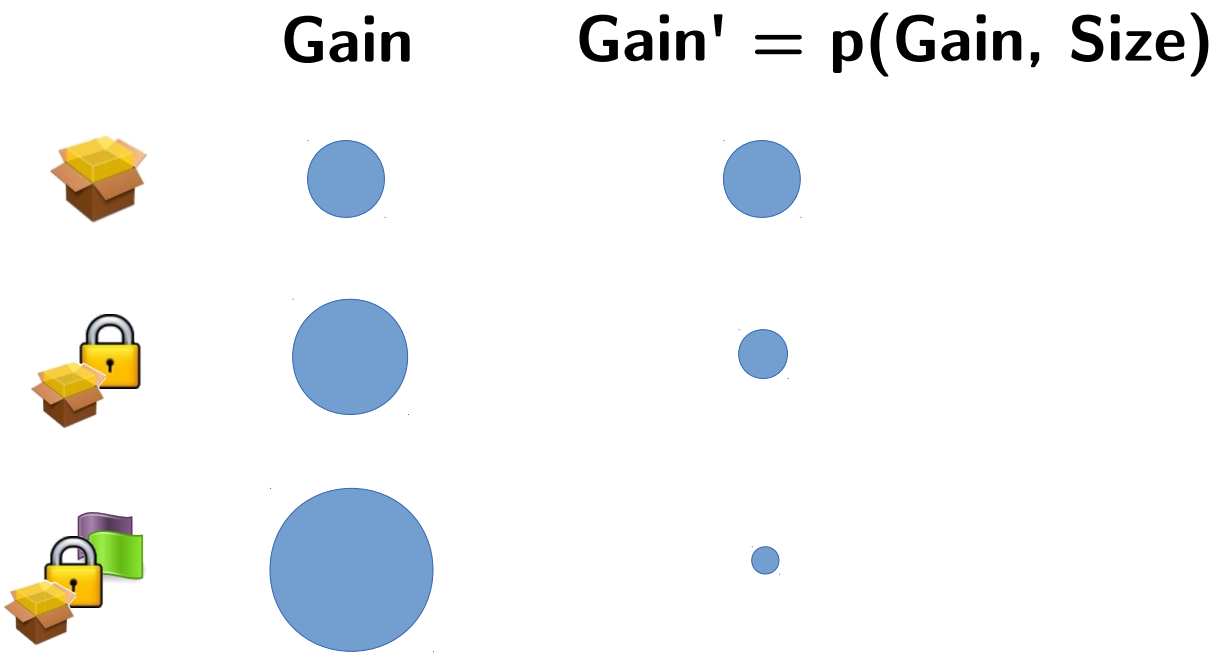

Algorithm Variants

1. **Non-hierarchical** **no penalty** **(nh)**
2. **Non-hierarchical** **penalty** **(nh,p)**
3. **Hierarchical** **no penalty** **(h)**
4. **Hierarchical** **penalty** **(h,p)**

Subject Systems

System	Domain	$\|O\|$	$\|C\|$	Performance
AJSTATS	stat. analysis	20	30 256	analysis time
APACHE	web-server	9	192	response rate
BDB-C	DBMS	18	2 560	I/O time
BDB-J	DBMS	26	180	I/O time
CLASP	ASP solver	19	700	solving time
DUNE	stencil code	31	2 304	solving time
HSMGP	stencil code	32	3 456	solving time
LLVM	compiler	11	1 024	optimization time
LRZIP	archiving tool	19	432	compression time
X264	video codec	16	1 152	encoding time

Results of the Experiments

System	Error				Complexity				Compute Time			
	nh	h	nh,p	h,p	nh	h	nh,p	h,p	nh	h	nh,p	h,p
AJSTATS	1.79	1.79	1.79	1.79	2	2	2	2	2.33 *h*	0.21	2.34 *h*	0.21
APACHE	3.4	3.41	3.71	3.75	19	13	14	8	0.32	0.02	0.27	0.06
BDB-C	58.11	35.67	58.33	42.31	39	12	31	41	2.53 *h*	3.83 *h*	1.98 *h*	47.25
BDB-J	2.58	2.58	2.58	2.58	5	5	5	5	16.19	0.04	16.17	0.04
CLASP	1.28	1.24	4.61	2.98	35	46	12	17	3.1 *h*	31.58	57.83	16.66
DUNE	6.05	6.58	8.56	9.8	91	76	25	15	9.52 *d*	18.28 *h*	2.47 *d*	7.5 *h*
HSMGP	6.61	7.07	12.73	12.73	111	96	7	7	24.82 *d*	1.43 *d*	1.62 *d*	9.54 *h*
LLVM	1.26	1.23	2.8	2.23	23	23	7	5	4.57	1.26	1.07	0.53
LRZIP	2.97	2.97	3.8	3.8	49	49	39	39	1.74 *h*	1.31	1.38 *h*	8.23
X264	0.77	0.55	1.24	1.51	31	39	18	17	1.26 *h*	13.56	45.36	7.15

Conclusion

- Tradeoffs among performance model accuracy, complexity, and compute time
- Empirical study on 10 highly-configurable real-world systems
- Low complexity and fast compute time with minor accuracy loss