

Software Analysis and Testing

Green Software


João Saraiva

2020/2021

Departamento de Informática
Universidade do Minho, Portugal



Going Green



2

Global energy system is unsustainable

The image shows a screenshot of a Bloomberg Business article. The article title is "Inside the Arctic Circle, Where Your Facebook Data Lives" by Janine Vance, dated October 04, 2013. The main image is a snowy landscape with a reindeer. Below the image, there is a caption: "Every year, computing giants including Hewlett-Packard (HP), Dell (DELL), and Cisco Systems (CSCO) sell north of \$100 billion in hardware. That's the total for the basic iron—". To the left of the article is an icon of a laptop and server rack, and to the right is an icon of a server rack.



3

Green Computing

- ▶ Caught the attention of many companies allowing them to save:



"close to 50% of the energy costs of an organization can be attributed to the IT departments"

- [PICMET, 2009]

"up to 90% of energy used by ICT hardware can be attributed to software"

- [The Greenhouse Gas Protocol Report, 2013]

4

Green Software

- ▶ Reducing energy consumption through software analysis and optimization
- ▶ Problem:
 - ▶ How to analyze
 - ▶ How to interpret
 - ▶ How to improve



Green Software

- ▶ Problems (extended to programmers):
 - ▶ How to analyze
 - ▶ How to interpret
 - ▶ How to improve



Mining questions about software energy consumption
- [MSR'14]

Integrated energy-directed test suite optimization
- [ISTA'14]

Seeds: A software engineer's energy-optimization decision support framework
- [ICSE'14]

5



6

Energy vs. Power

- ▶ *Power (w)* - rate (or effort) at which that work is done
- ▶ *Energy (J)* - amount of work done
- ▶ *Power* can change constantly while *Energy* is the accumulation



$$\text{Energy} = \text{Power} \times \text{Seconds}$$



$$360,000 \text{ J} = 100\text{W} \times 3,600\text{s}$$



7



Which languages are the most energy efficient?

6

Motivation

- ▶ Understanding the energy efficiency of programming languages

What paradigms are most energy efficient? Is a faster language always a more energy efficient one?

Are there languages which run slower while consuming less energy?

How much energy consumption is attributed to CPU? How much energy does memory usage consume?

Which are the most energy efficient languages?



Comparing Programming Languages



The Computer Language Benchmarks Game

* Formerly known as *The Great Computer Language Shootout*



The Computer Language Benchmarks Game (CLBG)

- ▶ Project to compare programming languages

- ▶ 28 different programming languages
- ▶ 13 different benchmarks

- ▶ Experts compete to code the fastest solution

- ▶ Same underlying algorithms

- ▶ Publicly available:

- ▶ Source Code
- ▶ Compiler Versions
- ▶ Optimization Flags

- ▶ Comparable and easily replicable programs/solutions!

Benchmark	Description	Input
n-body	Double precision N-body simulation	50M
fannkuch-redux	Indexed access to tiny integer sequence	12
spectral norm	Eigenvalue using the power method	5,500
mandelbrot	Generate Mandelbrot set portable bitmap file	16,000
pidigits	Streaming arbitrary precision arithmetic	10,000
regex-redux	Match DNA 8mers and substitute magic patterns	fasta output
fasta	Generate and write random DNA sequences	25M
k-nucleotide	Hashable update and k-nucleotide strings	fasta output
reverse-complement	Read DNA sequences, write their reverse-complement	fasta output
binary-trees	Allocate, traverse and deallocate many binary trees	21
chameneos-redux	Symmetrical thread rendezvous requests	6M
meteor-contest	Search for solutions to shape packing puzzle	2,098
thread-ring	Switch from thread to thread passing one token	50M



Design

Benchmark	Description	Input
n-body	Double precision N-body simulation	50M
fannkuch-redux	Indexed access to tiny integer sequence	12
spectral norm	Eigenvalue using the power method	5,500
mandelbrot	Generate Mandelbrot set portable bitmap file	16,000
pidigits	Streaming arbitrary precision arithmetic	10,000
regex-redux	Match DNA 8mers and substitute magic patterns	fasta output
fasta	Generate and write random DNA sequences	25M
k-nucleotide	Hashable update and k-nucleotide strings	fasta output
reverse-complement	Read DNA sequences, write their reverse-complement	fasta output
binary-trees	Allocate, traverse and deallocate many binary trees	21



Paradigm	Languages
Functional	Erlang, F#, Haskell, Lisp, Ocaml, Perl, Racket, Ruby, Rust;
Imperative	Ada, C, C++, F#, Fortran, Go, Ocaml, Pascal, Rust;
Object-Oriented	Ada, C++, C#, Chapel, Dart, F#, Java, JavaScript, Ocaml, Perl, PHP, Python, Racket, Rust, Smalltalk, Swift, TypeScript;
Scripting	Dart, Hack, JavaScript, JRuby, Lua, Perl, PHP, Python, Ruby, TypeScript;

Paradigm	Languages
Functional	Erlang, F#, Haskell, Lisp, Ocaml, Perl, Racket, Ruby, Rust;
Imperative	Ada, C, C++, F#, Fortran, Go, Ocaml, Pascal, Rust;
Object-Oriented	Ada, C++, C#, Chapel, Dart, F#, Java, JavaScript, Ocaml, Perl, PHP, Python, Racket, Rust, Smalltalk, Swift, TypeScript;
Scripting	Dart, Hack, JavaScript, JRuby, Lua, Perl, PHP, Python, Ruby, TypeScript;



Measurements

Energy

- ▶ Running Average Power Limit (RAPL)
- ▶ Designed by intel for i5/i7 architectures (SandyBridge, IvyBridge, Haswell, etc)
- ▶ Monitors energy consumption info for Machine-Specific Registers (MSRs)
- ▶ Allows very precise and fine-grain measurements through function calls
 - ▶ DRAM/GPU
 - ▶ CPU
 - ▶ Package


Peak Memory

- ▶ Unix's time tool



Execution

- ▶ Linux Ubuntu 16.0 4.8.0-22-generic
- ▶ 16GB Ram
- ▶ Intel(R) Core(TM) i5-4460 CPU @ 3.40GHz
- ▶ Compiler versions: <https://sites.google.com/view/energy-efficiency-languages/setup>
- ▶ Source Code: <http://benchmarksgame.alioth.debian.org/>




- ▶ Compile -> Run -> Measure

- ▶ 27 Languages x 10 Benchmarks = 270 Results



Results



Energy vs. Time

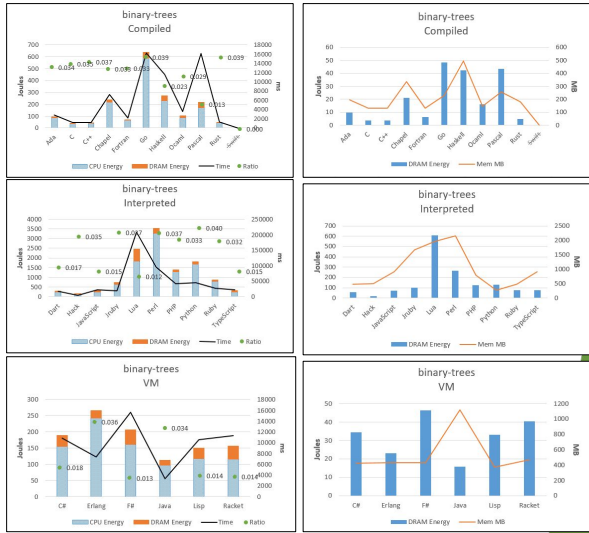
Energy vs. Memory

Energy vs. Time vs. Memory



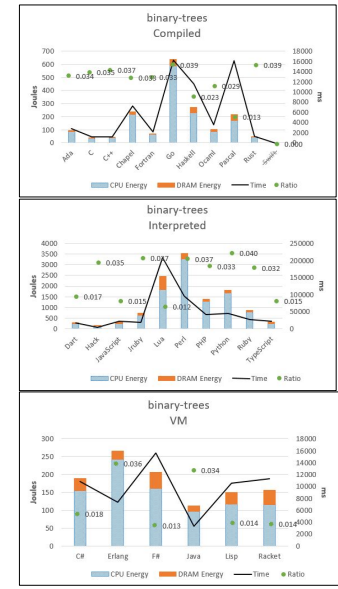
Energy vs. Time

binary-trees				
	Energy	Time	Ratio	Mb
(c) C	39.80	1125	0.035	131
(c) C++	41.23	1129	0.037	132
(c) Rust ↓ ₂	49.07	1263	0.039	180
(c) Fortran ↑ ₁	69.82	2112	0.033	133
(c) Ada ↓ ₁	95.02	2822	0.034	197
(c) Ocaml ↓ ₁ ↑ ₂	100.74	3525	0.029	148
(v) Java ↑ ₁ ↓ ₁₆	111.84	3306	0.034	1120
(v) Lisp ↓ ₃ ↓ ₃	149.55	10570	0.014	373
(v) Racket ↓ ₄ ↓ ₆	155.81	11261	0.014	467
(i) Hack ↑ ₂ ↓ ₉	156.71	4497	0.035	502
(v) C# ↓ ₁ ↓ ₁	189.74	10797	0.018	427
(v) F# ↓ ₃ ↓ ₁	207.13	15637	0.013	432
(c) Pascal ↓ ₃ ↑ ₅	214.64	16079	0.013	256
(c) Chapel ↑ ₅ ↑ ₄	237.29	7265	0.033	335
(v) Erlang ↑ ₅ ↑ ₁	266.14	7327	0.036	433
(c) Haskell ↑ ₂ ↓ ₂	270.15	11582	0.023	494
(i) Dart ↓ ₁ ↑ ₁	290.27	17197	0.017	475
(i) JavaScript ↓ ₂ ↓ ₂	312.14	21349	0.015	916
(i) TypeScript ↓ ₂ ↓ ₂	315.10	21686	0.015	915
(c) Go ↑ ₃ ↑ ₁₃	636.71	16292	0.039	228
(i) Jruby ↑ ₂ ↓ ₃	720.53	19276	0.037	1671
(i) Ruby ↑ ₅	855.12	26634	0.032	482
(i) PHP ↑ ₃	1,397.51	42316	0.033	786
(i) Python ↑ ₁₅	1,793.46	45003	0.040	275
(i) Lua ↓ ₁	2,452.04	209217	0.012	1961
(i) Perl ↑ ₁	3,542.20	96097	0.037	2148
(c) Swift		n.e.		



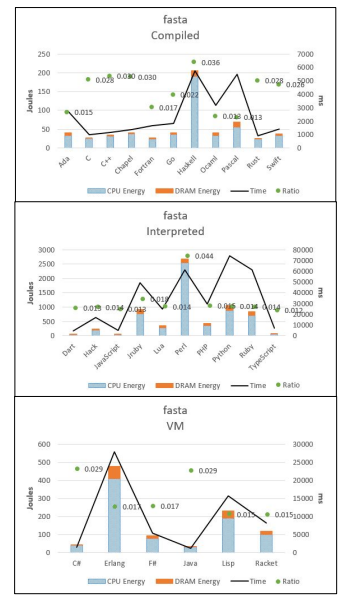
Energy vs. Time

binary-trees				
	Energy	Time	Ratio	Mb
(c) C	39.80	1125	0.035	131
(c) C++	41.23	1129	0.037	132
(c) Rust ↓ ₂	49.07	1263	0.039	180
(c) Fortran ↑ ₁	69.82	2112	0.033	133
(c) Ada ↓ ₁	95.02	2822	0.034	197
(c) Ocaml ↓ ₁ ↑ ₂	100.74	3525	0.029	148
(v) Java ↑ ₁ ↓ ₁₆	111.84	3306	0.034	1120
(v) Lisp ↓ ₃ ↓ ₃	149.55	10570	0.014	373
(v) Racket ↓ ₄ ↓ ₆	155.81	11261	0.014	467
(i) Hack ↑ ₂ ↓ ₉	156.71	4497	0.035	502
(v) C# ↓ ₁ ↓ ₁	189.74	10797	0.018	427
(v) F# ↓ ₃ ↓ ₁	207.13	15637	0.013	432
(c) Pascal ↓ ₃ ↑ ₅	214.64	16079	0.013	256
(c) Chapel ↑ ₅ ↑ ₄	237.29	7265	0.033	335
(v) Erlang ↑ ₅ ↑ ₁	266.14	7327	0.036	433
(c) Haskell ↑ ₂ ↓ ₂	270.15	11582	0.023	494
(i) Dart ↓ ₁ ↑ ₁	290.27	17197	0.017	475
(i) JavaScript ↓ ₂ ↓ ₂	312.14	21349	0.015	916
(i) TypeScript ↓ ₂ ↓ ₂	315.10	21686	0.015	915
(c) Go ↑ ₃ ↑ ₁₃	636.71	16292	0.039	228
(i) Jruby ↑ ₂ ↓ ₃	720.53	19276	0.037	1671
(i) Ruby ↑ ₅	855.12	26634	0.032	482
(i) PHP ↑ ₃	1,397.51	42316	0.033	786
(i) Python ↑ ₁₅	1,793.46	45003	0.040	275
(i) Lua ↓ ₁	2,452.04	209217	0.012	1961
(i) Perl ↑ ₁	3,542.20	96097	0.037	2148
(c) Swift		n.e.		

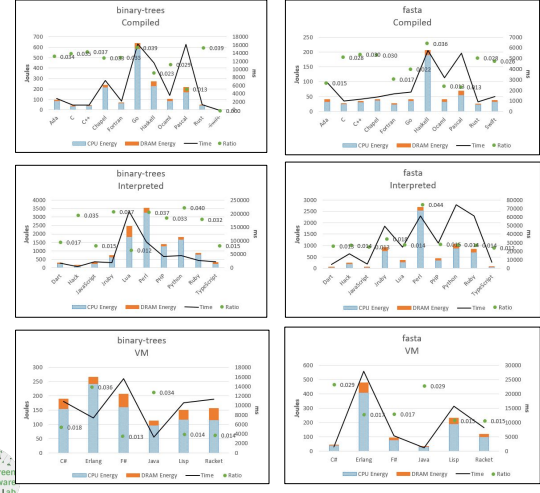


Energy vs. Time

fasta				
	Energy	Time	Ratio	Mb
(c) Rust ↓ ₉	26.15	931	0.028	16
(c) Fortran ↓ ₆	27.62	1661	0.017	1
(c) C ↑ ₁ ↓ ₁	27.64	973	0.028	3
(c) C++ ↑ ₁ ↓ ₂	34.88	1164	0.030	4
(v) Java ↑ ₁ ↓ ₁₂	35.86	1249	0.029	41
(c) Swift ↓ ₉	37.06	1405	0.026	31
(c) Go ↓ ₂	40.45	1838	0.022	4
(c) Ada ↓ ₂ ↑ ₃	40.45	2765	0.015	3
(c) Ocaml ↓ ₂ ↓ ₁₅	40.78	3171	0.013	201
(c) Chapel ↑ ₅ ↓ ₁₀	40.88	1379	0.030	53
(v) C# ↑ ₄ ↓ ₅	45.35	1549	0.029	35
(i) Dart ↓ ₆	63.61	4787	0.013	49
(i) JavaScript ↓ ₁	64.84	5098	0.013	30
(c) Pascal ↓ ₁ ↑ ₁₃	68.63	5478	0.013	0
(i) TypeScript ↓ ₂ ↓ ₁₀	82.72	6909	0.012	271
(v) F# ↑ ₂ ↑ ₃	93.11	5360	0.017	27
(v) Racket ↓ ₁ ↑ ₅	120.90	8255	0.015	21
(c) Haskell ↑ ₂ ↓ ₈	205.52	5728	0.036	446
(v) Lisp ↓ ₂	231.49	15763	0.015	75
(i) Hack ↓ ₂	237.70	17203	0.014	120
(i) Lua ↑ ₁₈	347.37	24617	0.014	3
(i) PHP ↓ ₁ ↑ ₁₃	430.73	29508	0.015	14
(v) Erlang ↑ ₁ ↑ ₁₂	477.81	27852	0.017	18
(i) Ruby ↓ ₁ ↑ ₂	852.30	61216	0.014	104
(i) JRuby ↑ ₁ ↓ ₂	912.93	49509	0.018	705
(i) Python ↓ ₁ ↑ ₁₈	1,061.41	74111	0.014	9
(i) Perl ↑ ₁ ↑ ₈	2,684.33	61463	0.044	53



Energy vs. Time



CPU-based Energy			
	Avg	Min	Max
Compiled	88.94%	85.27%	91.75%
Interpreted	87.98%	81.57%	92.90%
VM	88.94%	86.10%	92.43%



Energy vs. Time

Total			
	Energy	Time	Mb
(c) C	1.00	1.00	1.00
(c) Rust	1.03	1.04	1.05
(c) C++	1.34	1.17	1.17
(c) Ada	1.70	1.24	1.24
(v) Java	1.98	1.34	1.34
(c) Pascal	2.14	1.47	1.47
(c) Chapel	2.18	1.54	1.54
(v) Lisp	2.27	1.92	1.92
(c) Ocaml	2.40	2.45	2.45
(c) Fortran	2.52	2.57	2.57
(c) Swift	2.79	2.71	2.71
(c) Haskell	3.10	2.80	2.80
(v) C#	3.14	2.82	2.82
(c) Go	3.23	2.85	2.85
(i) Dart	3.83	3.34	3.34
(v) F#	4.13	3.52	3.52
(i) JavaScript	4.45	3.97	3.97
(v) Racket	7.91	4.00	4.00
(i) TypeScript	21.50	4.25	4.25
(i) Hack	24.02	4.59	4.59
(i) PHP	29.30	4.69	4.69
(v) Erlang	42.23	6.01	6.01
(i) Lua	45.98	6.62	6.62
(i) Jruby	46.54	6.72	6.72
(i) Ruby	69.91	7.20	7.20
(i) Python	75.88	8.64	8.64
(i) Perl	79.58	19.84	19.84

Average	
Joules	ms
Compiled	120
VM	576
Interpreted	2365

Average	
Joules	ms
Imperative	125
OO	879
Functional	1367
"Scripting"	2320

Energy vs. Time

	Energy
(c) C	1.00
(c) Rust	1.03
(c) C++	1.34
(c) Ada	1.70
(v) Java	1.98
(c) Pascal	2.14
(c) Chapel	2.18
(v) Lisp	2.27
(c) Ocaml	2.40
(c) Fortran	2.52
(c) Swift	2.79
(c) Haskell	3.10
(v) C#	3.14
(c) Go	3.23
(i) Dart	3.83
(v) F#	4.13
(i) JavaScript	4.45
(v) Racket	7.91
(i) TypeScript	21.50
(i) Hack	24.02
(i) PHP	29.30
(v) Erlang	42.23
(i) Lua	45.98
(i) Jruby	46.54
(i) Ruby	69.91
(i) Python	75.88
(i) Perl	79.58

	regex-redux			
	Energy	Time	Ratio	Mb
(i) TypeScript ↓ ₂ ↓ ₁₂	24.65	2009	0.012	587
(c) C ↑ ₁	24.83	805	0.031	151
(i) JavaScript ↓ ₂ ↓ ₉	25.68	2096	0.012	525
(i) PHP ↑ ₂ ↓ ₁	34.57	1667	0.021	182
(c) Pascal ↓ ₁ ↑ ₄	35.20	2282	0.015	106
(i) Hack ↑ ₂ ↓ ₂	38.96	2052	0.019	268
(c) Rust	40.26	2287	0.018	218
(c) Chapel ↓ ₁₁	97.19	4534	0.021	1055
(c) Ada ↑ ₅	148.66	5157	0.029	157
(i) Python ↓ ₁	161.62	7116	0.023	429
(c) Ocaml ↓ ₃ ↓ ₆	172.43	12978	0.013	948
(c) C++ ↓ ₁ ↑ ₆	176.24	10656	0.017	216
(i) Ruby ↓ ₄ ↑ ₄	192.88	14282	0.014	305
(v) Java ↑ ₄ ↓ ₆	194.65	5694	0.034	1225
(i) Dart ↓ ₁ ↑ ₄	197.92	13485	0.015	459
(i) Perl ↑ ₄ ↑ ₁₃	236.24	7164	0.033	154
(i) Jruby ↑ ₂ ↓ ₄	348.44	13477	0.026	1369
(v) Racket ↓ ₁	358.20	26152	0.014	983
(v) C# ↑ ₁ ↑ ₃	522.59	14723	0.035	851
(c) Swift ↑ ₅	538.11	41703	0.013	677
(v) F# ↑ ₇	650.51	46905	0.014	667
(c) Haskell		n.a.		
(c) Fortran		n.a.		
(c) Go		n.e.		
(i) Lua		n.e.		
(v) Erlang		n.a.		
(v) Lisp		n.e.		

Average	
Joules	ms
Compiled	120
VM	576
Interpreted	2365

Average	
Joules	ms
Imperative	125
OO	879
Functional	1367
"Scripting"	2320




Energy vs. Memory

Total			
	Energy	Time	Mb
(c) C	1.00	1.00	1.00
(c) Rust	1.03	1.04	1.05
(c) C++	1.34	1.56	1.17
(c) Ada	1.70	1.85	1.24
(v) Java	1.98	1.89	1.34
(c) Pascal	2.14	2.14	1.47
(c) Chapel	2.18	2.83	1.47
(v) Lisp	2.27	3.02	1.92
(c) Ocaml	2.40	3.09	2.45
(c) Fortran	2.52	3.14	2.57
(c) Swift	2.79	3.40	2.71
(c) Haskell	3.10	3.55	2.80
(v) C#	3.14	4.20	2.82
(c) Go	3.23	4.20	2.85
(i) Dart	3.83	6.30	3.34
(v) F#	4.13	6.30	3.52
(i) JavaScript	4.45	6.67	3.97
(v) Racket	7.91	11.27	4.00
(i) TypeScript	21.50	26.99	4.25
(i) Hack	24.02	27.64	4.59
(i) PHP	29.30	36.71	4.69
(v) Erlang	42.23	43.44	6.01
(i) Lua	45.98	46.20	6.62
(i) Jruby	46.54	59.34	6.72
(i) Ruby	69.91	65.79	7.20
(i) Python	75.88	71.90	8.64
(i) Perl	79.58	82.91	19.84

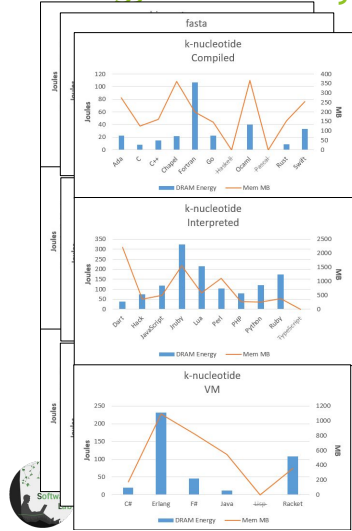
Average	
Mb	
Compiled	125
VM	285
Interpreted	426

Average	
Mb	
Imperative	116
OO	249
Functional	251
"Scripting"	421

Energy vs. Memory




Energy vs. Memory



► Spearman rank-order correlation coefficient

► Spearman $\rho = 0.2091$



Energy vs. Time vs. Memory

Total			
	Energy	Time	Mem
(e) C	1.00	1.00	(e) Pascal 1.00
(e) Rust	1.03	1.04	(e) Go 1.05
(e) C++	1.34	1.56	(e) C 1.17
(e) Ada	1.70	1.85	(e) Fortran 1.24
(v) Java	1.98	1.89	(e) C++ 1.34
(e) Pascal	2.14	2.14	(e) Ada 1.47
(e) Chapel	2.18	2.83	(e) Rust 1.54
(v) Lisp	2.27	3.02	(v) Lisp 1.92
(e) Ocaml	2.40	3.09	(e) Haskell 2.45
(e) Fortran	2.52	3.14	(v) C# 2.57
(e) Swift	2.79	3.40	(e) Swift 2.71
(e) Haskell	3.10	3.55	(i) PHP 2.80
(v) C#	3.14	4.20	(e) Ocaml 2.82
(e) Go	3.23	4.20	(v) C# 2.85
(i) Dart	3.83	6.30	(i) Hack 3.34
(v) F#	4.13	6.52	(v) Racket 3.52
(i) JavaScript	4.45	6.67	(i) Ruby 3.97
(v) Racket	7.91	11.27	(e) Chapel 4.00
(i) TypeScript	21.50	26.99	(v) F# 4.25
(i) Hack	24.02	27.64	(i) JavaScript 4.59
(i) PHP	29.30	36.71	(i) TypeScript 4.69
(v) Erlang	42.23	43.44	(v) Java 6.01
(i) Lua	45.98	46.20	(v) Perl 6.62
(i) JRuby	46.54	59.34	(i) Lua 6.72
(i) Ruby	69.91	65.79	(v) Erlang 7.20
(i) Python	75.88	71.90	(i) Dart 8.64
(i) Perl	79.58	82.91	(i) JRuby 19.84



Energy vs. Time vs. Memory (Pareto Optimization)

Time & Memory	Energy & Time	Energy & Memory	Energy & Time & Memory
C • Pascal • Go	C	C • Pascal	C • Pascal • Go
Rust • C++ • Fortran	Rust	Rust • C++ • Fortran • Go	Rust • C++ • Fortran
Ada	C++	Ada	Ada
Java • Chapel • Lisp • Ocaml	Ada	Java • Chapel • Lisp	Java • Chapel • Lisp • Ocaml
Haskell • C#	Java	Ocaml • Swift • Haskell	Swift • Haskell • C#
Swift • PHP	Pascal • Chapel	C# • PHP	Dart • F# • Racket • Hack • PHP
F# • Racket • Hack • Python	Lisp • Ocaml • Go	Dart • F# • Racket • Hack • Python	JavaScript • Ruby • Python
JavaScript • Ruby	Fortran • Haskell • C#	JavaScript • Ruby	TypeScript • Erlang
Dart • TypeScript • Erlang	Swift	TypeScript	Lua • JRuby • Perl
JRuby • Perl	Dart • F#	Erlang • Lua • Perl	
Lua	JavaScript	JRuby	
	Racket		
	TypeScript • Hack		
	PHP		
	Erlang		
	Lua • JRuby		
	Ruby		

Wrap Up

- General ranking of programming languages
- Unexpected results
- Time and Energy are not always proportional
- C is still the King (Rust is close behind)



Wrap Up

- ▶ General ranking of programming languages
- ▶ Unexpected results
- ▶ Time and Energy are not always proportional
- ▶ C is still the King (Rust is close behind)

Future Work

- ▶ Add new languages
- ▶ Update certain compilers
- ▶ Measure continuous RAM usage
- ▶ Suggestions?



29



If you're coding in any language other than C you're murdering polar bears. And drowning Florida.

Frank Denis @jedist1
Energy Efficiency across Programming Languages
[greenlab.di.uminho.pt/wp-content/upl...](http://greenlab.di.uminho.pt/wp-content/uploads/2017/09/paperSLE.pdf)



Replying to @hyc_symas @ISCdotORG

OTOH, if you're coding in C you're responsible for all the buffer overflows and resulting exploits. Polar bears and Florida, or security?



so if you like birds and trees and all the stuff then better avoid Python :)

<http://greenlab.di.uminho.pt/wp-content/uploads/2017/09/paperSLE.pdf> ...



sites.google.com/view/energy-ef ... is unfair towards python. For fair comparison they should also account for energy consumption needed for coding

30

Can we save energy by refactoring Java programs to use different data structure implementations?



31



Research Questions

- ▶ (RQ1) Can we define an energy consumption quantification of Java data structures and their methods?
- ▶ (RQ2) Can we use such quantification to decrease the energy consumption of software systems?

32

Towards a Ranking of Java data structures

Design

► Simple scenario of:

- Storing
- Retrieving
- Deleting

► *String* values



Towards a Ranking of Java data structures

Design - Data Structures

► *Java Collections Framework (JCF)* library

Sets	Lists	Maps
<u><i>ConcurrentSkipListSet</i></u>	<u><i>ArrayList</i></u>	<u><i>ConcurrentHashMap</i></u>
<u><i>CopyOnWriteArraySet</i></u>	<u><i>AttributeList</i></u>	<u><i>ConcurrentSkipListMap</i></u>
<u><i>HashSet</i></u>	<u><i>CopyOnWriteArrayList</i></u>	<u><i>HashMap</i></u>
<u><i>LinkedHashSet</i></u>	<u><i>LinkedList</i></u>	<u><i>Hashtable</i></u>
<u><i>TreeSet</i></u>	<u><i>RoleList</i></u>	<u><i>IdentityHashMap</i></u>
	<u><i>RoleUnresolvedList</i></u>	<u><i>LinkedHashMap</i></u>
	<u><i>Stack</i></u>	<u><i>Properties</i></u>
	<u><i>Vector</i></u>	<u><i>SimpleBindings</i></u>
		<u><i>TreeMap</i></u>
		<u><i>UIDefaults</i></u>
		<u><i>WeakHashMap</i></u>



Towards a Ranking of Java data structures

Design - Methods

Sets	Lists	Maps
<u><i>add</i></u>	<u><i>add</i></u>	<u><i>clear</i></u>
<u><i>addAll</i></u>	<u><i>addAll</i></u>	<u><i>containsKey</i></u>
<u><i>clear</i></u>	<u><i>add(index)</i></u>	<u><i>containsValue</i></u>
<u><i>contains</i></u>	<u><i>addAll(index)</i></u>	<u><i>entrySet</i></u>
<u><i>containsAll</i></u>	<u><i>clear</i></u>	<u><i>get</i></u>
<u><i>iterateAll</i></u>	<u><i>contains</i></u>	<u><i>iterateAll</i></u>
<u><i>iterator</i></u>	<u><i>containsAll</i></u>	<u><i>keySet</i></u>
<u><i>remove</i></u>	<u><i>get</i></u>	<u><i>put</i></u>
<u><i>removeAll</i></u>	<u><i>indexOf</i></u>	<u><i>putAll</i></u>
<u><i>retainAll</i></u>	<u><i>iterator</i></u>	<u><i>remove</i></u>
<u><i>toArray</i></u>	<u><i>lastIndexOf</i></u>	<u><i>values</i></u>
	<u><i>listIterator</i></u>	
	<u><i>listIterator(index)</i></u>	
	<u><i>remove</i></u>	
	<u><i>removeAll</i></u>	
	<u><i>Continues...</i></u>	
	<u><i>remove(index)</i></u>	
	<u><i>retainAll</i></u>	
	<u><i>set</i></u>	



Towards a Ranking of Java data structures

Design - Benchmark


Test description of Set methods

Method	Description of the test for the method
<u><i>add</i></u>	<u>add popsize/10 elements. half existing, half new</u>
<u><i>addAll</i></u>	<u>addAll of secondaryCol 5 times</u>
<u><i>clear</i></u>	<u>clear 5 times</u>
<u><i>contains</i></u>	<u>contains popsize/10 elements. half existing, half new</u>
<u><i>containsAll</i></u>	<u>containsAll of secondaryCol 5 times</u>
<u><i>iterateAll</i></u>	<u>iterate and consult popsize values</u>
<u><i>iterator</i></u>	<u>iterator popsize times</u>
<u><i>remove</i></u>	<u>remove popsize/10 elements. half existing, half new</u>
<u><i>removeAll</i></u>	<u>removeAll of secondaryCol 5 times</u>
<u><i>retainAll</i></u>	<u>retainAll of secondaryCol 5 times</u>
<u><i>toArray</i></u>	<u>toArray 5 times</u>



Towards a Ranking of Java data structures

Execution - Specifications




- ▶ Linux 3.13.0-74-generic OS
- ▶ 8GB Ram
- ▶ Intel(R) Core(TM) i3-3240 CPU @ 3.40GHz
- ▶ Java Interpreter/Compiler versions 1.8.0_66
- ▶ RAPL/jRAPL



Towards a Ranking of Java data structures

Execution - For every test



- ▶ Warm-up
 - ▶ Instantiated
 - ▶ Populated w/ **popsize**
 - ▶ Performed simple actions on the data structure
- ▶ Each test x20
 - ▶ Extracted time/Joules consumed
 - ▶ Removed lowest/highest 20%
- ▶ 336 different test (Collection.method) configurations
- ▶ 6720 executions for each **popsize**
- ▶ 20,000+ different executions



Towards a Ranking of Java data structures

Results (25k pop)

Methods	Concurrent SkipListSet	HashSet	Linked SkipListSet	TreeSet
add	1.6822	87	1.7749	87
addAll	1.4546	93	1.4771	89
clear	1.4901	78	1.3288	64
contains	1.4333	88	2.6845	78
containsAll	1.8317	96	1.9000	77
iterateAll	1.9225	99	1.4554	92
iterator	1.6296	83	1.9598	75
remove	1.8787	78	1.2633	75
removeAll	1.8072	85	2.1359	77
retainAll	3.2607	206	2.4092	200
toArray	1.4188	85	1.8833	80

Methods	Concurrent HashMap	Concurrent SkipListMap	HashMap	HashMap	Linked HashMap	Properties	Simple LinkedHashMap	TreeMap	UIDefaults	Weak HashMap
clear	2.0276	94	2.2961	88	1.8195	104	1.5783	94	1.5929	97
containsKey	2.3138	105	2.1051	113	2.1883	101	1.8581	94	1.8725	103
containsValue	21.5613	2305	7.4033	643	8.3615	683	8.4957	765	6.3326	462
entrySet	2.2878	93	2.2383	116	1.8533	108	2.1332	107	1.8362	111
get	2.3396	103	1.9972	119	1.8120	102	1.8078	100	1.8252	116
isEmpty	2.1041	95	1.8333	115	2.0673	100	1.9383	91	1.6462	113
keySet	1.7287	95	2.4868	124	1.6413	114	2.2226	99	1.8328	103
put	1.8591	104	2.2888	102	2.4628	92	1.3123	95	2.0338	108
putAll	1.8187	95	2.2852	122	1.7964	100	1.9349	105	1.8068	113
remove	1.8574	91	2.2133	105	1.9156	109	1.6886	91	2.2380	108
values	1.8272	85	2.4662	116	2.2528	109	2.2286	94	2.0009	107



Towards a Ranking of Java data structures

Methods	Concurrent SkipListSet		HashSet		Linked HashSet		TreeSet	
	J	ms	J	ms	J	ms	J	ms
add	14.0472	824	17.5243	1072	15.0643	876	14.1021	758
addAll	19.5092	1518	17.8589	1100	16.5155	983	13.5737	983
clear	11.5958	747	12.0199	764	12.2874	770	11.5565	758
contains	13.6576	870	16.6950	1014	15.6210	880	11.2337	682
containsAll	16.9809	1212	17.2110	1038	15.8865	886	12.3979	844
iterateAll	13.0184	785	18.1706	1091	15.4155	865	11.2088	684
iterator	13.2534	752	16.7433	1013	15.5284	850	11.0499	641
remove	12.7444	789	15.5699	949	13.6615	799	11.2653	675
removeAll	17.2849	1293	17.0514	998	14.5821	841	13.2071	937
retainAll	3621.9872	346898	3912.0129	384829	3584.3529	346337	4111.2397	408297
toArray	14.8120	875	17.8458	1070	14.3511	848	13.1271	750



Towards a Ranking of Java data structures

Methods	ArrayList		AttributeList		CopyOn Write ArrayList		LinkedList		RoleList		Role Unresolved List		Stack		Vector	
	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms
add	0.9773	71	1.1510	67	1.7839	117	1.8016	86	1.4801	76	1.1865	74	1.5659	76	1.5177	69
addAll	1.3353	76	1.0492	88	1.3586	82	1.1043	88	1.6661	76	1.8672	88	1.1015	88	1.7903	73
addAlli	1.7855	86	1.6035	68	1.1789	86	1.7272	99	1.5980	81	1.2497	85	1.2962	72	1.6268	90
addi	1.7125	93	1.3849	87	1.6558	119	1.6404	96	1.2718	85	1.3124	86	1.5287	83	1.4554	86
clear	1.1284	76	1.2409	75	1.7155	68	1.6497	74	1.6705	76	1.4304	80	1.6199	73	1.0574	71
contains	2.7568	166	2.4228	165	3.1768	167	3.1552	193	2.1751	162	2.4688	164	2.0128	166	2.1558	168
containsAll	1.5993	87	1.8053	92	2.1889	92	2.2887	118	1.3244	100	1.3930	96	1.2054	89	1.5091	87
get	2.0029	83	1.1171	78	1.4918	77	2.0168	109	2.2110	81	1.6613	71	1.8956	86	1.4978	73
indexOf	1.4447	76	2.0325	84	1.5682	70	2.6289	101	1.5674	79	1.1944	81	1.8090	81	2.0788	75
iterateAll	2.0701	79	1.0473	77	1.0103	73	2.6401	107	1.3605	85	1.7822	71	1.6036	81	1.1336	87
iterator	1.4893	84	1.1589	84	1.3922	72	1.7666	108	1.9760	73	1.3300	79	2.1895	84	1.6505	83
lastIndexOf	1.7750	99	1.7666	98	2.0383	94	2.5019	127	1.8914	92	1.4211	95	1.2260	84	1.2296	96
listIterator	1.4457	76	1.6190	84	1.3737	71	2.5003	106	1.3380	80	1.5176	85	1.6354	69	1.2746	81
listIteratori	1.7356	78	1.1552	81	1.5160	77	2.1996	105	1.7588	79	1.0334	80	1.8799	85	1.7545	78
remove	1.1308	96	1.4480	85	2.1946	162	1.6924	98	1.4560	84	1.1368	85	1.2663	96	1.4973	82
removeAll	8.0905	671	7.8108	697	7.3237	666	8.3150	752	7.6148	692	7.9911	664	7.3824	654	7.1281	665
removei	1.9135	85	1.3534	92	2.2858	118	1.7174	100	1.6308	85	1.6369	89	1.5850	81	1.5486	90
retainAll	2.7037	193	2.7845	200	2.6052	198	2.5982	205	3.0973	197	2.4172	200	2.7635	242	3.4019	245
set	0.9476	64	1.5943	70	1.9669	110	2.0474	112	1.5249	76	1.2312	73	1.4938	75	1.4957	72
sublist	1.3108	76	1.6021	80	1.4792	80	1.8457	98	1.4910	85	1.5117	71	1.7082	75	0.9414	75
toArray	1.6418	84	1.5024	84	2.0934	73	1.6739	106	1.5418	79	1.7455	83	1.5694	69	2.0213	80



Towards a Ranking of Java data structures

Methods	Concurrent HashMap		Concurrent SkipListMap		HashMap		Hashtable		Linked HashMap		Properties		Simple Bindings		TreeMap		UIDefaults		Weak HashMap	
	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms	J	ms
clear	2.0276	94	2.2961	88	1.8395	104	1.5761	94	1.5025	97	2.0777	98	2.1401	106	1.6706	98	1.8143	105	1.9941	95
containsKey	2.3132	105	2.1693	123	2.1343	103	1.8582	94	1.8726	103	1.6018	107	1.8055	99	1.9452	100	2.3366	89	1.9675	108
containsValue	21.5611	2305	7.8032	643	8.3615	683	8.4957	765	6.1326	462	7.3755	692	7.9912	678	9.1771	847	7.9341	714	6.7072	562
entrySet	2.2878	93	2.2363	116	1.8531	108	2.1332	107	1.8362	113	1.7800	97	2.1557	102	2.1617	115	1.7087	105	1.4666	102
get	2.3106	103	1.9972	119	1.8120	102	1.4071	100	1.8252	116	1.7851	97	1.5359	100	2.2331	115	1.5252	89	1.7185	103
iterateAll	2.1041	96	1.8353	115	2.6673	100	1.5343	91	1.6462	111	1.6362	100	2.0472	116	1.9122	111	1.6574	95	1.7139	106
keySet	1.7287	95	2.4889	124	1.6813	114	2.2226	99	1.8328	103	1.4866	92	2.0630	106	2.1680	110	1.5547	99	1.8749	105
put	1.8591	104	2.2888	102	2.4628	92	1.3123	96	2.0338	108	1.7038	107	2.1646	102	1.4355	91	2.1204	93	2.5784	105
putAll	1.4147	95	2.2852	122	1.7564	100	1.5949	105	1.8608	113	1.3097	95	2.1461	112	1.8914	116	2.3094	87	2.0750	108
remove	1.8574	92	2.2131	105	1.9256	109	1.6067	97	2.2300	106	1.9660	98	2.2178	106	1.8133	101	1.6888	92	2.4103	103
values	1.8279	85	2.4690	116	2.5755	109	2.2266	94	2.0009	107	1.9120	111	2.0692	108	1.4467	105	1.6533	100	2.4628	111

Is faster, Greener?!

Slower

Faster

Methods	Properties		Simple Bindings	
	J	ms	J	ms
clear	2.0777	98	2.1401	106
containsKey	1.6018	107	1.8055	99
containsValue	7.3755	692	7.9912	678
entrySet	1.7800	97	2.1557	102
get	1.7851	97	1.5359	100
iterateAll	1.6362	100	2.0472	116
keySet	1.4866	92	2.0630	106
put	1.7038	107	2.1646	102
putAll	1.3097	95	2.1461	112
remove	1.9660	98	2.2178	106
values	1.9120	111	2.0692	108

Is slower, Greener?!

Research Questions

- (RQ1) Can we define an energy consumption quantification of Java data structures and their methods?
- (RQ2) Can we use such quantification to decrease the energy consumption of software systems?



Optimizing Energy Consumption of Java Programs

Methodology

1. Compute which implementations/methods are used in the programs
2. Look up the appropriate energy tables for the used implementations/methods
3. Choose the most energy efficient alternative



Optimizing Energy Consumption of Java Programs

Applying the methodology - Data acquisition

- ▶ Obtained Java projects of a Journalism support platform
 - ▶ First year OO course
 - ▶ Collaborators, Journalists, Readers, Editors
 - ▶ Write chronicles and reports
 - ▶ Give likes and comments
 - ▶ Etc.
- ▶ Average of
 - ▶ 36 classes
 - ▶ 104 methods
 - ▶ 2000 lines of code



Optimizing Energy Consumption of Java Programs

Applying the methodology - Data acquisition

- ▶ Obtained 7 test cases to simulate usage
 - ▶ Size varied between 2000-10000 for each test case/each entity
- ▶ Chosen popsize – 25,000 (smallest)
- ▶ Applied methodology on 5 projects
 - ▶ Detected usage of any JCF implementation
 - ▶ Detected which methods were used for each implementation
 - ▶ Chose the most efficient implementation for each project
 - ▶ Measured the changes before and after



Optimizing Energy Consumption of Java Programs

Applying the methodology - Example

Projects	Data Structures		Methods					
	Original	Optimized	{containsKey, get, put, values}					
1	TreeMap	Hashtable	Linked		Simple		Weak	
	Concurrent	Concurrent	CopyOn Write		Role	Unresolved	Stack	Vector
Methods	ArrayList	AttributeList	ArrayList	LinkedList	RoleList	List	J	J
add	0.9773	1.1510	1.7839	1.8016	1.4801	1.1865	1.5659	1.5177
listiterator	1.4457	1.6190	1.3737	2.5003	1.3380	1.5176	1.6354	1.2746
Total	2.4229	2.7700	3.1575	4.3018	2.8181	2.7041	3.2012	2.7923



Optimizing Energy Consumption of Java Programs

Applying the methodology - Results

Results of pre and post optimization

Projects	Data Structures				Improvement	
	Original		Optimized		J	ms
1	23.744583	1549	22.7071302	1523	4.37%	1.68%
2	24.6787895	1823	23.525123	1741	4.67%	4.50%
3	25.0243507	1720	22.259355	1508	11.05%	12.33%
4	17.1994425	1258	16.2014997	1217	5.80%	3.26%
5	19.314512	1372	18.3067573	1245	5.22%	9.26%

- ▶ Between 4.37% - 11.05%
- ▶ Average of 6.2%



FW + Conclusion

- ▶ Presented detailed study of the energy consumption of Sets, Lists, and Maps
 - ▶ Quantification of the energy spent by each method (RQ1 Answer)
- ▶ Introduced a very simple methodology to optimize Java programs (RQ2 Answer)
- ▶ Consider other object types (int, objects, etc.)
- ▶ Implement an automatic refactoring plugin
- ▶ Found @ our page: <http://greenlab.di.uminho.pt/>



Can we detect energy hotspots in source code?

Spectrum-based Fault Localization (SFL)

The screenshot shows a code editor with the following C code:

```

1 #include <stdio.h>
2
3 int main(int argc, const char* argv[]) {
4     int factorialTotal = 0;
5
6     for (int i = 1; i < argc; i++) {
7         int value = atoi(argv[i]);
8         if (value % 2 == 0) {
9             int factorial = factorialCalculate(value);
10            factorialTotal += factorial;
11        }
12    }
13
14    printf("The sum of all numbers' factorial is: %d", factorialTotal);
15 }
    
```

The 'Tests' table has columns t1, t2, t3, t4, t5. The 'Energy Oracle' on the right shows a vertical stack of colored circles: a green circle (Low Energy) at the top, a yellow circle, a red circle (High Energy) in the middle, and a green circle at the bottom. An arrow labeled 'Energy Error vector' points to the bottom of the stack.

Problem: Energy cannot be defined as binary values



SPELL

Hit Spectrum

Hit Spectrum = $\begin{pmatrix} \text{Energy} \\ \text{Number} \\ \text{Time} \end{pmatrix}$

where Energy = $\{E_{CPU}, E_{DRAM}, E_{MEM}, E_{DISK}, E_{GPU}, E_{NO-FI}, E_{HD}\}$

Global Error Vector

$global_e(i) = \sum_{j=1}^m global_{e_j}(\lambda_{i,j})$

where $global_{e_j}(\lambda_{i,j}) = ER_{i,j} \times T_{i,j} \times N_{i,j}$

Component	Tests					ϕ	ψ
	t1	t2	t3	t4	t5		
c1	$\begin{pmatrix} 37 \\ 1 \\ 75 \end{pmatrix}$	$\begin{pmatrix} 38 \\ 3 \\ 77 \end{pmatrix}$	$\begin{pmatrix} 36 \\ 1 \\ 73 \end{pmatrix}$	$\begin{pmatrix} 37 \\ 3 \\ 74 \end{pmatrix}$	$\begin{pmatrix} 39 \\ 2 \\ 75 \end{pmatrix}$	0.2314 0.3125 0.3104	0.2466
c2	$\begin{pmatrix} 61 \\ 2 \\ 102 \end{pmatrix}$	$\begin{pmatrix} 50 \\ 1 \\ 103 \end{pmatrix}$	$\begin{pmatrix} 58 \\ 1 \\ 102 \end{pmatrix}$	$\begin{pmatrix} 66 \\ 2 \\ 105 \end{pmatrix}$	$\begin{pmatrix} 54 \\ 3 \\ 100 \end{pmatrix}$	0.3577 0.2815 0.4249	0.4678
c3	$\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 34 \\ 2 \\ 42 \end{pmatrix}$	$\begin{pmatrix} 35 \\ 1 \\ 43 \end{pmatrix}$	$\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 61 \\ 2 \\ 43 \end{pmatrix}$	0.1485 0.2188 0.1203	0.1450
c4	$\begin{pmatrix} 42 \\ 1 \\ 34 \end{pmatrix}$	$\begin{pmatrix} 44 \\ 1 \\ 37 \end{pmatrix}$	$\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 61 \\ 2 \\ 43 \end{pmatrix}$	$\begin{pmatrix} 65 \\ 2 \\ 60 \end{pmatrix}$	0.2623 0.1875 0.1444	0.1406
e	$\begin{pmatrix} 140 \\ 4 \\ 211 \end{pmatrix}$	$\begin{pmatrix} 166 \\ 7 \\ 259 \end{pmatrix}$	$\begin{pmatrix} 129 \\ 3 \\ 218 \end{pmatrix}$	$\begin{pmatrix} 164 \\ 7 \\ 222 \end{pmatrix}$	$\begin{pmatrix} 209 \\ 11 \\ 295 \end{pmatrix}$		
global _e	5659.98	6260.08	3416.66	9288.8	14310.6		

Component Similarity

$\phi_j = (\alpha_1(A(j), e), \alpha_2(A(j), e), \alpha_3(A(j), e))$

where $\alpha_x(A(j), e) = \frac{\sum_{i=1}^n \min(A(j,x), e(e)_i)}{\sum_{i=1}^n \max(A(j,x), e(e)_i)}$

Global Similarity

$\psi_j = \alpha(global_e, global_{e_j})$


where $\alpha(c, e) = \frac{\sum_{i=1}^n \min(c_i, e_i)}{\sum_{i=1}^n \max(c_i, e_i)}$

Error Vector

$e = \sum_{j=1}^m \lambda_{1,j} \sum_{i=1}^m \lambda_{2,j} \dots \sum_{i=1}^m \lambda_{m,j} T_i^j$

where $\sum_{j=1}^m \lambda_{i,j} = (\sum_{j=1}^m ER_{i,j} \sum_{i=1}^m N_{i,j} \sum_{j=1}^m T_{i,j})$

SPELL - initial studies



- ▶ Developers found the information to be very useful
 - ▶ Spent on average 43% less time
 - ▶ Produced more energy efficient programs (26% less energy on average)
- ▶ SPELL is:
 - ▶ Language independent
 - ▶ Multi level analysis
 - ▶ Multi hardware component analysis
 - ▶ Points to probable hot spots in source code

