

Benchmark::Perl::Formance for the masses

Steffen "renormalist" Schwigon

Dresden Perl Mongers

16 August 2011

Abstract

Motivation

Measuring

Workloads

Perl::Formance

Challenges

Infrastructure

Summary

Abstract

- **Abstract**

Abstract

- Perl lacks benchmarking
- I work on a full-coverage solution
 - benchmark implementation
 - complete "build - execute - evaluate" lifecycle
 - benchmark result database
 - graphically present numbers and details
- Present the vision and state

Motivation

- Motivation

Perl standard benchmark suite?

The screenshot shows a Google search results page. The search query "Perl standard benchmark suite" has been entered into the search bar, resulting in 0 results found in 0.19 seconds. The sidebar on the left includes links for Alles, Bilder, Videos, News, Shopping, and Mehr. It also features sections for "In der Nähe von ... suchen" with a Standort button and Festlegen button, and "Das Web" with options for Seiten auf Deutsch, Seiten aus Deutschland, and Übersetzte Seiten. The bottom of the page shows a "Zeitleiste" section and a set of navigation icons at the bottom right.

Web Bilder Videos Maps News Shopping Google Mail Mehr +

Anmelden

Google

Perl standard benchmark suite

Ungefähr 0 Ergebnisse (0,19 Sekunden)

Erweiterte Suche

Alles 0 results

Bilder

Videos

News

Shopping

Mehr

In der Nähe von
... suchen

Standort

Festlegen

Das Web

Seiten auf Deutsch

Seiten aus Deutschland

Übersetzte Seiten

Alle Ergebnisse

Zeitleiste

Perl standard benchmark suite?

Web Bilder Videos Maps News Shopping Google Mail Mehr + Anmelden

Google standard benchmark suite

Ungefähr 0 Ergebnisse (0,19 Sekunden)

Erweiterte Suche

Alles 0 results

Bilder

Videos

News

Shopping

Mehr

In der Nähe von ... suchen

Standort Festlegen

Das Web

Seiten auf Deutsch

Seiten aus Deutschland

Übersetzte Seiten

Alle Ergebnisse

Zeitleiste

Perl standard benchmark suite?

Web Bilder Videos Maps News Shopping Google Mail Mehr + Anmelden 

Google Perl standard benchmark suite
Ungefähr 0 Ergebnisse (0,19 Sekunden) Erweiterte Suche

Alles 0 results

Bilder

Videos

News

Shopping

Mehr

In der Nähe von ... suchen Standort Festlegen

Das Web
Seiten auf Deutsch
Seiten aus Deutschland
Übersetzte Seiten

Alle Ergebnisse
Zeitleiste

Google is really sorry about Perl missing a benchmark solution.

But, dear user with

IP 145.12.3.123, geo location 12°65'43",
cookie ID 0xAFDEADBEEF, MAC address 1c:a4:56:f8:90

could you please stop entering the same search again
and again and please, please, please JFDI by yourself?

Thanks.



Navigation icons: back, forward, search, etc.

Measuring

- **Measuring**

The most important benchmark question

- **What do you want to measure?**
- Primary goal: perl5, the interpreter
 - real world workloads
 - micro benchmarks
 - usethreads vs. no usethreads
 - use64bitall vs. no use64bitall
- Secondary goal: misc Perl5 topics
 - subs vs. methods
 - Moose
 - regex engines

Workloads

- Workloads

Where is Perl?

- Everywhere and nowhere
- Benchmarks are a corporate topic
- Perl's usual corporate visibility problem

"The One Workload"

- No dedicated single heavy applications
- Except SpamAssassin
 - important app
 - heavy for a reason
 - hit by 5.12 deprecations
 - handholding on new CPAN release

Self-referential - Perl for Perl

- Refer to Perl itself
 - CPAN.pm
 - Perl::Critic
 - Perl6 STD.pm (gimme5, viv)

Nothing else!

- Not much more

Writing my own benchmark suite

- Goals

- Have benchmark workloads
 - "server" workloads (MEM! + TIME!)
 - "desktop" workloads (small + fast)
 - "micro" benchmarks (language features)
 - "multi purpose" workloads (average, not fitting in others)
- Provide general benchmarking umbrella

Perl::Formance

- **Perl::Formance**

The Perl::Performance benchmark suite

- Collect existing heavy code
- Collect existing benchmark snippets
- Write new benchmark code from scratch
- Measure execution time ("smaller is better")
- <http://xrl.us/cpanperformance>

Perl::Performance features (1)

- Heavy workloads
 - The one: SpamAssassin
 - Type: "server"
 - Example corpus from spamassassin.org
 - Run sa-learn

Perl::Formance features (2)

- Wrap existing benchmarks
 - "Language Shootout" on alioth.debian.org
 - Type: "multi purpose"
 - `binarytrees`
 - `fasta`
 - `nbody`
 - `pidigits`
 - `regexdna`
 - `revcomp`
 - `spectralnorm`
 - `fannkuch (!)`
 - `knucleotide (!)`
 - `mandelbrot (!)`
 - Mostly copy'n'paste into Perl::Formance plugins

Perl::Performance features (3)

- Known heavy code from CPAN
 - Heavy test suites
 - Type: "desktop"
 - Moose test suite
 - Regexp::Common
 - Perl6 STD.pm
 - Type: "desktop"
 - gimme5 STD.pm6
 - viv STD.pm6
 - Data::DPath

Perl::Performance features (4)

- Stress of language features
 - Stress recursion (use Fibonacci numbers as vehicle)
 - Type: "micro"
 - subs - plain Perl
 - methods - plain Perl
 - methods - Moose
 - methods - MooseX::Declare

Perl::Performance features (5)

- Pathological Regular Expressions

- Type: "micro"
- Known pathological issues
- `my $re = ("a?" x $n) . ("a" x $n);`
- See

<http://swtch.com/~rsc/regexp/regexp1.html>

Perl::Formance features (6)

- Compare different regex engines
 - Pluggable regex engines since Perl 5.10
 - Again pathological regular expressions
 - Type: "micro"
- Regex engines
 - native
 - POSIX::Regex
 - **re::engine::Plan9**
 - **re::engine::PCRE**
 - **re::engine::Lua**
 - **re::engine::LPEG** (confusion bonus)
 - **re::engine::Oniguruma**

Different output styles (1)

- Human readable

```
$ benchmark-perlperformance --fastmode [...]  
Rx.regexes.fieldsplit1 : 1.267907  
Rx.regexes.fieldsplit2 : 2.106220  
Rx.regexes.pathological : 1.000129  
Shootout.binarytrees    : 1.046751  
Shootout.fasta          : 2.270553  
Shootout.nbody          : 1.685537  
Shootout.spectralnorm   : 1.855935
```

Different output styles (2)

- Evaluation friendly (`-outstyle=yaml`)

```
---
```

```
perlformance:
  config:
    fastmode: 1
  results:
    Shootout:
      fasta:
        Benchmark:
          - 0.263123989105225
          - 0.25
          - 0
          - 0
          - 1
      count: 1
      goal: 5000
```

Output style

- Augment YAML with surrounding TAP

```
--tapdescription="some description"
```

- Additional Codespeed data structure

```
--codespeed
```

- Why TAP?

- TAP is my hammer
- Later embed into Tapper infrastructure

Plugin API

- Very lax, just namespace + sub main()

```
Benchmark::Perl::Formance::Plugin::Skeleton
Benchmark::Perl::Formance::Plugin::*  
sub main($options) {  
    ...  
    return { Benchmark => $t,  
            ...  
        };  
}
```

Challenges

- Challenges

Automatic building

Automatic building

- sigh

Automatic building



Automatic building

- Working throughout Perl's git history (5.8..blead)
- Cherry-pick fixes from the future (5.8.x)
- Inconsistent version tags
- perlbrew?
 - No git
 - No bootstrapping CPAN, distroprefs (ANDK++), etc.

The Art of Benchmarking (1)

- Stable benchmark environment
 - Perl <-> OS interaction
 - Flush Caches
 - Address Space Layout Randomisation (ASLR)
 - Disable Core Performance Boost
 - Provide stable set of CPAN dependencies
 - Own CPAN mirror
 - Only sync once in a time
 - Stable execution environment
 - Machine, Memory, Harddisk
 - Operating System
 - Compile toolchain
 - Lazily load benchmark plugins+dependencies after fork

The Art of Benchmarking (2)

- I/O everywhere
 - depends on the workload
 - SPECcpu 2006 (Perl 5.8.7) avoids I/O
 - I do **not** fight that battle

Infrastructure

- CPAN module obviously not enough
- I should solve all other challenges, too
- Set up complete infrastructure

Infrastructure

- **Infrastructure**

Infrastructure

- perl64.org - **Benchmark** machine
 - Rent dedicated machine
 - 6-core AMD Opteron 4180
 - Without any running services
 - No private data, for relaxed access
- performance.net - "**Tapper**" application
- speed.performance.net - "**Codespeed**" application

Single Run

- Builds Perl
 - from git
 - inject CPAN toolchain
 - Perl 5.8 .. blead
- Installs Benchmark::Perl::Performance
- Installs Tapper::TestSuite::Benchmark::Perl::Performance
- Runs tapper-testsuite-benchmark-perl-performance
 - runs benchmark-performance
 - --outstyle=yaml
 - --tapdescription="performance"
 - --codespeed
 - adds meta information for Tapper
 - sends report to Tapper server

Multiple runs

- Tapper

Tapper

- Test infrastructure, open sourced by AMD in 2011
 - Overview:
 - <http://developer.amd.com/zones/opensource/AMDTapper>
 - Source:
 - <http://github.com/amd>
 - Productized:
 - Starterkit/Deployment, Docs, Wiki

Tapper

- Why Tapper?

- Tapper == TAP database + automation + scheduler
- Query API for detailed data forensics
 - see my YAPC::EU 2009 presentation
 - "Cinderella 'TAP - The lazy evaluation sisters of TAP::Parser"
 - <http://xrl.us/lazytap> (PDF)

Tapper

- Why Tapper?
 - Schedule / intermix different use-cases
 - auto-rerun for blead/threads
 - auto-rerun for blead/notthreads
 - specific commits
 - completely different stuff I add later (Perl6 benchmarks?, smoke tests?)
 - Have bandwidth ratio for each,
 - eg. more non-threaded than threaded
 - Advanced automation support, timeout handling, etc.
 - Frontends to start single runs
 - cmdline + webGUI
 - with requested options (-commitid, -threads, -64bit)
 - Pass-through data chunks to Codespeed application

Codespeed application

- <http://github.com/tobami/codespeed>
- Web application
- Render benchmark graphs
- Show commit/meta information
- Comparison graphs, baselines, etc.
- Can handle git repos

Own CPAN mirror

- <http://perlformance.net/CPAN>
- To guarantee no unexpected changes
 - slower/faster/break

Summary

- Summary

Summary

- I have benchmarks
- I have automatic Perl + CPAN setup
- I have applications for complete infrastructure
- I have dedicated benchmarking server
- I have dedicated CPAN mirror
- Polished in many details
- **TODO :**
 - 1. Re-setup after server went mental
 - 2. Allow users to request benchmark runs

Bonus Screenshots

- **Bonus Screenshots**

Tapper at <http://performance.net>

tapper

Start Testruns Reports Testplans Metareports Manual

72: benchmark-performance-fast

[report id: 72] [2011-08-12 22:10:53 GMT] [Host: ssd-work]

Test Execution Context

Test results

PASSED

Test file	Test results	%
metainfo		100.0%
results		100.0%

2 files 2 tests, 2 ok, 0 failed, 0 todo, 0 skipped, 0 parse errors
exit status: 0, wait status: 0
elapsed time: 0 wallclock secs (0.10 usr 0.00 sys + 0.00 cusr 0.01 csys = 0.11 CPU)

[raw TAP report](#)

reports by date

- today
- 2 days
- 1 week
- 2 weeks
- 3 weeks
- 1 month
- 2 months
- 4 months
- 6 months
- 12 months

reports by suite

reports by host

Copyright © 2008-2011 AMD Operating System Research Center.

Navigation icons: back, forward, search, etc.

Tapper - Overview lists

Mon Aug 23, 2010

ID	DateTime (GMT)	Suite	Machine	Success	Ratio	Grouped by	Owner
r124555	2010-08-23 01:49	benchmark-performance	muabrek	FAIL			
r124720	2010-08-23 11:23	benchmark-performance	muabrek	PASS			
r124661	2010-08-23 08:02	benchmark-performance	muabrek	PASS			
r124583	2010-08-23 04:47	benchmark-performance	muabrek	PASS			

Sun Aug 22, 2010

ID	DateTime (GMT)	Suite	Machine	Success	Ratio	Grouped by	Owner
r124305	2010-08-22 13:09	benchmark-performance	muabrek	FAIL			
r124486	2010-08-22 22:43	benchmark-performance	muabrek	PASS			
r124412	2010-08-22 19:22	benchmark-performance	muabrek	PASS			
r124371	2010-08-22 16:06	benchmark-performance	muabrek	PASS			
r124242	2010-08-22 10:13	benchmark-performance	muabrek	PASS		muebarek-3-d797762b03e04b842f6f7cd9c5f1fb4	

Sat Aug 21, 2010

ID	DateTime (GMT)	Suite	Machine	Success	Ratio	Grouped by	Owner
r123612	2010-08-21 00:57	benchmark-performance	muabrek	PASS			
r124011	2010-08-21 22:44	benchmark-performance	muabrek	FAIL			
r123985	2010-08-21 19:46	benchmark-performance	muabrek	FAIL			
r123895	2010-08-21 16:44	benchmark-performance	muabrek	PASS			
r123844	2010-08-21 13:37	benchmark-performance	muabrek	PASS			
r123802	2010-08-21 10:29	benchmark-performance	muabrek	PASS			
r123745	2010-08-21 07:15	benchmark-performance	muabrek	PASS			
r123683	2010-08-21 04:09	benchmark-performance	muabrek	FAIL			

Fri Aug 20, 2010

ID	DateTime (GMT)	Suite	Machine	Success	Ratio	Grouped by	Owner
r123200	2010-08-20 02:35	benchmark-performance	muabrek	PASS			
r123559	2010-08-20 21:34	benchmark-performance	muabrek	PASS			

Tapper - Metainfo

tapper

Start Testruns Reports Testplans Metareports Manual

271038: Perl-Performance-Dummy

(Report ID: 271038 | 2011-08-12 09:40:04 GMT | 2 cores [AMD Athlon(tm) 64 X2 Dual Core Processor 6000+] | Ubuntu 10.10 | Host: perl64.org)

Reportgroup ()

ID	Date/Time (GMT)	Suite	Machine	Success	Ratio	Grouped by	Owner
r271038	2011-08-12 09:45	Perl-Performance-Dummy	perl64.org	PASS	100%	cdd8e4c48e7ab7636bc15e5ba4b36e7f1	
r271042	2011-08-12 10:01	Perl-Performance-Dummy	perl64.org	PASS	100%		

Test Execution Context

Perl-Performance-Dummy (r271038)

section 000

ram: 2007
cpuinfo: 2 cores [AMD Athlon(tm) 64 X2 Dual Core Processor 6000+]
uname: Linux bascha 2.6.35-30-generic #56-Ubuntu SMP Mon Jul 11 20:01:08 UTC 2011 x86_64 GNU/Linux
osname: Ubuntu 10.10
flags: root=UUID=6990cb5e-1a77-40b8-ba05-919f6c928607 ro quiet splash
kern: 2.6.35-30-generic
changeset: 98765
ticket_url: <http://speed.performance.net/changes/?rev=98765>

reports by date

today
2 days
1 week
2 weeks
3 weeks
1 month
2 months
4 months
6 months
12 months

Test results

PASSED

Test file	Test results	%
section-000		100.0%
1 files	4 tests, 4 ok, 0 failed, 0 todo, 0 skipped, 0 parse errors exit status: 0, wait status: 0 elapsed time: 0 wallclock secs (0.04 usr + 0.00 sys = 0.04 CPU)	100.0%

reports by suite

reports by host

Tapper - Metainfo

Perf-Performance-Dummy ([#271038](#))

[reports by host](#)

section-000
ram: 2007
cpuinfo: 2 cores [AMD Athlon(tm) 64 X2 Dual Core Processor 6000+]
uname: Linux bascha 2.6.35-30-generic #56-Ubuntu SMP Mon Jul 11 20:01:08 UTC 2011 x86_64 GNU/Linux
osname: Ubuntu 10.10
flags: root=UUID=6990cb5e-1a77-40b8-ba05-919fc928607 ro quiet splash
kernel: 2.6.35-30-generic
changeset: 98765
ticket_url: <http://speed.performance.net/changes/?rev=98765>

Test results

PASSED

Test file	Test results	%
section-000		100.0%
1 files	4 tests, 4 ok, 0 failed, 0 todo, 0 skipped, 0 parse errors exit status: 0, wait status: 0 elapsed time: 0 wallclock secs (0.04 usr + 0.00 sys = 0.04 CPU)	100.0%

[raw TAP report](#)

Attachments

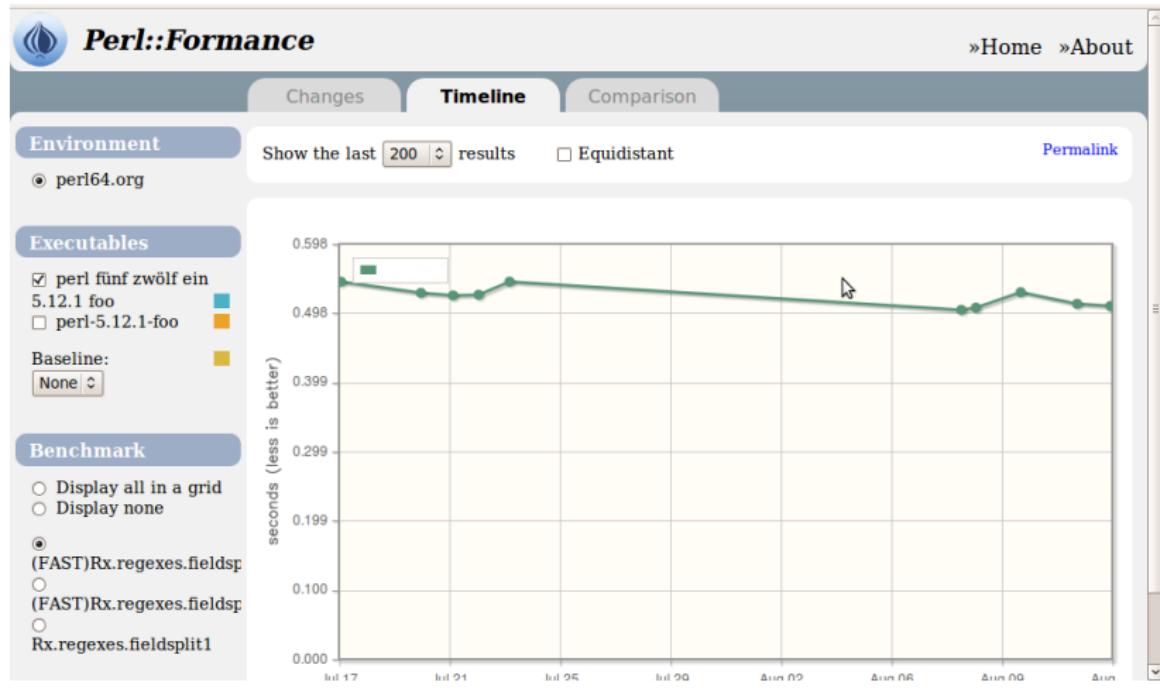
.perfomance-dummy.sh	287 Bytes	view inline ansi-colored	plain	2011-08-12 09:45:18 GMT
/proc/cpuinfo	1466 Bytes	view inline ansi-colored	plain	2011-08-12 09:45:25 GMT
/proc/devices	481 Bytes	view inline ansi-colored	plain	2011-08-12 09:45:32 GMT
/proc/version	146 Bytes	view inline ansi-colored	plain	2011-08-12 09:45:39 GMT

« [perfomance-dummy](#) [perfomance-dummy](#) »

Tapper - Tracked data

```
- 0
- 0
- 0.03
- 5
goal: STD.pm6
viv:
Benchmark:
- 0.0236821174621582
- 0
- 0
- 0
- 0
- 5
goal: STD.pm6
PerlCritic:
PLUGIN_VERSION: 0.001
bundled:
Benchmark:
- 29.1310901641846
- 0
- 0
- 28.96
- 0.16
- 2
count: 2
perl_critic_version: 1.108
upstream:
Benchmark:
- 30.2070469856262
- 0
- 0
- 29.71
- 0.13
- 2
count: 2
perl_critic_version: 1.108
```

Codespeed at <http://speed.performance.net>



Thanks.

- **Thanks.**