

Performance Report July 14 2020

Tester: me@rochus-keller.ch

Subject: **St80LjVirtualMachine 0.5.6 and 0.6.0** on Linux i386 on Intel Core Duo L9400 1.86GHz with 4GB RAM

Using benchmark expression <https://github.com/rochus-keller/Smalltalk/blob/master/benchmark/Benchmark.st>

Using LuaJIT 2.0.x commits 7eb96843ff9d4bed and d518511e4ffe742

Categorized by behaviour, ordered by number of observations, all times in milliseconds

One fast, then slow:

<i>run</i>	<i>finished</i>	<i>duration</i>
1	33'735	11'010
2	104'873	23'195
1	259'477	12'498
2	319'056	21'260
1	19'580	9'631
2	60'609	20'207
3	137'235	20'143
1	18'482	9'594
2	53'065	24'229
3	123'030	24'184
1	129'433	12'762
2	180'316	22'184
1	62'834	10'809
2	100'257	19'538
3	171'045	19'505
1	34'651	13'387
2	68'949	19'654
3	198'654	19'593

Two or more fast, then slow:

<i>run</i>	<i>finished</i>	<i>duration</i>
1	33'098	13'385
2	92'817	13'175
3	141'666	15'145
4	205'490	21'463
5	386'153	21'454
6	452'859	22'699
1	23'913	11'055
2	70'502	10'862
3	130'342	19'981
4	207'344	20'380
5	782'182	20'217
1	11'011	10'657
2	57'230	11'396
3	100'356	19'432
1	9'367	10'125
2	39'914	12'267
3	71'415	20'382
4	119'398	20'364

All same medium speed:

<i>run</i>	<i>finished</i>	<i>duration</i>
1	14'410	14'677
2	52'810	14'321
3	87'134	14'576
4	119'521	14'443
5	163'304	14'441
6	197'662	14'492
7	250'612	14'324
8	299'042	14'278
9	419'603	14'331
10	532'863	14'267

All same high speed:

(this was the first set after new Benchmark expression)

<i>run</i>	<i>finished</i>	<i>duration</i>
1	?	10'676
2	?	10'587
3	?	10'615

(changed the expression to show finished)

4	329'540	10'669
5	436'904	10'622

(wasn't aware this is exceptional so I stopped here)

The geomean of the equivalent C++ app is 5'137

One medium, then slow:

<i>run</i>	<i>finished</i>	<i>duration</i>
1	15'413	16'072
2	63'183	23'606

Geomean fast state:
11'231

Geomean slow state:
21'071

Min duration:
9'594

Max duration:
24'229

Speed-down factor vs. C++
1.87

Speed-down factor vs. C++
4.72

Note: St80LjVirtualMachine is restarted for each new sequence of runs; a sequence of runs happens in the same LuaJIT session; remember that the actual running Lua program is the Smalltalk interpreter; the benchmark runs on top of the interpreter; the interpreter keeps running between benchmark runs; the *finished* column is the time since its start.

Conclusions:

- LuaJIT proved to be able to run the Benchmark in 9.6 sec.
- Sequences of similar durations show that performance is stable over a period of time and then changes abruptly to another (worse) stable level.
- The difference between the fast and slow sequences tends to be around 10 secs.
- This 10 secs difference was already observed in earlier, less efficient versions of the VM, when fast was 42 secs and slow 52 secs in geomean.
- There is no obvious reason why LuaJIT decides to drop fast traces in favour of slow ones.