СЕКЦИОННЫЕ ДОКЛАДЫ

An automated system for program parameters fine tuning in the cloud

S. A. Smirnov[a], A. S. Tarasov

Institute for Information Transmission Problems of the Russian Academy of Science, Kharkevich Institute,
Bolshoy Karetny per. 19, build.1, Moscow 127051 Russia

E-mail: [a] sasmir@gmail.com

The paper presents a software system aimed at finding best (in some sense) parameters of an algorithm. The system handles both discrete and continuous parameters and employs massive parallelism offered by public clouds. The paper presents an overview of the system, a method to measure algorithm's performance in the cloud and numerical results of system's use on several problem sets.

Keywords: algorithmic parameter optimization, parameter tuning, cloud computing

Автоматическая облачная система подстройки параметров алгоритмов

С. А. Смирнов, А. С. Тарасов

*Институт проблем передачи информации им. А. А. Харкевича Российской академии наук, Россия,
127051, г. Москва, Большой Каретный переулок, д. 19, стр. 1*

В работе представлена система, обеспечивающая подбор наилучших в смысле времени выполнения настроек алгоритма. В качестве алгоритма был взят пакет решения задач частично-целочисленного линейного и нелинейного программирования SCIP. Возможность параллельного перебора множества вариантов настроек обеспечивается кластером из виртуальных машин, автоматически создаваемых в облаке. Представлены результаты работы системы на нескольких наборах задач.

Ключевые слова: оптимизация параметров алгоритмов, облачные вычисления

Introduction

Growing number of Infrastructure as a service (IaaS) providers we observe today is a direct effect of computation costs getting cheaper and of infrastructure automatization levels getting higher. Cloud services make it possible to automate more programmer's work making him more productive. It may be considered as another step in continuous process of adding more abstraction levels to a computer system: high-level programming languages, interactive debugging, automatic build systems, etc. For developers it allows for rapidly creating development and test sandboxes, quickly provisioning virtual machines with needed software, testing load and scalability.

There are lots of problems that can be automated by the use of clouds. One of such problems is fine tuning an algorithm to make it work better in some sense, for example, faster. It can be done in many ways: modifying hard coded parameters inside a program, smart analysis of the program's source code, adjusting parameters inside configuration files of program's modules. In our study we have chosen the last described way: fine tuning configuration parameters for the SCIP (Solving Constraint Integer Programs) solver. SCIP [Achterberg Tobias, 2009] is currently one of the fastest non-commercial solvers for mixed integer programming (MIP) and mixed integer nonlinear programming (MINLP). It is also a framework for constraint integer programming and branch-cut-and-price. It allows for total control of the solution process and the access of detailed information down to the guts of the solver. Although SCIP is a very fast solver even with default parameters, it should be possible to fine tune the parameters for one's work. It is quite simple if there are a couple of parameters and not many test problems. However SCIP is very configurable having more than a thousand parameters. Apparently having such a large set of configuration parameters makes fine tuning it quite time consuming. That is why we tried to automate this process making a system choosing best configuration setting for a set of problem instances. Due to vast number of SCIP runs needed we had to use a cloud to make the process quick.

In this paper we use the following terminology. A program has configuration *parameters* controlling how the it works. Every parameter has a *value* assigned. A set of parameter values is called the *settings*. When we run a solver it is given settings in form of a configuration file and a problem instance. The goal of our system may be thought as finding the settings yielding the shortest running time on a set of problem instances.

Other works on general optimization of algorithmic parameters include Selection Tool for Optimization Parameters (STOP) [Baz et al., 2007] based on intelligent sampling of settings throughout the space and OPAL framework [Audet, Dang, Orban, 2012] based on mesh adaptive direct search. The former tool works with a small set of parameters having discrete values. Our aproach allows working with large numbers of parameters and thair values.

Implementation

Let us begin with a brief overview of the system. One begins using it by specifying the number of computing hosts in the Vagrant's [Vagrant] configuration file and then starting the system by running init-virtualbox.sh or init-digitalocean.sh script. After a while one has a cluster of a master host and the specified number of slave hosts where Simple Linux Utility for Resource Management (SLURM) [Yoo, Morris, and Mark, 2003] and other essential software are installed and running. Then one connects to the master host by issuing *vagrant ssh* command where one can manage the system with *optctl.py* command.

Currently the settings optimization process consists of three phases: time check, big step and inter step. During the first phase every problem's instance is evaluated once on each computing node with default settings. The main goal of this step is to get an estimate of maximum time allowed for a problem instance to run until it's killed by SLURM. The big step phase is the most computationally intensive one. On this step huge number of settings with only one parameter different from defaults is evaluated. As a result the big step allows us to sort parameter values based on their impact on solving time.

Next, on the inter step phase, four best parameter values from the head of the big step's sorted list are chosen and all their possible combinations are evaluated. After this step we have the best settings in terms of running time. This step is not very time consuming and can be repeated multiple times.

Measuring running time in the cloud

Measuring programs's running time in the cloud reliably is not very simple. Naive approach like wall-clock time or processor cycles are not reliable due to computer resources overcommit by a cloud provider. Depending on the load other virtual machines express on the hypervisor host, program's running time can change dramatically. There is a better approach: one can measure the number of instructions executed by the CPU while running the program. Of course, different instructions may need different numbers of cycles to complete so it may be hard to correlate running time to the number of instructions executed. Instruction count becomes handy when comparing performance the same program expresses with the same input but with different settings.

In x86 CPUs instructions can be counted in hardware by the Performance Monitoring Unit (PMU). One can use PAPI or perf tool to set up and access the hardware counters. Not every hypervisor supports PMU virtualization, e.g. VirtualBox does not. However modern KVM releases has such support.

In our system we used perf tool to measure user space instruction count which gives very stable results independent of the hypervisor host's load.

Here is a sample run of SCIP under perf-stat. Six runs of SCIP were made, average counter values and their standard deviations can be observed:

```
$ perf stat -r 6 -e cpu-clock,task-clock,cycles,instructions,instructions:u,\
instructions:k scipampl TSP_Uniform_50_10.nl
 Performance counter stats for 'scipampl TSP_Uniform_50_10.nl' (6 runs):

    81154.175629 cpu-clock                                        ( +-  2.32% )
    81154.063870 task-clock                # 0.999 CPUs utilized  ( +-  2.32% )
 175,626,392,898 cycles                    # 2.164 GHz            ( +-  0.13% )
 267,235,503,611 instructions              # 1.52  insns per cycle ( +-  0.00% )
 265,101,243,265 instructions:u            # 1.51  insns per cycle ( +-  0.00% )
   2,134,260,346 instructions:k            # 0.01  insns per cycle ( +-  0.10% )


    81.224668399 seconds time elapsed                            ( +-  2.32% )
```

Instaructions:u counter gives much more stable results than software counters or cycles counted in hardware.

Same single CPU virtual machine with two SCIP instances running simultaneously:

```
 Performance counter stats for 'scipampl TSP_Uniform_50_10.nl':

    82580.457064 cpu-clock
    82579.334255 task-clock                # 0.493 CPUs utilized
 181,566,274,355 cycles                    # 2.199 GHz
 267,300,821,128 instructions              # 1.47  insns per cycle
 265,099,385,783 instructions:u            # 1.46  insns per cycle
   2,201,435,345 instructions:k            # 0.01  insns per cycle


   167.578326122 seconds time elapsed


 Performance counter stats for 'scipampl TSP_Uniform_50_10.nl':

    82581.195083 cpu-clock
    82580.104484 task-clock                # 0.493 CPUs utilized
 181,589,302,031 cycles                    # 2.199 GHz
 267,299,923,846 instructions              # 1.47  insns per cycle
 265,099,381,995 instructions:u            # 1.46  insns per cycle
   2,200,541,851 instructions:k            # 0.01  insns per cycle


   167.589704033 seconds time elapsed
```
Again, instructions:u are much more accurate.

As we can see from this examples, instructions:u is the most stable event counter at least with SCIP. It even allows for running multiple solver instances simultaneously with acceptable timing accuracy.

Numerical results

We have performed testing with two different sets of problem instances. One of the sets was tested with two version of SCIP: 3.0.2 and 3.1.0. Throughout the tests, 48 computing nodes with identical virtual machines were used.

First problem set consisted of ten randomly generated traveling salesman problem instances of the same size. SCIP 3.1.0 was used. Big step for this set consisted of 28810 jobs and took six hours and a half to complete while total CPU time consumed was 296 hours, as if 46 machines were used. After one inter step optimal settings were obtained. Second inter step showed no improvement. If we compare the sums of running times for default settings and for optimized ones, we oserve 3x speedup with the latter (see Table 1). Optimized settings consisted of only one parameter value:
lp/scaling = FALSE

Table 1. Traveling salesman problem, learning data set

Problem instance	Default, sec.	Optimized, sec.
TSP_Uniform_50_1	3,0	2,5
TSP_Uniform_50_2	10,4	6,4
TSP_Uniform_50_3	29,4	16,1
TSP_Uniform_50_4	6,8	8,7
TSP_Uniform_50_5	39,9	36,6
TSP_Uniform_50_6	37,9	7,6
TSP_Uniform_50_7	105,8	31,3
TSP_Uniform_50_8	102,8	13,6
TSP_Uniform_50_9	31,4	10,9
TSP_Uniform_50_10	80,7	13,2

For testing purposes more TSP instances were generated and run with the same optimized parameters (see Table 2), the speedup is just 1,41x here.

Table 2. Traveling salesman problem, control data set

Problem instance	Default, sec.	Optimized, sec.
TSP_Uniform_50_11	52,6	27,3
TSP_Uniform_50_12	11,5	9,2
TSP_Uniform_50_13	1,4	2,2
TSP_Uniform_50_14	7,0	5,2
TSP_Uniform_50_15	270,8	230,6
TSP_Uniform_50_16	64,0	10,5
TSP_Uniform_50_17	4,0	3,7
TSP_Uniform_50_18	27,1	14,4
TSP_Uniform_50_19	5,0	12,0

Second problem set consisted of five instances which solved quickly with SCIP 3.0.2 and very slowly with SCIP 3.1.0.

An attempt was made to find parameters making SCIP 3.1.0 working on the problem as good as 3.0.2. We took all parameters that changed their default values, were renamed or added in 3.1.0, which resulted in 186 parameters (against 1547 total parameters). One of the instances

(w6_t19_test_8) was dropped after the time check phase due to hitting memory limit (512 MB RAM in VM). Big step consisted of 1260 jobs for the first four instances and took five hours and a half to complete. Total CPU time spent in SCIP was 237 hours which equals to 43 hosts working. After one inter step optimized settings were obtained. Second inter step showed no improvement. If we compare the sums of running times for default settings and for optimized ones, we oserve 1,65x speedup when the latter is used. It should be noted that optimized settings also improved time for the problem w6_t19_test_8 that was not involved in the tests due to memory limitation. As we can see, our system was not able to find settings making SCIP 3.1.0 perform as good as SCIP 3.0.2 for this problem, how-ever a noticeble speedup was obtained. Optimized settings:

```
heuristics/rins/minnodes = 25
lp/checkdualfeas = FALSE
lp/disablecutoff = 1
```

Table 3. Load balancing problem

Problem	3.1.0, def., sec.	3.1.0, opt., sec.	3.0.2, def., sec.	3.0.2, opt., sec.
w6_t15_test_4	5,95	3,11	1,97	0,92
w6_t18_test_4	586,18	186,8	70,3	32,5
w6_t19_test_4	1420,4	823,9	223,9	178,3
w6_t19_test_5	941,7	737,6	138,7	119,3
w6_t19_test_8	11596,3	7077,7	382,6	319,7

We also tried optimizing settings for this problem set in SCIP 3.0.2 on all its parameters. During big step 13200 jobs were run in 13 hours and a half, 594 hours were spent in the solver as if 44 hosts were working. After two inter steps optimized settings were obtained, third interstep yielded no im-provement. Here 1,26x speedup was obtained. Optimized settings after first inter step:

```
constraints/linear/upgrade/setppc = FALSE
lp/solvefreq = 0
```

After the second inter step:

```
constraints/linear/upgrade/setppc = FALSE
lp/solvefreq = 0
conflict/preferbinary = TRUE
heuristics/fracdiving/freqofs = 1
heuristics/veclendiving/freq = -1
```

Conclusion

As a result of the study the system described was made. It uses Vagrant for virtual machine man-agement, SLURM for batch job processing, Python [Sanner, 1999] for automation, Virtualbox [Oracle VM…] for debugging and DigitalOcean [DigitalOcean…] as a cloud provider. It was tested on a num-ber of problem classes and noticeable speedup was shown.

It is possible to extend the system on other solvers e.g. CBC or Ipopt. Another possible im-provement may be made by making the system accessible on the Web. It is also planned to publish the source code on GitHub after some cleanup.

In conclusion we expect that the service may become popular among SCIP users. Another con-clusion is that cloud computing is very convenient and cheap nowadays which is definitely a good driver for developing new and nonconventional approaches.

References

Achterberg Tobias SCIP: solving constraint integer programs // Mathematical Programming Computa-tion. — 2009. — Vol. 1, No. 1. — P. 1–41.

Audet C., Dang K. C., & Orban D. Optimization of algorithms with OPAL // Mathematical Programming Computation. — 2012. — P. 1–22.

Baz M., Hunsaker B., Brooks P., & Gosavi, A. Automated tuning of optimization software parameters // University of Pittsburgh Department of Industrial Engineering Technical Report. — 2007. — 7.

DigitalOcean cloud hosting, https://digitalocean.com/

Oracle VM VirtualBox, https://www.virtualbox.org/

Sanner Michel F. Python: a programming language for software integration and development // J. Mol. Graph. Model. — 1999. — 17.1. — P. 57–61.

Vagrant, http://www.vagrantup.com/

Yoo Andy B., Morris A. Jette, and Mark Grondona. SLURM: Simple linux utility for resource management // Job Scheduling Strategies for Parallel Processing. — Springer Berlin Heidelberg, 2003.