CHEP 2018
9-13 July 2018
National Palace of Culture
Sofia, Bulgaria

Modeling allocation utilization strategies on Supercomputers

Alexey Poyda[1], **Mikhail Titov**[1], Alexei Klimentov[1,2], Jack Wells[3], Sarp Oral[3], Kaushik De[4], Danila Oleynik[4], Shantenu Jha[2,5]

[1] National Research Centre «Kurchatov Institute» (RU), [2] Brookhaven National Laboratory (US), [3] Oak Ridge National Laboratory (US), [4] University of Texas at Arlington (US), [5] Rutgers University (US)

Abstract

Most supercomputers provide computing resources that are shared between users and projects, with utilization determined by predefined policies, load and quotas. The efficiency of the utilization of resources in terms of user/project depends on particular supercomputer policy and dynamic workload of supercomputer based on users' activities. The load on a resource is characterized by the number and parameters of jobs: the number of required nodes per job, required execution time called walltime per job, and jobs generation rate. We introduce the concept of an **execution strategy**, which is defined as the set of values of denoted parameters that uniquely define the group of jobs to be executed.

The aim is to find execution strategies that maximize the probability of utilizing a certain allocated resources, as well as the execution strategy that will optimize the utilization of a given number of core-hours on a resource.

The found execution strategy is applied to the upcoming jobs over a particular time period, and it is called static for that period. Our working hypothesis, is that in the most cases the well-defined static strategy gives a higher probability of successful application than, for example, a random dynamic strategy.

Methods

A **quantitative model** to estimate the probability of a given number of core-hours (i.e., allocation time) being utilized

- trained by the previous processes of utilization of allocation time;
- represented by the equation which calculates the probability that utilization U during the time interval T_0 will reach or exceed the predefined value U_0 (equation 1)
  - $f(x, \mu, \sigma^2)$ is a function of probability density of the normal distribution $N(\mu, \sigma^2)$;
  - μ and σ^2 are expected value and variance of a random variable describing duration of waiting time in the queue for jobs correspondingly;
  - μ_U and σ_U^2 - the same as previous, but for a random variable describing utilization of one job.

$$P(U > U_0) = \sum_{n=100}^{\infty} \left[\int_{U_0}^{\infty} f\left(x, n\mu_U, n\sigma_U^2\right) dx \left(\int_{-\infty}^{T_0} f(x, n\mu, n\sigma^2)\, dx - \int_{-\infty}^{T_0} f(x, (n+1)\mu, (n+1)\sigma^2)\, dx \right) \right]$$

Equation 1. Probability that utilization will reach the defined utilization value during the defined time period (cumulative distribution function).

Both the quantitative model and the simulator are statistical innature and are reliable for large job counts over long duration.

A **simulator** [1] that emulates the load on a supercomputer and produces (synthetic) job traces for a given partial workload (has the possibility to set requirements and restrictions of a particular supercomputer)

- based on queueing theory;
- used for the quantitative model validation and adjustment;
- provides a job state model (figure 1a) with the following states: generated, holding, pending, starting, executing, finished.

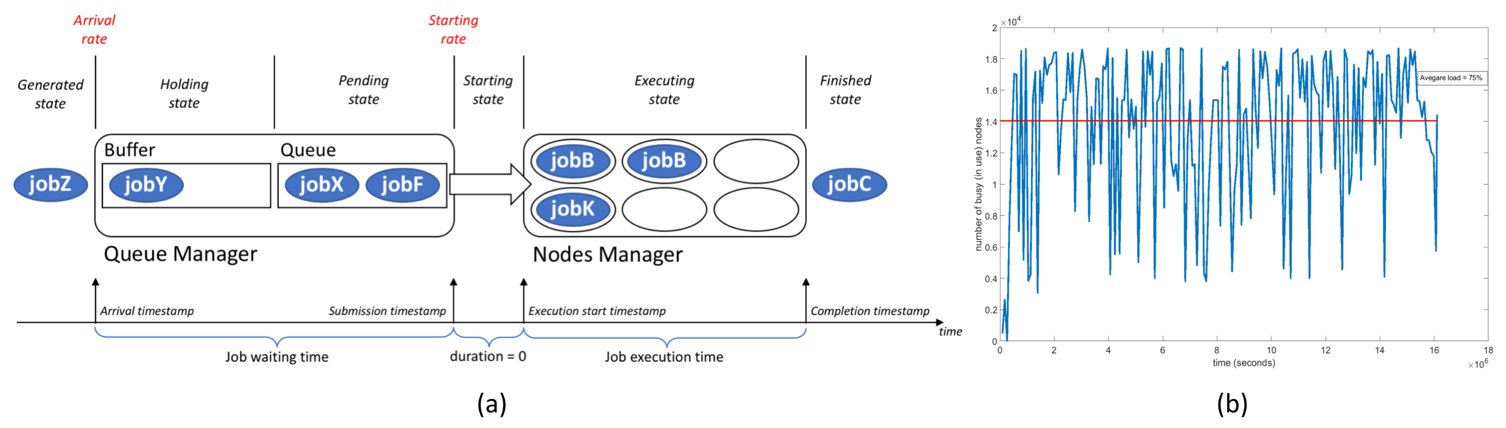

Figure 1. The simulator that emulates the workload of a supercomputer
(a) job state transitions; (b) simulated load (18 688 nodes).

Experiments

Quantitative model testing with synthetic data

- Common parameters for the quantitative model and for the simulator
  - number of nodes / cores per job = 1
  - job waiting time and execution time characteristics (same values for both parameters)
    - expected value, variance = 1, 1
  - the total processing time = 5000 (time units / hours)
- Specific parameters for a simulation process
  - job waiting time is defined according to the Poisson distribution
  - job execution time is defined according to the Normal distribution
  - job launching scheme: one stream and there is always one job in the queue
  - the total number of simulation runs = 100

Figure 2 shows the plot with two lines that represent the probability that a given utilization will be achieved in a given time interval. The blue line corresponds to the results obtained on the simulator, while the red line corresponds to calculations with the quantitative model.

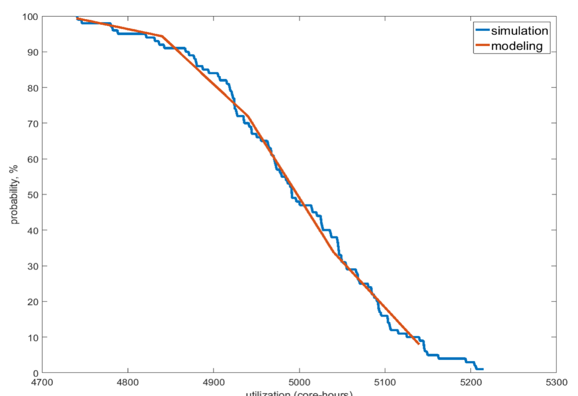

Figure 2. Probability (axis Y) that utilization will reach the corresponding utilization value (axis X) during the time of 5000 hours.

Titan log data analysis

- Log data characteristics
  - collected for 6 months (from aug'17 to jan'18)
  - contains information about job processing
    - job arrival timestamp (to the queue)
    - execution start timestamp
    - completion timestamp
    - the number of required nodes
      - 1 node = 16 cores at Titan
    - requested walltime
- Analysis actions
  - all jobs are divided into categories according to the number of required nodes and the volume of walltime requested (every category corresponds to a particular Titan's bin, where bin is a group of jobs that are treated equally)
  - for each category the following values are calculated: the expected value and variance of the random variables describing waiting time in the queue and the utilization achieved by a single job
  - obtained values were used as input data in equation 1 to calculate the probability that jobs of a given category will be able to utilize provided allocation in 3 months
  - job launching scheme: one stream and there is always one job in the queue

Figure 3 shows the resulted plots with probability distributions for all our categories. Every plot is associated with a certain Titan's bin. The most successful jobs are from the third bin, that required from 313 to 400 nodes and requested from 4 to 12 hours of walltime, as well as jobs from the first and second bins.

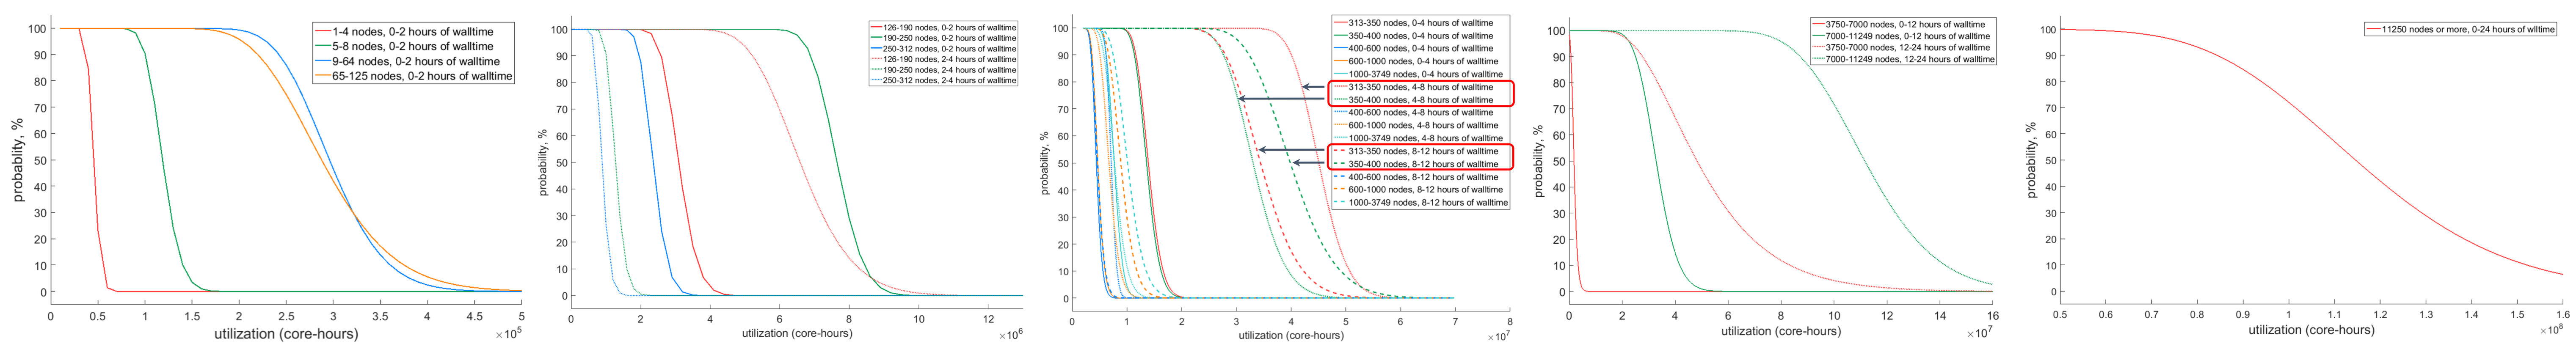

Figure 3. Probability distribution of utilization of allocation time during 3 months (based on full Titan log data for 6 months)
(a) bin #5; (b) bin #4; (c) bin #3; (d) bin #2; (e) bin #1.

Conclusions

The outcome of our research brings designed tools that are used to estimate the potential utilization of allocated resources on a particular supercomputer according to its workload. It assists in job parameters adjustment to reach the required utilization value or the maximum one among others. The benefit of this modeling process is the possibility to regulate the time of the utilization process. It still requires more fine-tune and understanding the accuracy of the model and simulator (e.g., estimation of the time scale for periods of historical data dump and the strategy time period). The designed solution is applied to large scale processing that extends towards exascale. This is early work and it is being developed to consider different kinds of workflows as well as different types of workloads and heterogeneous resources.

National Research Centre Kurchatov Institute
BROOKHAVEN NATIONAL LABORATORY
U.S. DEPARTMENT OF ENERGY | Office of Science
RUTGERS THE STATE UNIVERSITY OF NEW JERSEY
UNIVERSITY OF TEXAS ARLINGTON
OAK RIDGE National Laboratory

Acknowledgements

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. We also would like to express our appreciation to the Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC) and to all involved in this project, especially to OLCF working group. This work was funded in part by the Russian Ministry of Science and Education under contract No. 14.Z50.31.0024.

References

1. https://github.com/ATLAS-Titan/allocation-modeling