

SIMULATION OF ENERGY-EFFICIENT DEMAND RESPONSE FOR HIGH PERFORMANCE COMPUTING SYSTEMS

Kishwar Ahmed

Department of Computer Science
University of South Carolina Beaufort
1 University Blvd
Bluffton, SC 29909, USA

Jason Liu

School of Computing and Information Sciences
Florida International University
11200 SW 8th St
Miami, FL 33199, USA

ABSTRACT

Energy consumption is a critical issue for high-performance computing (HPC) systems. Demand response is a program offered by power service providers to reduce energy consumption. In this paper, we present the simulation of an energy-efficient economic demand-response model based on contract theory. Simulation is developed to examine the effectiveness of the demand-response model in a variety of real-world scenarios. Results from simulation show that the model can achieve energy efficiency through the rewarding mechanism of the contract-based design. In particular, the HPC users can be properly compensated to ensure their willing participation in energy reduction.

1 INTRODUCTION

The demand for high performance computing to support large-scale scientific applications has steadily increased over the last two decades or so. Recently, there has been an increasing effort to surpass the computation capability of the current high-performance computing systems by at least one order of magnitude to reach exascale. A major challenge for building future supercomputers is the energy consumption (SFP Division Plasma 2014). There has been significant effort to reduce supercomputer energy consumption. In particular, the HPC community in the U.S. is faced with the challenge of maintaining the energy consumption of the future exascale systems within 20-40 MWs. Energy cost contributes significantly to the total operational cost of HPC. For example, (Ludwig, and Dolz 2014) show that the power consumption is the second-most contributing factor to the total supercomputer cost, right after the hardware cost (energy cost contributes to 28% of the overall cost). A future exascale supercomputer with approximately 20 – 40 MWs power consumption is expected to incur \$20 – \$40 million cost per year.

A number of programs have been offered from power grid to reduce energy cost. Demand response is a popular program, which provides a means to reduce energy consumption during peak electricity price period. The global market for demand response program has been projected to grow to \$9.7 billion by 2023, a 400% increase since 2014 (National Renewable Energy Laboratory 2018). An economic demand response is a program that enables voluntary participation from the participants, without any prior commitment. Economic demand response participation has also been steadily increasing over the years. For example, PJM Interconnection, a regional electric transmission organization in the United States, observed a 4.7% increase in demand response participation from \$2.17 million in first nine months of 2017 to \$2.27 million in the first nine months of 2018 (PJM Interconnect 2018). Participation from different sectors, including data centers and smart buildings, has significantly increased to reduce their energy bills.

This paper presents a simulation study of a contract-based demand-response participation model designed specifically for HPC systems. Contract theory is a tool from microeconomics, where principle (or employer) offers contracts to the agents (or employees). Since participation in economic demand response program may cause “inconvenience” to the HPC users (with increased job execution time), we use to contract theory

to design appropriate rewards to HPC users to ensure their willing participation in the demand response program. The main contribution of this paper can be summarized as follows:

- We developed a trace-based simulator for the contract-based demand-response model, aimed at examining the HPC system's participation in economic demand response program.
- We studied several economic HPC demand-response participation scenarios. The simulation scenarios represent different types of HPC users' behavior (e.g., greediness, apathetic to participation). We demonstrate that the economic demand response model is capable of ensuring HPC system's participation for energy reduction through willing contribution from HPC users.

The rest of this paper is organized as follows. Section 2 presents the background and related work. Section 3 describes the framework for HPC resource allocation based on economic demand response participation. In Section 4, we present the simulator and a number of simulation scenarios which we use to examine effectiveness of demand-response model in ensuring contract theory property. Section 5 presents our conclusions.

2 BACKGROUND AND RELATED WORK

In this section, we provide the background and discuss existing work related to the demand-response simulation model for HPC systems. We first introduce demand response and discuss existing work on demand response for data center and HPC systems. Next, we describe existing approaches for job scheduling simulation for HPC systems and provide a brief background on a parallel discrete-event simulation engine used to develop our scheduling and resource allocation model. Finally, we present a brief introduction of contract theory and its applications.

2.1 Demand Response

Demand response programs are designed to help the energy service providers to stabilize the power grid during peak electricity period or during emergency incidents (e.g., extreme bad weather). Demand response can be broadly categorized into two types: economic demand response and emergency demand response. In economic demand response, participants voluntarily enroll in the programs (without the need of prior commitment) and willingly reduce the load based on economic incentives offered by the supplier. Emergency demand response requires prior commitment from the participants; once enrolled, it is mandatory for the participants to reduce the energy consumption to requested levels when supply shortage situations or emergency conditions occur.

Demand response has been noted as important policy goals to achieve power grid efficiency by both the the U.S. Department of Energy (DoE) and the National Institute of Standards and Technology (NIST) (Holmberg et al. 2014; Federal Energy Regulatory Commission 2016). Overall, demand response has become increasingly popular among power utility companies, and as such has become a significant source of revenue earnings. Much of this growth is comes from economic demand response participation, where power grid signals participants to adapt their power consumption at various time granularity and provides incentives based on participation.

Energy cost reduction in data centers and smart buildings with consideration of demand response scheme has been studied quite extensively. Load shifting in time, geographical load balancing, speed-scaling, server consolidation, power-capping are some of the approaches proposed in the literature for data center's demand response (Wierman et al. 2014; Liu et al. 2013; Wang et al. 2016). These approaches however are applicable for internet transaction-based data center workload (e.g., map-reduce applications or network transaction based applications), not for HPC applications. For data center workload, the service time is typically assumed to be uniform and delay intolerant (most jobs need to be serviced within the hour from which they are submitted). HPC jobs are much less uniform both in terms of service time and job size (requested resources in the number of processors). Also, HPC users usually achieve for high

performance, while data center users may not have such constraint. As such, energy saving is only a secondary consideration to HPC users.

In our prior studies (Ahmed et al. 2017; Ahmed et al. 2018), we developed *emergency* demand response participation models for HPC systems. To enable emergency demand response, the HPC operators must establish an agreement *a priori* with the utility companies. HPC systems obligatorily reduce energy consumption in response to signals received from utility companies in the event of transient surge in power demand or other emergency events. Our demand response model considers HPC job scheduling and resource provisioning algorithms (based on frequency scaling, power capping, and node scaling) to allow power-bound energy-conservation during the critical demand response periods, while maintaining the traditional performance-optimized job scheduling during normal operation. We also developed an economic demand response model for HPC system using contract theory (Ahmed et al. 2018). Unlike *emergency* demand response, *economic* demand response allows participants to voluntarily participate in the program to earn rewards without prior commitment.

This paper focuses on the economic demand response model. The focus of this study is using simulation methods to create different scenarios to examine the effectiveness of the contract-based design for HPC demand response.

2.2 Job Scheduling Simulation

There exist many job scheduling simulators particularly focusing on HPC systems. For example, PYSS (Python Scheduler Simulator) is an open-source HPC workload scheduling simulator written in Python (Parallel Systems Lab 2010). The simulator was developed by the Experimental System Lab at the Hebrew University, and has been used to study various scheduling algorithms (e.g., (Georgiou et al. 2015; Liu and Weissman 2015)). CQSim is another event-based simulator to study the detailed queuing behavior of job schedulers using real system workload. The simulator was developed by Illinois Institute of Technology and has been used to evaluate fault-aware utility-based job scheduling (Tang et al. 2009), adaptive metric-aware job scheduling (Tang et al. 2012), and so on.

Our job scheduler simulator is developed based on Simian, which is an open-source, process-oriented, parallel discrete-event simulation engine (Santhi, Eidenbenz, and Liu 2015). Simian has several unique design features that make it more attractive for us to build our scheduler simulator (e.g., simple API, minimalist design approach). Also, there has been a significant efforts in developing models for HPC architectures and applications using Simian (e.g., (Ahmed et al. 2016; Ahmed et al. 2016; Chennupati et al. 2017)). Our job scheduler can take advantage of these models.

2.3 Contract Theory

A contract is an agreement between involved parties that stipulates various actions that the parties are expected to take. A major problem studied by contract theory is how parties contract in order to deal with the presence of hidden or asymmetric information. In the *adverse selection* model, an uninformed “principal” offers contracts to a privately-informed “agent” who decides whether to accept or refuse the offered contracts. Akerlof studied this in the context of a used-car market where sellers had private information about the quality of their cars in 1970 (Akerlof 1970), and since then the economics literature has applied this to many types of applications in which the agent has private information about her productive ability or cost of effort. Contract theory has been applied to various areas, such as economics and communication systems. In communications, contract theory has been applied in radio networks (Zhao et al. 2015), vehicular networks (Gao et al. 2013; Wang et al. 2014), cellular networks (Liu et al. 2017), wireless networks (Zhang et al. 2018; Le et al. 2017), smart grids (Zhang et al. 2018), and Internet of Things (Hou et al. 2017)).

3 CONTRACT-BASED DEMAND RESPONSE MODEL

In this section, we briefly discuss the overall framework of our demand-response aware scheduling and resource allocation model for HPC systems. A more detailed discussion of the contract-based design can be found in (Ahmed et al. 2018).

We consider an HPC system with one HPC operator and a set of job types $\Theta = \{\theta_1, \theta_2, \dots, \theta_i, \dots, \theta_n\}$. Our model determines per-job resources (e.g., processor frequency, power-capping values) and rewards to be provided to the HPC users for their willing participation. We assume that a job i runs at frequency f_i on the requested processors. We denote the processor's power consumption at frequency f_i for job i by $p(f_i)$ and the execution time for job of i at frequency f_i by $t(f_i)$. To reduce power consumption, the HPC operator determines the optimal processor frequency f'_i for running job i . After changing the frequency to f'_i , the power consumption and execution time are denoted by $p(f'_i)$ and $t(f'_i)$, respectively. We assume that by default the jobs are run at highest frequency to achieve the best performance. Energy consumption for job i at maximum frequency can be represented as following: $e_i = n_i \cdot p(f_{max}) \cdot t_s$, where t_s denotes duration of job execution and n_i denotes number of processors required for job type i . We assume job i is allocated to homogeneous processors at running at same frequencies. However, our model can be easily extended to heterogeneous processors by including number and type of the processors as decision variable. Energy consumption for job i at frequency f'_i can be represented as following: $e'_i = n_i \cdot p(f'_i) \cdot t_s$. Therefore, energy reduction is $\Delta e_i = e_i - e'_i$ and the percentage change in execution time is $\Delta t_i = 100 \cdot [t(f_{max}) - t(f'_i)] / t(f_{max})$.

To ensure willing participation of HPC users, they need to be compensated for application performance changes. θ_i denotes the flexibility of jobs of type i to participate in demand response. A job with higher value of θ_i are less flexible to participate in demand response (i.e., can tolerate little delays in job execution). The inconvenience cost for jobs of type i is a function of changes in the job's execution time, and can be represented as: $v(\theta_i, \Delta t_i) = \theta_i \cdot c(\Delta t_i)$. Here, we assume that $c(\Delta t_i) = c_0 \cdot (\Delta t_i)^2$ (Zhao et al. 2015; Kordali and Cottis 2015). We define the utility for jobs of type i as follows: $u_i = r_i - v(\theta_i, \Delta t_i), \forall i \in \{1, 2 \dots n\}$, and define the HPC operator's utility function for jobs of type- i as follows: $u_{op}^i = m_i \cdot (\phi \cdot \gamma \cdot \Delta e_i - r_i)$. Here m_i denotes the number of jobs for job type i , ϕ denotes electricity price, r_i denotes reward awarded for job type i and γ is the factor of power usage effectiveness (PUE), which is used to capture the power consumption of non-IT parts of the HPC systems. The total utility for all jobs is thus $u_{op} = \sum_{i=1}^n u_{op}^i$.

The optimization objective of the HPC operator can be given as following:

$$\max_{(\Delta e, r)} m_i \cdot \left[\sum_{i=1}^n \phi \cdot \gamma \cdot \Delta e_i - r_i \right] \quad (1)$$

$$\text{s.t.} \quad r_i - \theta_i \cdot c(\Delta t_i) = 0, \forall i \quad (2)$$

$$r_i - \theta_i \cdot c(\Delta t_i) \geq r_{i'} - \theta_{i'} \cdot c(\Delta t_{i'}), \forall i, i', \text{ and } i \neq i' \quad (3)$$

$$f_{min} \leq f'_i \leq f_{max} \quad (4)$$

Equation 2 represents the individual rationality (IR) constraint, that ensures participants in contract mechanism achieve non-negative pay-off or utility. Equation 3 denotes the incentive compatibility (IC) constraint and represents the fact that a participants utility is maximized when it chooses the contract intended (or designed) for its type over those intended for the other contract.

The problem in Equations 1– 4 is intractable to solve due to large number of IR and IC constraints. However, we can resort to constraint reductions to reduce the number of constraints and make the problem solvable. Using contract theory property, both IR and IC constraints can be reduced and represented as

following:

$$\begin{aligned}
 & \max_{(\Delta e, r)} m_i \cdot \left[\sum_{i=1}^n \phi \cdot \gamma \cdot \Delta e_i - r_i \right] \\
 & \text{s. t. } r_n - \theta_n \cdot c(\Delta t_n) = 0 \\
 & r_{i-1} - \theta_{i-1} \cdot c(\Delta t_{i-1}) = r_i - \theta_{i-1} \cdot c(\Delta t_i) \\
 & f_{min} \leq f_i \leq f_{max}
 \end{aligned}$$

We can solve the reduced optimization problem using standard optimization solving methods, such as the sequential least squares programming algorithm using the Han-Powell quasi-Newton method, and determine optimal frequency and reward amount for the HPC users. In the literature, other methods (Zhang et al. 2015; Zhang et al. 2016) have also been proposed to solve the same problem, such as taking Lagrangian multipliers and then solve the Lagrangian formulation.

4 ECONOMIC DEMAND RESPONSE SIMULATION

In this section, we describe our simulation study of the economic demand-response model. First, we describe the trace-based simulator we created for this purpose. Next, we present the data set we used in simulation. Finally, we describe the simulation scenarios, along with the experiment results.

For the simulation study, we consider scenarios. They are summarized as follows:

- **Scenario #1:** This is the baseline model. In this scenario, we assume that HPC users act accordingly, that is accepts contracts given to them based on their classification to different types.
- **Scenario #2:** We assume that HPC users may not select contract intended for them. We consider two cases: in first case, the HPC users *always* select contract from the same contract type; in the second case, the users *randomly* select a contract from the contract set.
- **Scenario #3:** We assume that some HPC users do not select any of the given contracts, and decides not to participate in the demand response program.
- **Scenario #4:** We consider a different job trace, to demonstrate effectiveness of the demand-response model for a different trace log.

4.1 Trace-Driven Simulator

Our trace-driven simulator is developed based on Simian, a parallel discrete-event simulation engine. The simulator consists of five major components: a job dispatcher, a job executioner, scheduling policies, application models, and a resource manager. The job dispatcher can handle different types of events, such as job arrival, job departure, job eviction (e.g., when an executed job is interrupted and removed in the middle of the run), and power demand change (e.g., changes in power limit from power grid during a demand response event). After a job is submitted, it is entered to the job waiting queue and job dispatcher is invoked. The job dispatcher determines which jobs to run from the waiting queue based on the scheduling policies, application models (that describes the power and performance characteristics of the jobs), and available resources from the resource manager. When a job is scheduled to run the job dispatcher removes the job from waiting queue and put it in the list of running jobs. The job executioner allocates the resources using the resource manager to represent the occupied processors with associated power consumption for running the job. The job executioner then simulates the job execution. When a job completes its execution, the job executioner removes the job from the list of running jobs and reclaims the resources occupied by the job. The job scheduler simulator has been validated against real-life data (Ahmed et al. 2017).

Our simulator has also been augmented to handle both emergency and economic demand response (Ahmed et al. 2017; Ahmed et al. 2018). We can schedule an event to indicate the power demand change, with a lower power limit upon the arrival of a emergency demand response event, or an another when the power

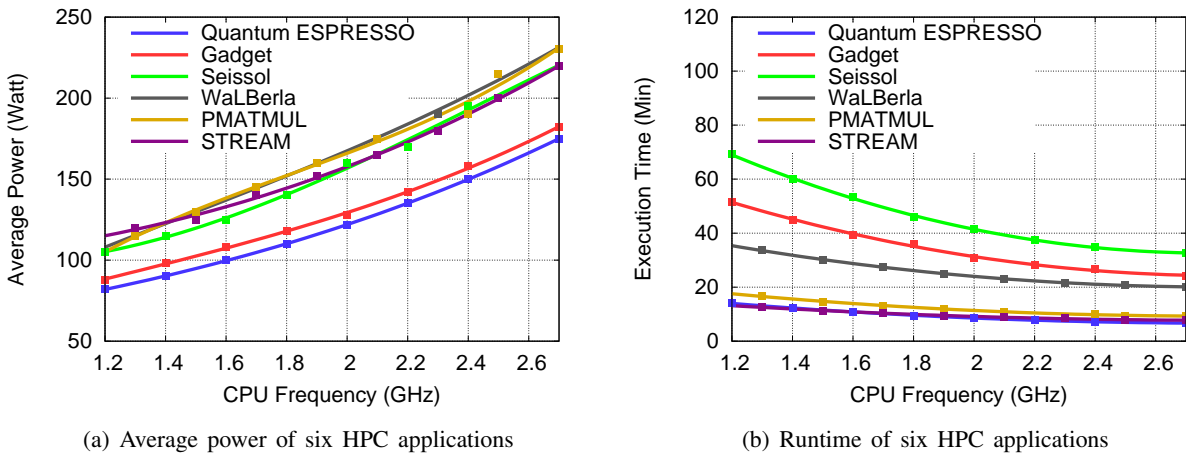


Figure 1: Data set used for different simulation scenarios.

limit returns to level for normal operations. In the former case, the job executioner may evict jobs if the current power level is no longer sufficient to support all running jobs. In the latter case, the job scheduler may start new jobs to run. Our simulator can also simulate economic demand response participation from HPC system. It can determine optimal resource allocation values to incoming jobs, as well as users' reward amount to enable economic demand response participation.

4.2 Data Set

Our considered simulation setup includes an HPC system with one HPC operator and six job types. We collected electricity price data from PJM Interconnect for a 24-hour period on a day of October 2017. We use 1.3 as the power usage effectiveness (PUE) of the target HPC center. PUE is a ratio of the total energy of the HPC center facility to the energy consumed by computing equipment.

We collected and used the power and performance data for six HPC applications from an existing study (Auweter et al. 2014). More specifically, we used the performance and power data at different frequencies for four scientific applications (including Quantum ESPRESSO (Giannozzi et al. 2009), Gadget (Springel 2005), Seissol (Käser et al. 2008) and WaLBerla (Feichtinger et al. 2011)) and two synthetic benchmarks (PMATMUL, a parallel benchmark for dense matrix multiplication, and STREAM (McCalpin, John D 2002), a benchmark for measuring sustainable memory bandwidth). We denote the job types as Type#1 to Type#6 correspondingly.

We used discrete frequency values for the processors, ranging from 1.2GHz to 2.4GHz at 0.2GHz intervals and 2.7GHz. That is, the minimum and maximum frequency are 1.2GHz and 2.7GHz, respectively. Figure 1 shows the empirical and fitted polynomial function for collected data. Figure 1(a) and Figure 1(b) present the power and runtime, respectively. The peak power of the processors was set to 240W (determined from the power consumption of the six HPC applications when running at the maximum frequency).

4.3 Simulation Scenario #1

This is an ideal scenario, where we assume that users behaves accordingly; that is, the users select contract intended for their own job type.

We collected real-life HPC workload trace from the parallel workload archive (Feitelson, Dror and others 2007). The workload traces contain runtime information collected at the Blue Gene/P system at Argonne National Laboratory from January 2009 to September 2009. The trace contains 68,936 jobs. The trace includes information about job start time, job run time, job wait time, requested processors, etc. We

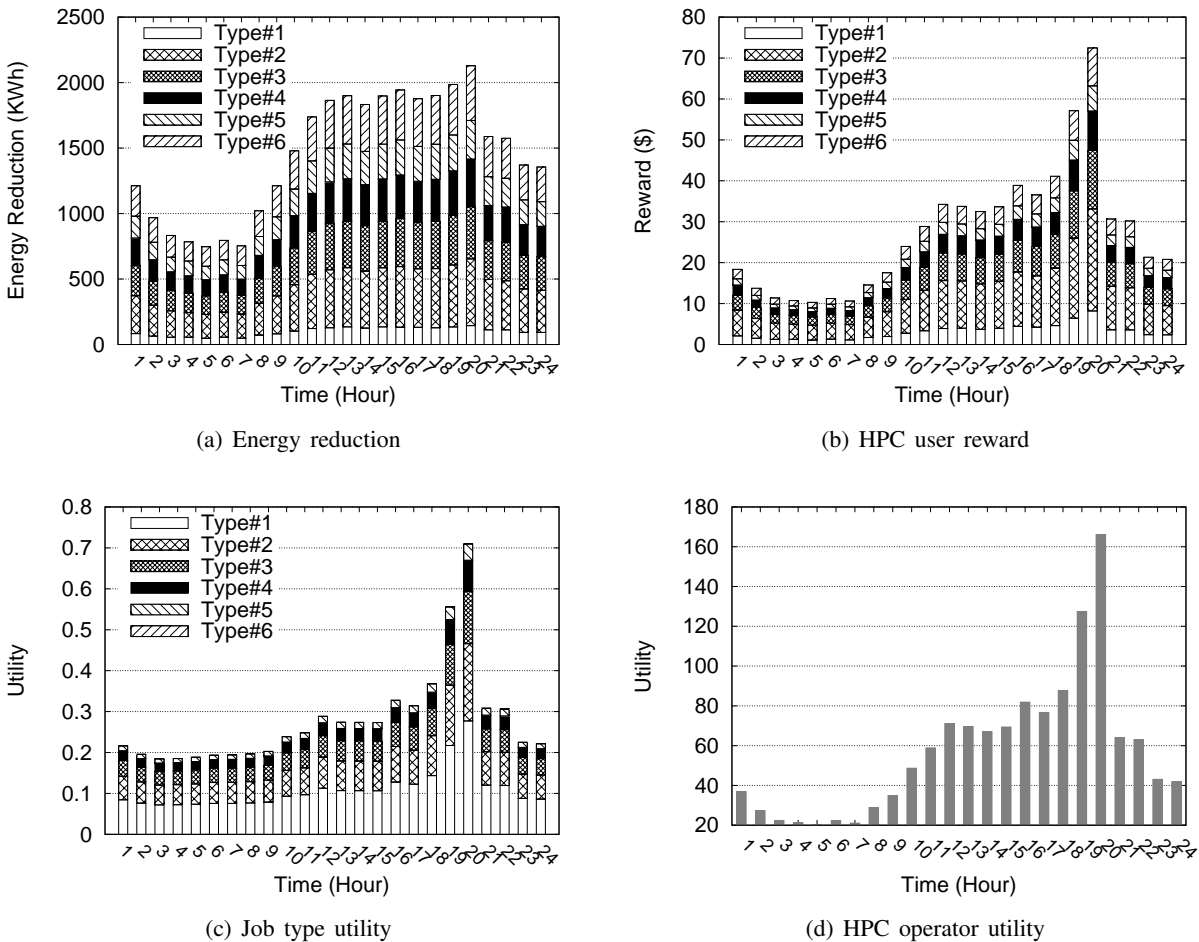



Figure 2: Scenario#1: simulation of demand-response model.

processed the data and collected the number of job arrivals during each hour to determine percentage job arrivals throughout a day.


Figure 2 presents results for energy reduction, users' rewards, and users' and operator's utilities. Figure 2(a) shows energy reduction by all jobs of different types. As can be seen in Figure 2(a), different jobs participate in energy reduction at varying level to ensure demand response participation. Figure 2(b) demonstrates rewards awarded to all the jobs of various job types. During high electricity price periods, the users contribute more to the energy reduction, and therefore are rewarded more. Figure 2(c) shows utility (i.e., the difference of reward and inconvenience cost) of each job for different job types. As evident from the figure, different job types have non-negative utility. Figure 2(d) shows the HPC operator's utility (i.e., the participation benefit of HPC system after rewarding the HPC users). As can be seen in Figure 2(d), the operator will have positive utility at different time period, which means that the demand response participation will present participation benefits to the HPC operator, even after rewarding the HPC users for their participation.

4.4 Simulation Scenario #2

In this scenario, we assume that HPC users select contract that is not intended for the type to which they belong to. First, we consider a case where all the users select contract offered to one particular type (e.g., Type#1 in our simulation). Figure 3(a) shows utility of each job for each job type in such case. As can be

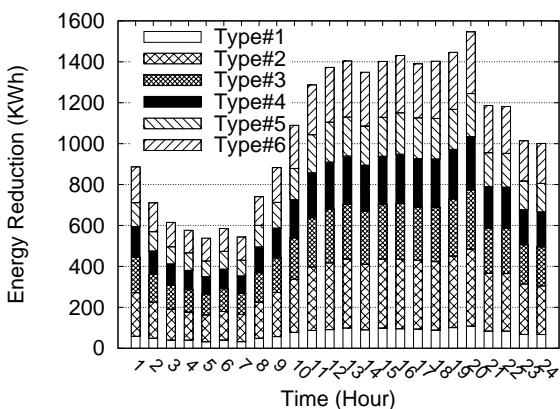


(a) Always select same contract




(b) Randomly select contract


Figure 3: Scenario#2: HPC users select other contracts.




(a) Energy reduction



(b) HPC user reward



(c) Job type utility



(d) HPC operator utility

Figure 4: Scenario#3: some HPC users reject the contracts.

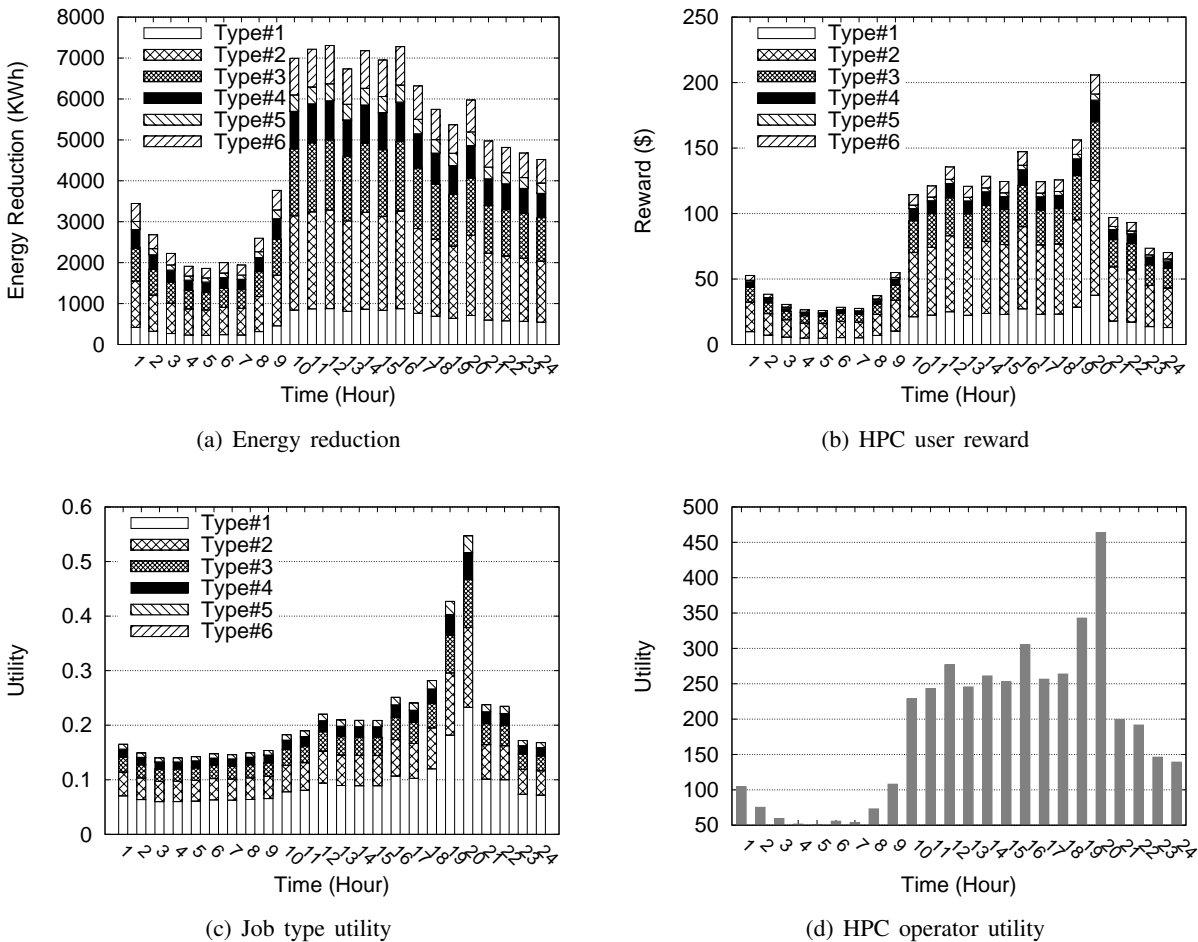


Figure 5: Scenario#4: simulation with workload collected from MetaCentrum Czech National Grid.

seen in the figure, if the users do not select contract for their own type, they may incur negative utility and therefore incur loss. Moreover, users have lower utility throughout the entire time period compared to the results from scenario #1 (Figure 2(c)).

In second case, the users randomly select contract from another type. Figure 3(b) presents per-job utility for different types when users randomly select contract from the uniformly distributed given contract set. As can be seen in the figure, the users achieve lower (and sometimes negative) utility for such participation compared to scenario #1 (Figure 2(c)).

Overall, this simulation scenario proves that our designed contract-based demand-response model ensures that each user will select contract that are offered to them to achieve the maximum participation benefit.

4.5 Simulation Scenario #3

In this scenario, we assume that some users do not accept any of the contracts given in the contract set. More specifically, we consider a situation where approximately 25% of HPC users did not select any of the contracts, and therefore did not participate in the demand response at all.

Figure 4 presents experiment results for energy reduction, users' rewards, and users' and operator's utilities. Figure 4(a) shows energy reduction by all jobs of different types. The energy reduction amount is lower than scenario #1 (where all the users participated in the demand response). However, HPC operator

can resort to energy storage device (e.g., batteries) to compensate low participation of users. We plan to investigate this case in our future study. Figure 4(b) demonstrates rewards awarded to all the jobs of various job types. As expected, the rewards earned by HPC users are lower compared to scenario #1 due to their low participation. However, per-job utility remains the same as scenario #1, as evident from Figure 4(c). Therefore, overall lower participation of HPC users does not effect the participation benefit of users who participate in the program. Figure 4(d) shows the HPC operator's utility. As can be seen in the figure, HPC operator achieves non-negative utility throughout the entire time period. That is, participation in demand response will always be beneficial to the HPC operator, even after rewarding the users for their participation. The utility is higher during the time period with higher job arrivals or with higher electricity price.

4.6 Simulation Scenario #4

We collected and used a different set of trace data for this scenario. The job trace is collected from MetaCentrum Czech National Grid from January 2013 to April 2015. The workload contains 5,731,100 jobs.

Figure 5 shows the results for running this workload. With increased number of job arrivals per hour, the possibility of energy reduction increases for this workload. As can be seen in Figure 5(a), various amount of energy reduction is possible throughout time period. Figure 5(b) shows reward earned by different job types. The rewards fluctuate in accordance with the job arrival rate. With a higher number of job arrivals, the demand response algorithm offers higher rewards to users to incentivize them to take on low-energy execution, as it tries to lower the overall energy consumption of the entire system. Figure 5(c) and Figure 5(d) show utility for different job types and operator, respectively. Utility are non-negative for both cases, showing that demand response participation is beneficial for both HPC users and HPC operator with different job trace.

5 CONCLUSION

In this paper, we present a simulation study of economic demand response participation model for HPC systems. Our model is designed based on contract theory mechanism to reward HPC users for their willing participation in demand response. We created a trace-based simulator and designed various simulation scenarios to show the effectiveness of the contract-based demand response participation model. The various simulation scenarios replicate different behaviors of HPC users (e.g., greediness, apathetic to participation). The trace-based simulation studies show that the contract-based economic demand-response model for HPC systems can ensure willing participation from HPC users, and therefore prove that the economic demand response participation model can be realized for HPC systems in practice to allow energy reduction.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation EPSCoR Program under NSF Award # OIA-1655740. Any Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation.

REFERENCES

- Ahmed, K., J. Bull, and J. Liu. 2018. "Contract-Based Demand Response Model for High Performance Computing Systems". In *Proceedings of the 2018 IEEE Intl Conf on Parallel & Distributed Processing with Applications (ISPA)*, 580–589. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Ahmed, K., J. Liu, S. Eidenbenz, and J. Zerr. 2016. "Scalable Interconnection Network Models for Rapid Performance Prediction of HPC Applications". In *Proceedings of the 2016 IEEE 18th International Conference on High Performance Computing and Communications (HPCC)*, 1069–1078. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Ahmed, K., J. Liu, and X. Wu. 2017. "An Energy Efficient Demand-Response Model for High Performance Computing Systems". In *Proceedings of the 2017 IEEE 25th International Symposium on Modeling, Analysis, and Simulation of*

- Computer and Telecommunication Systems (MASCOTS)*, 175–186. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Ahmed, K., J. Liu, and K. Yoshii. 2018. “Enabling Demand Response for HPC Systems through Power Capping and Node Scaling”. In *Proceedings of the 2018 IEEE 16th International Conference on High Performance Computing and Communications (HPCC)*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Ahmed, K., M. Obaida, J. Liu, S. Eidenbenz, N. Santhi, and G. Chapuis. 2016. “An Integrated Interconnection Network Model for Large-Scale Performance Prediction”. In *Proceedings of the 2016 annual ACM Conference on SIGSIM Principles of Advanced Discrete Simulation*, 177–187. Association for Computing Machinery.
- Akerlof, G. 1970. “The market for lemons”. *Quarterly Journal of Economics* 89:488–500.
- Auweter, A., A. Bode, M. Brehm, L. Brochard, N. Hammer, H. Huber, R. Panda, F. Thomas, and T. Wilde. 2014. “A Case Study of Energy Aware Scheduling on SuperMUC”. In *International Supercomputing Conference*, 394–409. Springer.
- Chennupati, G., N. Santhi, R. Bird, S. Thulasidasan, A.-H. A. Badawy, S. Misra, and S. Eidenbenz. 2017. “A scalable analytical memory model for CPU performance prediction”. In *International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems*, 114–135. Springer.
- Federal Energy Regulatory Commission 2016. “Assessment of Demand Response and Advanced Metering”. <https://www.ferc.gov/legal/staff-reports/2016/DR-AM-Report2016.pdf>. accessed 2nd April 2017.
- Feichtinger, C., S. Donath, H. Köstler, J. Götz, and U. Rüde. 2011. “WalBerla: HPC Software Design for Computational Engineering Simulations”. *Journal of Computational Science* 2(2):105–112.
- Feitelson, Dror and others 2007. “Parallel Workloads Archive”. <http://www.cs.huji.ac.il/labs/parallel/workload/>. accessed 14th August 2016.
- Gao, Y., Y. Chen, C. Wang, and K. J. R. Liu. 2013. “A Contract-Based Approach for Ancillary Services in V2G Networks: Optimality and Learning”. In *Proceedings of the 2013 IEEE INFOCOM*, 1151–1159. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Georgiou, Y., D. Glesser, K. Rzadca, and D. Trystram. 2015. “A Scheduler-Level Incentive Mechanism for Energy Efficiency in HPC”. In *Proceedings of the Cluster, Cloud and Grid Computing (CCGrid), 2015 15th IEEE/ACM International Symposium on*, 617–626. Institute of Electrical and Electronics Engineers, Inc.
- Giannozzi, P., S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo et al. 2009. “QUANTUM ESPRESSO: A Modular and Open-Source Software Project for Quantum Simulations of Materials”. *Journal of physics: Condensed matter* 21(39):395502.
- Holmberg, D. G., S. T. Bushby, and D. B. Hardin. 2014. “Facility Smart Grid Interface and a Demand Response Conceptual Model”. Technical Report NIST Technical Note 1832, NIST.
- Hou, Z., H. Chen, Y. Li, Z. Han, and B. Vucetic. 2017. “A Contract-Based Incentive Mechanism for Energy Harvesting-Based Internet of Things”. In *Proceedings of the 2017 IEEE International Conference on Communications (ICC)*, 1–6. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Käser, M., J. d. a. Puente, C. Castro, V. Hermann, and M. Dumbser. 2008. “Seismic Wave Field Modelling Using High Performance Computing”. In *SEG Technical Program Expanded Abstracts 2008*, 2884–2888. Society of Exploration Geophysicists.
- Kordali, A. V., and P. G. Cottis. 2015. “A Contract-Based Spectrum Trading Scheme for Cognitive Radio Networks Enabling Hybrid Access”. *IEEE Access* 3:1531–1540.
- Le, T. H. T., N. H. Tran, P. L. Vo, Z. Han, M. Bennis, and C. S. Hong. 2017. “Contract-Based Cache Partitioning and Pricing Mechanism in Wireless Network Slicing”. In *Proceedings of the 2017 IEEE Global Communications Conference*, 1–6. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Liu, F., and J. B. Weissman. 2015. “Elastic Job Bundling: An Adaptive Resource Request Strategy for Large-Scale Parallel Applications”. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 33. Association for Computing Machinery.
- Liu, T., J. Li, F. Shu, M. Tao, W. Chen, and Z. Han. 2017. “Design of Contract-Based Trading Mechanism for A Small-Cell Caching System”. *IEEE Transactions on Wireless Communications* 16(10):6602–6617.
- Liu, Z., A. Wierman, Y. Chen, B. Razon, and N. Chen. 2013. “Data Center Demand Response: Avoiding the Coincident Peak via Workload Shifting and Local Generation”. *Performance Evaluation* 70(10):770–791.
- Ludwig, and Dolz 2014. “Total Cost of Ownership in High Performance Computing”. https://wr.informatik.uni-hamburg.de/_media/teaching/sommersemester_2014/tco-14-anna-lena.pdf. accessed 7th April 2019.
- McCalpin, John D 2002. “STREAM benchmark”. <http://www.cs.virginia.edu/stream/stream2>. accessed 14th August 2016.
- National Renewable Energy Laboratory 2018. “Potential Roles for Demand Response in High-Growth Electric Systems with Increasing Shares of Renewable Generation”. <https://www.nrel.gov/docs/fy19osti/70630.pdf>. accessed 7th April 2019.
- Parallel Systems Lab 2010. “Python Scheduler Simulator”. <https://code.google.com/archive/p/pyss/>. accessed 10th April 2017.
- PJM Interconnect 2018. “2018 Quarterly State of the Market Report for PJM: January through September”. https://www.monitoringanalytics.com/reports/PJM.State_of_the_Market/2018/2018q3-som-pjm-sec6.pdf. accessed 22nd February 2019.

- Santhi, n., S. Eidenbenz, and J. Liu. 2015. "The Simian Concept: Parallel Discrete Event Simulation with Interpreted Languages and Just-In-Time Compilation". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3013–3024. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- SFP Division Plasma 2014. "High-Performance Computing at Exascale: Challenges and Benefits". <https://hal.inria.fr/hal-01820511/document>. accessed 5th April 2019.
- Springel, V. 2005. "The Cosmological Simulation Code GADGET-2". *Monthly notices of the royal astronomical society* 364(4):1105–1134.
- Tang, W., Z. Lan, N. Desai, and D. Buettner. 2009. "Fault-Aware, Utility-Based Job Scheduling on Blue, Gene/P Systems". In *Proceedings of the 2009 IEEE International Conference on Cluster Computing and Workshops*, 1–10. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Tang, W., D. Ren, Z. Lan, and N. Desai. 2012. "Adaptive Metric-Aware Job Scheduling for Production Supercomputers". In *Proceedings of the 2012 41st International Conference on Parallel Processing Workshops*, 107–115.
- Wang, H., J. Huang, X. Lin, and H. Mohsenian-Rad. 2016. "Proactive Demand Response for Data Centers: A Win-Win Solution". *IEEE Transactions on Smart Grid* 7(3):1584–1596.
- Wang, Y., H. Nazari-pouya, C.-C. Chu, R. Gadh, and H. R. Pota. 2014. "Vehicle-to-Grid Automatic Load Sharing with Driver Preference in Micro-Grids". In *Proceedings of the Innovative Smart Grid Technologies Conference Europe (ISGT-Europe), 2014 IEEE PES*, 1–6. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Wierman, A., Z. Liu, I. Liu, and H. Mohsenian-Rad. 2014. "Opportunities and Challenges for Data Center Demand Response". In *International Green Computing Conference*, 1–10. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Zhang, B., C. Jiang, J.-L. Yu, and Z. Han. 2016. "A Contract Game for Direct Energy Trading in Smart Grid". *IEEE Transactions on Smart Grid* 9(4):2873–2884.
- Zhang, D., Z. Chang, T. Hmlinen, and W. Gao. 2018, April. "A Contract-Based Resource Allocation Mechanism in Wireless Virtualized Network". In *Proceedings of the IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 474–479. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Zhang, K., Y. Mao, S. Leng, Y. He, S. Maharjan, S. Gjessing, Y. Zhang, and D. H. Tsang. 2018. "Optimal Charging Schemes for Electric Vehicles in Smart Grid: A Contract Theoretic Approach". *IEEE Transactions on Intelligent Transportation Systems* 19(9):1–13.
- Zhang, Y., L. Song, W. Saad, Z. Dawy, and Z. Han. 2015. "Contract-Based Incentive Mechanisms for Device-to-Device Communications in Cellular Networks". *IEEE Journal on Selected Areas in Communications* 33(10):2144–2155.
- Zhao, N., M. Wu, W. Xiong, and C. Liu. 2015. "Optimal Contract Design for Cooperative Relay Incentive Mechanism Under Moral Hazard". *Journal of Electrical and Computer Engineering* 2015:56.

AUTHOR BIOGRAPHIES

KISHWAR AHMED is an Assistant Professor at the Department of Computer Science, University of South Carolina Beaufort. He received a Ph.D. degree from School of Computing and Information Sciences, Florida International University. His research interests include high-performance computing, energy-efficiency, and optimization. His email address is ahmedk@uscb.edu.

JASON LIU is an Associate Professor at the School of Computing and Information Sciences, Florida International University. He received a Ph.D. degree from Dartmouth College in Computer Science. His research focuses on parallel simulation and high-performance modeling of computer systems and communication networks. His email address is liux@cis.fiu.edu.