RGC Ref.: N_HKUST610/11 NSFC Ref. :61161160566
(please insert ref. above)

The Research Grants Council of Hong Kong
NSFC/RGC Joint Research Scheme
Joint Completion Report

(Please attach a copy of the completion report submitted to the NSFC by the Mainland researcher)

Part A: The Project and Investigator(s)

1. Project Title

The Minimized Energy Consumptions and Maximized Resource Utilizations in Large-scale Datacenters

2. Investigator(s) and Academic Department/Units Involved

	Hong Kong Team	Mainland Team
Name of Principal Investigator *(with title)*	Bo Li (Prof.)	Zhiyong Liu (Dr.)
Post	Professor	Professor
Unit / Department / Institution	CSE/HKUST	Institute of Computing Technology/Chinese Academy of Science
Co-investigator(s) *(with title)*		Fa Zhang (Dr.), Associate Professor

3. Project Duration

	Original	Revised	Date of RGC/ Institution Approval *(must be quoted)*
Project Start date	Jan. 1, 2012		
Project Completion date	Dec. 31, 2014		
Duration *(in month)*	36		
Deadline for Submission of Completion Report	Dec 31, 2015		

Part B: The Completion Report

5. Project Objectives

5.1 Objectives as per original application

1. To characterize and predict the application workload in datacenters.
2. To investigate the fundamental trade-off between the system performance and resource utilization in datacenters along with the energy consumption
3. To propose a multi-objective optimization framework for both resource and energy-efficient optimization.
4. To design dynamic task allocation and virtual machine migration algorithms towards resource and energy optimization.

5.2 Revised Objectives

N.A.

6. Research Outcome

Major findings and research outcome
(maximum 1 page; please make reference to Part C where necessary)

With the presence of unpredictable and time-varying application demands, on the one hand, we want to maximize the profit by accepting and processing as many application requests as possible, which can lead to maximal resource utilization; on the other hand, we seek to minimize the penalty from system congestion due to the potentially excessive requests and potentially excessive power consumption.

To balance such a power-performance tradeoff, we enforced the following three control decisions: 1) how many requests from diverse applications are to be admitted at any given time; 2) how to distribute the admitted requests from different applications across a large number of servers hosting the corresponding VMs; and 3) how to schedule each VM by switching between a running state for processing requests and an idle state for conserving the server power.

To address these challenges, in this project, we took advantage of the Lyapunov optimization techniques to design and analyze a new optimal online control framework, designed to independently and concurrently make all three decisions. Our findings based on the control decision led to three simple yet effective strategies, corresponding to the each decision that need to be made: 1) a threshold based admission control strategy to improve system throughput while avoiding system congestion; 2) a ''Join the Shortest Queue'' request routing strategy for balancing the server load, while reducing service delays for the admitted requests; and 3) a decentralized greedy strategy to optimally schedule VMs, i.e., which VMs are to process incoming requests, and which VMs are to stay idle for power conservation.

The upshot of our proposed framework is that it can be flexibly extended to explore various design choices and operational requirements of a realistic datacenter. From a design perspective, to improve the robustness of the datacenter system, some designers may introduce buffering facilities to mitigate bursty admitted requests. Through redefining the Lyapunov function to incorporate the buffering facilities, we show that our framework can be enhanced to accommodate such a design choice. From an operational perspective, as most real-world datacenters are operated within a certain power budget to improve the performance (dollar) per watt, our results demonstrates how this model and framework can be extended to accomplish this by achieving a desired performance level with a pre-specified power budget.

Potential for further development of the research and the proposed course of action *(maximum half a page)*

The rise of research interests in power-performance trade-off is largely fueled by the performance requirements of cloud applications, as well as the operational requirements of the datacenter network. There are a number of avenues for further study. (1) There have been demands from emerging applications such as big data analytics and mobile cloud applications that are increasingly relying on geographically distributed services from multiple datacenters. However, wide-area inter-datacenter links are costly to build and upgrade, and inter-datacenter bandwidth is bound to become a scarce resource in the near future, inherently shared between all cloud applications. Therefore, there is a pressing need to manage and allocate inter-datacenter bandwidth at a fine granularity to meet diverse performance targets from such applications. (2) With the development of software defined networking (SDN), datacenters of large-scale Internet service providers such as Google are interconnected by the Software Defined Wide Area Network (WAN). In such a context, the scope of bandwidth allocations is broadened to cover the scenario where applications compete for the bandwidth of inter-datacenter links. How to effectively manage the resource utilization in this scenario becomes a challenging problem. (3) To reduce power cost and carbon footprint, an increasing number of cloud service providers have now considered green datacenters which utilize renewable energy sources such as solar or wind. However, unlike the stable supply of grid energy widely used today, it is challenging to utilize and realize renewable energy due to the uncertain, intermittent and variable nature. We are currently exploring these three directions in our on-going research

7. The Layman's Summary

(describe in layman's language the nature, significance and value of the research project, in no more than 200 words)

This project studies the problems associated with the optimization of energy efficiency and resource utilization in datacenters. The results of this research and its outcomes are: (1) in the short term, we have examined the fundamental characteristics of application workloads and their impact on energy and resource consumptions; (2) in the long term, we have designed and evaluated new energy-efficient algorithms towards optimal resource utilization. The long-term significance of this project lies in the establishment of a solid theoretical framework, which provides practical guidelines in the design of dynamic task allocation and virtual machine migration algorithms in achieving the optimal energy and resource utilization.

Part C: Research Output

8. Peer-reviewed journal publication(s) arising directly from this research project

(Please attach a copy of each publication and/or the letter of acceptance if not yet submitted in the previous progress report(s). All listed publications must acknowledge RGC's funding support by quoting the specific grant reference.)

The Latest Status of Publications				Author(s) *(**bold** the authors belonging to the project teams and denote the corresponding author with an asterisk*)*	Title and Journal/ Book *(with the volume, pages and other necessary publishing details specified)*	Submitted to RGC *(indicate the year ending of the relevant progress report)*	Attached to this report *(Yes or No)*	Acknowledged the support of this Joint Research Scheme *(Yes or No)*	Accessible from the institutional repository *(Yes or No)*
Year of publication	Year of Acceptance *(For paper accepted but not yet published)*	Under Review	Under Preparation *(optional)*						
2015				Y. Wu, C. Wu*, **B. Li**, L. Zhang, Z.-P. Li and F. Lau	Scaling Social Media Applications into Geo-Distributed Clouds, *IEEE/ACM Transactions on Networking,23(3):689-702, June 2015*	2015	yes	yes	yes
2014				F. Xu,**F.-M. Liu***, L. Liu, H. Jin, **B. Li** and B. Li	iAware: Making Live Migration of Virtual Machines Interference-Aware in the Cloud, *IEEE Transactions on Computers,63(12):3012-3025, December 2014*	2015	yes	yes	yes

2014				L. Chen, B. Li and **B. Li**	Allocating Bandwidth in Datacenter Networks: A Survey, Journal of Computer Science and Technology,29(5):910-917, September 2014.	2015	Yes	Yes	Yes
2014				Y. Zhu, Y.-C. Wu and **B. Li**	Trajectory Improves Data Delivery in Urban Vehicular Networks, *IEEE Transactions on Parallel and Distributed Systems*, 25(4):1089-1100, April 2014.	2015	Yes	Yes	Yes
2014				Z. Zhou, **F.-M. Liu***. H. Jin, **B. Li**, B. Li and **H. Jiang**	On Arbitrating the Power-Performance Tradeoff in SaaS Clouds,*IEEE Transactions on Parallel and Distributed Systems*, 25(10):2648-2658, October 2014.	2013	No	Yes	Yes

2014				**Y. Feng**, B. Li and **B. Li***	Price Competition in an Oligopoly Cloud Market with Multiple IaaS Cloud Providers, *IEEE Transactions on Computer, 63(1):59-73, January 2014*	2013	No	Yes	Yes
2013				**F.-M. Liu**, S. Shen, **B. Li***, B. Li, and H. Yin	Cinematic-Quality VoD in a P2P Storage Cloud: Design, Implementation and Measurements, *IEEE Journal on Selected Areas in Communications, Special Issue on Emerging Technologies in Communications, 31(9):214-226, September 2013*	2013	No	Yes	Yes

2013				W. Deng, **F.-M. Liu,** H. Jin, **B. Li*** and D. Li	Harnessing Renewable Energy in Cloud Datacenters: Opportunities and Challenges, *IEEE Networks Magazine*, 20(3):14-22, June, 2013	2013	No	Yes	Yes
2013				**F.-M. Liu, B. Li***, B. Li and H. Jin	Peer-Assisted On-Demand Streaming: Characterizing Demands and Optimizing Supplies, *IEEE Transactions on Computers*, 62(2):351-361, February 2013.	2013	No	Yes	Yes
2012				**G. Y. Keung, B. Li*** and Q. Zhang,	Intrusion Detection in Mobile Sensor Networks, *IEEE/ACM Transactions on Networking*, 20(4): 1152-1161, August 2012.	2013	No	Yes	Yes

9. **Recognized international conference(s) in which paper(s) related to this research project was/were delivered** *(Please attach a copy of each delivered paper. All listed papers must acknowledge RGC's funding support by quoting the specific grant reference.)*

Month/Year/ Place	Title	Conference Name	Submitted to RGC *(indicate the year ending of the relevant progress report)*	Attached to this report *(Yes or No)*	Acknowledged the support of this Joint Research Scheme *(Yes or No)*	Accessible from the institutional repository *(Yes or No)*
November/2014/Florida, USA	LiveRender: A Cloud Gaming System Based on Compressed Graphics Streaming	ACM International Conference on Multimedia	2015	Yes	Yes	Yes
July/2014/Spain	Fuel Cell Generation in Geo-Distributed Cloud Services: A Quantitative Study	IEEE ICDCS 2014	2015	Yes	Yes	Yes
April/2014/ Toronto, Canada	Towards Performance-Centric Fairness in Datacenter Networks	IEEE Infocom 2014	2015	Yes	No (my oversight for not including this grant ack.)	Yes
April/2013/Italy	A Framework for Truthful Online Auction in Cloud Computing with Heterogeneous User Demands	IEEE Infocom 2013	2013	No	Yes	Yes
April/2013/Italy	On Arbitrating the Power-Performance Tradeoff in SaaS Clouds	IEEE Infocom 2013	2013	No	Yes	Yes
November/2012/Japan	Jetway: Minimizing Costs on Inter-Datacenter Video Traffic	ACM Multimedia, 2012	2013	No	Yes	Yes
Oct-Nov/2012/USA	Airlift: Video Conferencing as a Cloud Service Using Inter-Datacenter Network	IEEE ICNP	2013	No	Yes	Yes
March/2012/ USA	Bargaining Towards Maximized Resource Utilization in Video Streaming Datacenters	IEEE Infocom, 2012	2013	No	Yes	Yes

March/2012/ USA	Scaling Social Media Applications into Geo-Distributed Clouds	IEEE Infocom, 2012	2013	No	Yes	Yes

10. Student(s) trained *(Please attach a copy of the title page of the thesis.)*

Name	Degree registered for	Date of registration	Date of thesis submission/ graduation
Chen Chen	PhD	Fall 2014	July 2018 expected
Shiyao Ma	PhD	Fall 2013	July 2017 expected
Yi Zhang	PhD	Fall 2012	July 2016 expected
Jingjie Jiang	MPhil	Fall 2012	July 2016 expected
Zhan Hu	MPhil	Spring 2011	Dec. 2015 expected
Xueling Lin	PhD	Fall 2013	August 2015
Wei Sun	MPhil	Fall 2011	May 2015
Yifei Zhu	MPHil	Fall 2012	January 2015
Shengkai Zhang	MPhil	Spring 2011	October 2014

11. Other impact *(e.g. award of patents or prizes, collaboration with other research institutions, technology transfer, etc.)*