Preview of Award 1456656 - Annual Project Report

Cover

Federal Agency and Organization Element to Which Report is Submitted:	4900
Federal Grant or Other Identifying Number Assigned by Agency:	1456656
Project Title:	CRII: NeTS: Exploiting System and Network Dynamics in Mobile Clouds
PD/PI Name:	Wei Gao, Principal Investigator
Recipient Organization:	University of Tennessee Knoxville
Project/Grant Period:	06/15/2015 - 05/31/2017
Reporting Period:	06/15/2015 - 05/31/2016
Submitting Official (if other than PD\PI):	Wei Gao Principal Investigator
Submission Date:	05/27/2016
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	Wei Gao

Accomplishments

* What are the major goals of the project?

Mobile cloud computing (MCC) bridges the gap between the limited capabilities of mobile devices and the increasing complexity of mobile applications, by offloading the computational workloads from local devices to the remote cloud. However, the effectiveness of mobile cloud computing could be impaired by the dynamic nature of system and network contexts, which lead to heterogeneous mobile application behaviors and seriously reduce the appropriateness of workload offloading decisions. This project exploits these critical dynamics in mobile clouds that are indispensable to efficient, prompt, and reliable workload offloading. More specifically, it addresses three closely intertwined research issues in mobile cloud computing. The first part investigates how to analytically formulate the stochastic characteristics of run-time application executions, based on which the workload offloading decisions are probabilistically made and systematic techniques are developed to practically enforce such decisions. The second part incorporates the contexts and performance requirements of mobile cloud applications into the design of wireless networks, so as to adaptively balance between the wireless energy cost and application performance in mobile clouds through fundamental redesign of wireless transmission scheduling algorithms. The third part focuses on testbed development to automatically investigate the run-time system and network dynamics of mobile cloud applications in practice. This testbed consists of off-the-shelf smartphones and wearable devices, and enables in-field experiments for evaluating the performance of the proposed techniques and system designs.

* What was accomplished under these goals (you must provide information for at least one of the 4 categories below)?

Major Activities: The present MCC systems have been considered unsatisfactory to ensure efficient workload offloading due to their ignorance of the various system and network dynamics. To efficiently address these dynamics and ensure the performance and energy efficiency of MCC in practical wireless network scenarios, we had the following major activities in the past year.

First, we synergistically took both mobile system and wireless network dynamics into account when designing MCC techniques, by adaptively investigating and balancing the tradeoff between the performance and energy efficiency of MCC applications with respect to the wireless network contexts. More specifically, we experimentally verified that workload offloading in MCC wastes a large amount of energy during the “tail times” after wireless data transmissions. To efficiently eliminate such unnecessary energy consumption without impairing the mobile application performance, we developed application-aware wireless transmission scheduling algorithms. These algorithms take both causality and run-time dynamics of application method executions into account when deferring wireless transmissions to eliminate tail times between these transmissions, so as to minimize the wireless energy cost while satisfying the specified application delay constraints. These algorithms can also be operated in practical wireless network scenarios with different information availability about mobile application executions.

Second, we developed physical-layer wireless networking techniques that are able to ensure reliable and timely wireless transmissions between mobile devices and the remote cloud for remote executions of mobile programs, regardless of the specific dynamics of the wireless network. Being different from traditional schemes which ensure network QoS via traffic scheduling or flow control but cannot scale to the amount of network traffic, our approach builds a separate wireless side channel that operates concurrently with the existing wireless channel over the same spectrum but dedicates for MCC. We built this side channel by exploiting the Signal-to-Noise Ratio (SNR) margin of a wireless channel, which is measured as the difference between the actual channel SNR and the required SNR for the configured channel operations. Our design is able to reach a throughput of up to 2.5 Mbps in the side channel without noticeably impairing the performance or reliability of the existing wireless channel. Therefore, it could be applied to support a large collection of MCC applications, ranging from mobile gaming and image processing to future emerging applications of Virtual Reality.

Education Activities:

Two PhD students have worked on the project. Some of the research results have been integrated with the education curricula at University of Tennessee, Knoxville. For example, we have added smartphone techniques to the contents of our undergraduate course “ECE455: Embedded System Design” and provided the students with the opportunity of working with the Android OS and off-the-shelf mobile devices.

Specific Objectives: System and network dynamics have posed significant challenges in adapting mobile clouds to the uncertain changes of system conditions and environmental contexts. First, the run-time dynamics of mobile application executions could affect the correctness of the decisions of remote method executions. To efficiently incorporate such dynamics into MCC, an analytical framework is needed to formulate the stochastic characteristics of application execution paths at run-time and apply these characteristics into MCC operations. Existing techniques of remote code execution, however, have not reached the required level of adaptability or granularity for handling such run-time dynamics. Second, the efficiency and reliability of workload offloading depend on high-speed, energy-efficient, and stable wireless networks. Current designs of wireless networks in mobile clouds, however, are isolated from the workload offloading decisions and ignore the specific contexts or performance requirements of mobile applications. To address the above challenges, our work incorporates the following two objectives.

Our first objective is to simultaneously ensure both the mobile application performance and the energy efficiency during MCC operations. When mobile workloads are offloaded to the remote cloud via cellular networks, the frequent wireless data transmissions will incur a large amount of tail times at the cellular radio interface and hence lead to unnecessary energy waste. A common solution to reducing such tail energy is to defer wireless transmissions and send data as bundles, but may increase the response delay and impair the performance of mobile applications. Instead, we aim to fundamentally redesign the wireless network transmission strategies in mobile clouds with respect to the mobile application contexts, so as to efficiently balance between the energy efficiency and application performance in mobile clouds through appropriate scheduling of wireless data transmissions incurred by remote method executions.

Our second objective is to develop new wireless networking techniques for timely wireless data transmissions in MCC. In practice, such real-time wireless traffic of MCC applications may be seriously delayed when competing with other data traffic being transmitted concurrently over the same wireless channel. A straightforward solution to such network congestion is to allocate additional wireless spectrum that is exclusively used for real-time wireless traffic. For example, a dedicated spectrum is designated as the control plane in cellular networks. However, such exploitation of additional spectrum is infeasible due to the severe scarcity of wireless spectrum resources nowadays. Instead, another viable solution to removing this fundamental limitation on supporting real-time MCC traffic is to explore a wireless communication side channel, which operates over the same spectrum but dedicates for MCC. When the main channel is congested, MCC traffic is transmitted through the side channel. Hence, such MCC traffic will never be delayed by concurrent wireless traffic, and its latency only depends on the link propagation delay. The major challenge, though, is how to design such a side channel with sufficiently high throughput without impairing the functionality and performance of the main wireless channel.

Significant Results: We have developed multiple application-aware wireless transmissions scheduling algorithms. More specifically, our work minimizes the wireless energy cost of workload offloading while satisfying the application delay constraints, which are either specified prior to execution or flexibly adjusted at run-time. First, we developed efficient algorithms for offline transmission scheduling in MCC, based on a pre-known sequence of transmissions to be scheduled. These scheduling algorithms ensure global minimization of the tail times after wireless data transmissions through a Dynamic-Programming-based approach. The basic idea is to break the scheduling problem down to a collection of simpler subproblems, which are individually solved and then merged together for global results. Our algorithms also enable flexible balancing between the energy and delay aspects in MCC at run-time. Furthermore, we incorporate the run-time dynamics of mobile application executions into transmission scheduling by developing a stochastic framework, so as to extend the transmission scheduling algorithms to online operations by probabilistically predicting the application execution path in the future. We have evaluated the effectiveness of our proposed transmission scheduling approaches over a large collection of realistic mobile applications, by comparing our approaches with multiple existing MCC schemes that ignore the run-time characteristics of mobile application executions. The evaluation results have shown that our scheme could significantly reduce the overhead of workload offloading in MCC by more than 40% and improve the mobile application performance by more than 25%.

We have developed a practical high-throughput design of the wireless side channel that efficiently supports real-time wireless traffic for MCC, by exploiting the unique properties of modern digital modulation methods, particularly OFDM which will be the technical foundation of next-generation high-speed wireless networks (e.g., LTE-A and 5G). The key insight of such a side channel is that the SNR of a wireless channel is usually higher than the SNR required to support the data rate being used, due to inaccurate SNR estimation and conservative rate adaptation in wireless networks. This in-band SNR margin can be exploited to encode data as patterned interference over the main channel. The impact of such a side channel on packet decoding over the main channel, on the other hand, could be efficiently eliminated by limiting the amount of additional patterned interference within the scope of the main channel's SNR margin. As a result, our basic idea of the side channel design is to encode data as patterned interference by erasing the energy of specific subcarriers in the main channel's OFDM symbols. Since such energy erasure does not increase the RF transmit power, it can be used to encode data into every OFDM symbol in the main channel, hence dramatically increasing the side channel throughput. On the other hand, since OFDM modulates data into separate subcarriers in both time and frequency domains, the amount of patterned interference could be efficiently controlled by interfering only a small portion of OFDM subcarriers, without affecting the packet decoding in the main channel and its resistance to channel contention. We have implemented the proposed system design over practical Software-Defined Radio (SDR) hardware platforms, and evaluated the effectiveness of our design over realistic VoIP applications. The experimental results show that our system can provide a side channel throughput of up to 2.5 Mbps, which is 10 times higher than existing work with minimal impairment to the main channel performance. It also reduces the data transmission delay in commodity 802.11 networks by up to 90%, and significantly eliminates the chance of delay jitters in such networks.

Key outcomes or Other achievements:	The results of our work “Application-Aware Traffic Scheduling for Workload Offloading in Mobile Clouds,” has been accepted by the highly competitive *IEEE Conference on Computer Communications (INFOCOM 2016)*, which has an acceptance ratio of 18%. The results of our work “Supporting Real-Time Wireless Traffic through A High-Throughput Side Channel,” has been accepted by the highly competitive *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2016)*, which has an acceptance ratio of 18%.

*** What opportunities for training and professional development has the project provided?**

Two PhD students have worked on the project, and the research results have been published at various academic conferences.

*** How have the results been disseminated to communities of interest?**

Our research work in this project has resulted in two conference papers. These publications will help people better understand our novel techniques on exploiting system and network dynamics in mobile clouds, and further apply these techniques to improve the performance and energy efficiency of MCC applications in practice. We have also given seminar and summer camp talks to high school students to stimulate their interest in engineering majors.

*** What do you plan to do during the next reporting period to accomplish the goals?**

We will further investigate techniques to adapt the MCC decisions and operations to the fluctuating conditions of the wireless network channel and traffic. In particular, when the wireless link quality degrades, the amount of data being transmitted via wireless links will be limited to save energy. In this case, we will adaptively preserve the mobile application performance by only transmitting the most important program states and maximizing the amount of mobile programs being executed at the remote cloud. Furthermore, we will also explore the possibility of exploiting the heterogeneity of such wireless link quality over multiple co-located mobile users, which could potentially collaborate with each other to improve the energy efficiency of their MCC applications.

Products

Books

Book Chapters

Inventions

Journals or Juried Conference Papers

Licenses

Other Conference Presentations / Papers

Liang Tong and Wei Gao (2016). *Application-Aware Traffic Scheduling for Workload Offloading in Mobile Clouds*. in the Proceedings of the 35th IEEE Conference on Computer Communications (INFOCOM). San Francisco, CA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Haoyang Lu and Wei Gao (2016). *Supporting Real-Time Wireless Traffic through A High-Throughput Side Channel*. in Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc). Paderborn, Germany. Status = ACCEPTED; Acknowledgement of Federal Support = Yes

Other Products

Other Publications

Patents

Technologies or Techniques

Thesis/Dissertations

Websites

Project website
http://web.eecs.utk.edu/~weigao/reporting/mcc_dynamics.html

On this project website, we provide details regarding this specific project (personnel, papers, software, etc.).

Participants/Organizations

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
Gao, Wei	PD/PI	1
Lu, Haoyang	Graduate Student (research assistant)	3
Tong, Liang	Graduate Student (research assistant)	4

Full details of individuals who have worked on the project:

Wei Gao
Email: weigao@utk.edu
Most Senior Project Role: PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: Manage the project and the research team. Design and evaluate the application-aware transmission scheduling algorithms for the mobile cloud. Design the high-throughput wireless side channel for real-time MCC traffic.

Funding Support: This grant

International Collaboration: No
International Travel: No

Haoyang Lu
Email: hlu9@vols.utk.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 3

Contribution to the Project: Design, implement and evaluate the high-throughput wireless side channel for real-time MCC traffic.

Funding Support: This grant

International Collaboration: No
International Travel: No

Liang Tong
Email: ltong@vols.utk.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 4

Contribution to the Project: Designed and evaluated the application-aware transmission scheduling algorithms for MCC.

Funding Support: This grant

International Collaboration: No
International Travel: No

What other organizations have been involved as partners?
Nothing to report.

What other collaborators or contacts have been involved?
Nothing to report

Impacts

What is the impact on the development of the principal discipline(s) of the project?

Contextual awareness enabled by recent emergence of cognitive mobile applications and wearable devices fundamentally transforms the way people behave and interact with the environment, but also imposes serious challenges on the capabilities and battery lifetime of mobile devices. This project aims to completely rethink how mobile cloud computing could be practically realized to alleviate the local computational burden and adaptively support contextual awareness over heterogeneous mobile scenarios, by turning analytical formulations of the various system and network dynamics into actionable system design strategies. The results from wireless network scheduling and channel design are likely to foster further research along these directions. The research can also spawn a new area of research on application-aware wireless networking for mobile cloud computing. Finally, the analysis techniques, the evaluation methodology and systems developed in this research will be valuable for future undertakings.

What is the impact on other disciplines?

The mobile cloud is a typical example of mobile computing systems with complex dynamics rooted in the system's execution. Being able to precisely characterize and appropriately exploit these dynamics to improve the system efficiency and adaptability has a direct and immediate impact on a large variety of ubiquitous computing and cyber-physical systems.

What is the impact on the development of human resources?

Many of the research results have been integrated into the undergraduate curricula at the University of Tennessee, by adopting many perspectives of the research results for undergraduate students' course projects and senior design topics. The project has supported two PhD students working on their dissertations. The involvement of the graduate and undergraduate students into this research will prepare them for leadership roles in computer science research, academia, and industry.

What is the impact on physical resources that form infrastructure?
Nothing to report.

What is the impact on institutional resources that form infrastructure?
Nothing to report.

What is the impact on information resources that form infrastructure?
Nothing to report.

What is the impact on technology transfer?
Nothing to report.

What is the impact on society beyond science and technology?
Nothing to report.

Changes/Problems

Changes in approach and reason for change
Nothing to report.

Actual or Anticipated problems or delays and actions or plans to resolve them
Nothing to report.

Changes that have a significant impact on expenditures
Nothing to report.

Significant changes in use or care of human subjects
Nothing to report.

Significant changes in use or care of vertebrate animals
Nothing to report.

Significant changes in use or care of biohazards
Nothing to report.