Preview of Award 1526769 - Annual Project Report

Cover |
Accomplishments |
Products |
Participants/Organizations |
Impacts |
Changes/Problems

Cover

Federal Agency and Organization Element to Which Report is Submitted:	4900
Federal Grant or Other Identifying Number Assigned by Agency:	1526769
Project Title:	NeTS: Small: Collaborative Research: Network-Centric Mobile Cloud Computing
PD/PI Name:	Wei Gao, Principal Investigator
Recipient Organization:	University of Tennessee Knoxville
Project/Grant Period:	10/01/2015 - 09/30/2018
Reporting Period:	10/01/2015 - 09/30/2016
Submitting Official (if other than PD\PI):	Wei Gao Principal Investigator
Submission Date:	09/28/2016
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) has been used to address the resource limitation of mobile devices by migrating expensive local computations to the cloud. However, transmitting data wirelessly from mobile devices to the cloud also consumes energy. Hence, the key problem of MCC is how to minimize the energy consumption while preserving the mobile application performance. Different from traditional solutions which focus on reducing the cost of wireless transmission solely from the application perspective, this project focuses on designing MCC schemes from a network-centric perspective, by investigating, formulating, and mitigating the impact of special characteristics of wireless networks on the energy efficiency of MCC. The proposed research could benefit end users with various mobile devices by extending their battery lifetime and improving their performance.

This project aims to improve the performance of MCC by mitigating the impacts of two special characteristics of wireless networks: the long-tail problem at the wireless interface and the quality variations of the wireless link. More specifically, this project consists of three closely intertwined research thrusts: (i) reducing the amount of tail energy when transmitting the program states to the remote cloud, while ensuring that the performance requirements of mobile applications can be met; (ii) mitigating the impact of wireless link quality on both energy and performance, and minimizing the degradation of application performance when the wireless link quality is low; and (iii) exploiting the difference of wireless link quality among mobile users to further improve the energy efficiency of MCC via user cooperation. An experimental testbed will be developed to investigate the practical impact of wireless network characteristics on MCC and evaluate the proposed MCC schemes.

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

categories below)?

Major Activities:

Current MCC solutions have limitations in that they reduce the cost of wireless transmission solely from the application perspective, by reducing the size of program states being transmitted according to the network bandwidth and programs' computational complexity. They have minimal explorations into the special characteristics of wireless networks, which affect the energy consumption of wireless transmission and have a pivotal role in energy-efficient MCC. Ignorance of these wireless network characteristics is also the major factor hindering practical integration of mobile devices into the cloud. To efficiently investigate, formulate and mitigate the impact of special characteristics of wireless networks on the energy efficiency of remote program execution in MCC, we had the following major activities in the past year.

First, we took the 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, to efficiently eliminate the unnecessary energy consumption of workload offloading in MCC during the "tail times" after wireless data transmissions, while preventing any degradation on 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, appropriate scheduling over multiple concurrent communication links is crucial to the efficiency of MCC operations in wireless networks. Existing wireless network scheduling algorithms usually formulate the scheduling problem as solving the NP-hard Maximum Weighted Independent Set (MWIS) problem over a network conflict graph, but are designed to provide approximations to global optimality via distributed operations in wireless networks. Therefore, these algorithms will need to be frequently operated over the entire network in cases of network dynamics, regardless of the actual network area being affected by these dynamics. Instead, in order to reduce the computation and communication overhead when scheduling wireless networks in such dynamic environmental settings, we developed distributed algorithms that adaptively constrain network scheduling within the limited scope where network dynamics occur. The scheduling results from such limited operations are then combined with the previous scheduling results over the remaining portions of the network, hence still providing guaranteed network throughput. The performance of our proposed algorithms has been validated by formal analysis, and been verified by both numerical studies and real-world experiments.

Third, 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: The goal of this research is to design MCC schemes from a *network-centric* perspective, by investigating, formulating, and mitigating the impact of special characteristics of wireless networks on the energy efficiency of remote program execution in MCC. We consider these characteristics in two aspects: i) *long-tail problem at the wireless interface*, i.e., the cellular radio interface stays in the high-power state for some time after each data transmission, and ii) *quality of the wireless link*, which may be highly dynamic due to various environmental factors. We will incorporate both aspects to ensure that each program state is transmitted to the cloud in the most energy-efficient manner. However, such network-centric design is challenging due to the following reasons. First, due to the low bandwidth and long delay of cellular network links, most existing works focus on transmitting program states via WiFi instead of cellular networks, and remote execution of mobile programs would be impossible when WiFi is unavailable. Recent deployment of 4G networks such as LTE eliminates the bandwidth and delay barrier of exploiting cellular networks in MCC and hence enables uninterrupted access to the remote cloud, but incurs a large amount of additional energy cost due to the long-tail problem. Second, the condition of wireless link quality could be highly dynamic in practice, and should also be considered to ensure energy-efficient MCC. Most existing MCC techniques, however, inappropriately assume that the wireless link quality is stable over time. To address the above challenges, our work incorporates the following two objectives.

Our first objective is to reduce the amount of energy being consumed during the tail times of wireless data transmissions for MCC operations. In MCC, the tail energy should be reduced with respect to the performance requirements of mobile applications. Existing work reduces the tail energy by deferring transmissions of program states and sending them as bundles, but greatly increases the response time of mobile applications. To address this problem, we need to decide whether a mobile program should be executed locally or remotely, to minimize the energy consumption without deferring the program execution. Such decision will be made according to the amount of tail energy incurred by the transmission of program states. Afterwards, based on these decisions of remote program execution, we further reduce the tail energy by deferring the execution of a mobile application within its delay constraint and scheduling the transmissions of program states accordingly.

Our second objective is to mitigate the impact of wireless network dynamics, which affect the quality of wireless links, on both the energy efficiency and performance of MCC. 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 shall 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. Being different from existing work which determines the importance of a program state solely from the program’s computational complexity, we need to develop fine-grained program parsing techniques to evaluate the importance of each memory field in a program state to executing other programs. Application partitioning strategy will then be developed based on the importance of program states.

Significant Results: First, 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%.

Second, we have made theoretical efforts to reduce the computation and communication overhead of wireless network scheduling, without degrading the wireless communication performance over multiple concurrent wireless links. By experimentally investigating wireless network dynamics in different types of practical application scenarios, we observe that these dynamics in wireless networks usually affect only a small portion of the network at any specific moment, when the rest major portion of the network remains stable and unaffected. Based on this observation, we developed distributed algorithms that schedule wireless networks only within the limited scope where network dynamics occur. Our proposed algorithms combine the scheduling results from such limited operations with the previous scheduling results over the remaining portions of the network, hence still providing guaranteed network throughput with much lower overhead. We conducted formal analysis that demonstrates the correctness of our scheduling algorithms on avoiding network interference, and also shows that our algorithms provide a fixed approximation ratio to optimality. We evaluated the performance and overhead of our scheduling algorithms using both numerical studies with large-scale conflict graphs and real-world experiments over software-defined radio platforms that take practical wireless channel conditions into account.

In addition, 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 “Scheduling Dynamic Wireless Networks with Limited Operations,” has been accepted by the highly competitive *IEEE International Conference on Network Protocols (ICNP)*, which has an acceptance ratio of 20%.

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%.

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%.

*** 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 three 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, the energy efficiency of MCC could be further improved by cooperation among mobile users. Existing schemes focus on cooperative workload balancing among collocated mobile devices, but cannot reduce the total energy cost over all mobile devices. To address this problem, we will reduce the tail energy by aggregating the transmissions of program states from multiple co-located mobile devices to one device, called the proxy. Such a proxy will be selected based on the quality of wireless links at different mobile devices, to ensure that the program states are only transmitted via wireless links with good quality.

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 Proceedings of the 35th IEEE Conference on Computer Communications (INFOCOM). San Francisco, USA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Haoyang Lu and Wei Gao (2016). *Scheduling Dynamic Wireless Networks with Limited Operations*. in Proceedings of the 24th IEEE International Conference on Network Protocols (ICNP). Singapore. 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 = PUBLISHED; 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_network_centric.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
Li, Yong	Graduate Student (research assistant)	4
Lu, Haoyang	Graduate Student (research assistant)	3

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. Design the wireless network scheduling protocols with limited operations.

Funding Support: This grant.

International Collaboration: No
International Travel: No

Yong Li
Email: yli118@vols.utk.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 4

Contribution to the Project: Contribute to the design and implementation of various MCC schemes and protocols.

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. Design and implementation of the wireless network scheduling protocols with limited operations.

Funding Support: this grant.

International Collaboration: No
International Travel: Yes, Germany - 0 years, 0 months, 7 days

---

What other organizations have been involved as partners?

Name	Type of Partner Organization	Location
Guohong Cao	Academic Institution	University Park, PA

Full details of organizations that have been involved as partners:

Guohong Cao

Organization Type: Academic Institution
Organization Location: University Park, PA

Partner's Contribution to the Project:
Collaborative Research

More Detail on Partner and Contribution: Dr. Guohong Cao from Pennsylvania State University has been actively involved in discussing the design and implementation of our application-aware traffic scheduling work.

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?

MCC can greatly benefit a large variety of mobile devices ranging from smartphones to embedded sensor devices, and from smart watches to wearable medical devices. The transformative nature of the proposed research is to completely redesign MCC from a network-centric perspective, and to turn various system modeling and formal analysis techniques into actionable distributed algorithms and design strategies that apply to practical mobile cloud systems. It will investigate the essential factors determining the energy efficiency of MCC that are tightly coupled with the special characteristics of wireless networks, and will inspire further theoretical and systematic studies that can open up new areas of research. To summarize, the research output will have the following impacts. First, it leads to a new design philosophy for MCC that is pivotal to practical integration of mobile devices into the cloud. Second, it introduces novel solutions to deal with the long-tail problem related to the wireless radio interface that is vital to energy-efficient MCC. Third, this work can spawn a new area of research on mitigating the impact of wireless link quality on MCC. Fourth, the techniques of cooperative wireless link selection in MCC are likely to foster further research along this direction. Finally, the developed testbed as a unique research facility will benefit the whole research community.

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.