My Desktop
Prepare & Submit Proposals
Proposal Status
Proposal Functions
Awards & Reporting
Notifications & Requests
Project Reports
Submit Images/Videos
Award Functions
Manage Financials
Program Income Reporting
Grantee Cash Management Section Contacts
Administration
Lookup NSF ID

Preview of Award 1456656 - 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:	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 - 05/31/2018
Reporting Period:	06/01/2016 - 05/31/2017
Submitting Official (if other than PD\PI):	Wei Gao Principal Investigator
Submission Date:	05/13/2017
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 executions of mobile applications 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 target on the phenomenon that a mobile user nowadays is usually equipped with multiple types of mobile computing devices, and aim to fully utilize the capabilities of these mobile devices to maximize the performance of mobile applications. More specifically, our basic idea is to construct a *personal mobile cloud*, which incorporates and interconnects all the mobile devices owned by a user via wireless links. These devices are then able to flexibly share system resources with each other, augmenting the mobile computing capability provided to the user.

Second, we aim to reduce the computation and communication overhead when scheduling highly dynamic wireless networks with frequent changes of link status, by adaptively restricting the scope of network scheduling within the local proximity of where wireless link status changes. More specifically, existing scheduling algorithms are designed to provide approximations to global optimality of solving the Maximum Weighted Independent Set (MWIS) problem over a network conflict graph via distributed operations in wireless networks, but will frequently reschedule the entire network in cases of network dynamics regardless of the actual network area being affected by these dynamics. Instead, 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.

Third, we strive the address one key limitation of the current mobile cloud computing (MCC) schemes: these traditional schemes either restrict the scope of workload offloading to a specific set of system frameworks and mobile applications, or migrate a large amount of application contexts to the remote cloud regardless of the specific execution patterns of the application partition to be offloaded. These limitations seriously impair the efficiency of workload offloading in practical mobile cloud scenarios. Instead, we present a novel design of workload offloading system which addresses these challenges and performs automated method-level workload offloading with least context migration.

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 a generic system framework that allows heterogeneous types of mobile systems to efficiently interconnect with each other and form a personal mobile cloud. Being different from traditional solutions which are limited to interconnecting mobile devices with respect to an individual mobile application or a specific type of shared hardware, we aim to realize such a personal mobile cloud in a generic manner across heterogeneous mobile devices, without incurring repetitive reprogramming efforts to support any newly added types of mobile hardware. To efficiently support such generic interconnection and overcome the challenge of hardware and software heterogeneity in mobile systems, our basic idea is to develop the resource sharing framework as a middleware in the mobile OS, which exploits the existing mobile OS services to share resources between mobile devices. These services hide the low-layer details of device driver operations while providing unified data access APIs to user applications. Interconnection between mobile devices, then, could be realized via remote access and invocation of these OS services. Any new device can be incorporated into the mobile cloud by inserting our framework into its OS, without modifying the OS kernel, our framework itself, or the source code of any mobile application.

To efficiently reduce the overhead of network scheduling over highly dynamic wireless networks, we experimentally investigated the wireless network dynamics in different types of practical application scenarios, and 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.

To improve the efficiency of workload offloading and address the limitation of existing MCC schemes, we focus on developing systematic solutions that minimize the amount of memory contexts being migrated to the remote cloud. Our basic idea of achieving the least context migration while ensuring the offloading appropriateness is to identify the memory contexts that may be accessed by a specific application method prior to its execution, through offline parsing of the application executables. The parsing results will be stored as metadata along with the application executables at local mobile devices, and will be utilized by the run-time application execution to screen the thread stack and heap contexts to migrate only the relevant memory contexts to the remote cloud. In order to further improve the efficiency of such offline parsing and avoid unnecessary redundancy during parsing, we also pre-parse all the OS libraries that may be invoked by mobile application methods and then reuse these parsing results for different user applications.

Key outcomes or Other achievements:

The results of our work “Interconnecting Heterogeneous Devices in the Personal Mobile Cloud,” has been accepted by the highly competitive *IEEE Conference on Computer Communications (INFOCOM 2017)*, which has an acceptance ratio of 20.9%.

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 2016)*, which has an acceptance ratio of 20.1%.

The results of our work “Minimizing Context Migration in Mobile Code Offload” has been published by the highly prestigious *IEEE Transactions on Mobile Computing* in 2017.

*** 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 journals and conference proceedings.

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

Our research work in this project has resulted in one journal paper and 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

Haoyang Lu and Wei Gao (2016). Scheduling Dynamic Wireless Networks with Limited Operations. *Proceedings of the 24th IEEE International Conference on Network Protocols (ICNP)*. . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Yong Li and Wei Gao (2017). Interconnecting Heterogeneous Devices in the Personal Mobile Cloud. *Proceedings of the 36th IEEE Conference on Computer Communications (INFOCOM)*. . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Yong Li and Wei Gao (2017). Minimizing Context Migration in Mobile Code Offload. *IEEE Transactions on Mobile Computing*. 16 (4), 1005-1018. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Licenses

Other Conference Presentations / Papers

Other Products

Other Publications

Patents

Technologies or Techniques

Thesis/Dissertations

Websites

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)	3
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. Designed the wireless network scheduling protocols with limited operations. Design the workload offloading techniques with least context migration.

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: 3

Contribution to the Project: Design, implement and evaluate the mobile workload offloading techniques with least context migration.

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 wireless network scheduling protocol with limited operations.

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 mobile computing system design and wireless network scheduling protocols are likely to foster further research along these directions. The research can also spawn a new area of research on efficiently interconnecting heterogeneous mobile devices towards a personal mobile cloud. 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

The PI has requested a no-cost extension of this grant for one additional year due to the difficulty of recruiting graduate students in the first year of the project. This request has been approved by NSF.

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.