1st DFG/GENI Doctoral Consortium, San Juan, PR March 13th-15th, 2011

Secure Opportunistic Mobile Application Offload for Enterprise Networks

Aaron Gember and Aditya Akella
University of Wisconsin – Madison

Abstract

Application-independent offloading is a promising approach to allow resource-limited mobile devices to access intensive applications more frequently without performance or energy costs. Offloading migrates part of a running mobile application and executes it on an available desktop or server. Existing approaches have not been widely adopted because: 1) No mechanisms are in place to maintain the *security* and *privacy* of data used by an offload application; and 2) The focus is on *what* to offload instead of *where* to offload, and prior systems assume dedicated local resources are available for offloading. SOON is our mobile application offloading system designed to *opportunistically* leverage available resources and offer *security guarantees* and *latency/energy improvements* to smartphones. Expanding SOON beyond an enterprise network to utilize cloud resources requires careful decisions to ensure privacy and performance are preserved.

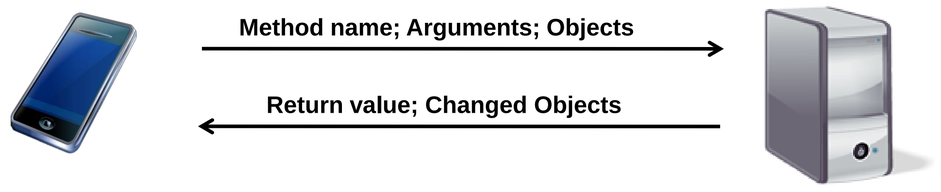

Research Objectives

- Preserve data privacy by appropriately securing execution state transfers and selecting offloading resources with sufficient trust
- Provide latency improvements, energy savings, or both to smartphone users by opportunistically leveraging compute resources with sufficient capacity and network overhead
- Identify network services required in the enterprise, cloud, or WAN to ensure security is preserved and benefits exceed overheard

Experiments

SOON has been evaluated in a small enterprise setting using the University of Wisconsin – Madison OpenFlow network. We run two different applications (chess and speech recognition) on12 Google Android phones with varying privacy levels and objects, and we use 4-6 desktops for offloaded execution. SOON provides significant latency improvements to all phones, but applications desiring energy savings run 50 – 60% slower than latency seeking applications. Phone energy usage also decreases by 24 – 44%, with more savings for phones seeking energy savings. Some phones (P10) are unable to offload due to trust requirements and resulting contention with only 4 desktops. Increasing the available resources to 6 desktops alleviates the issue and provides further latency and energy benefits to all phones.

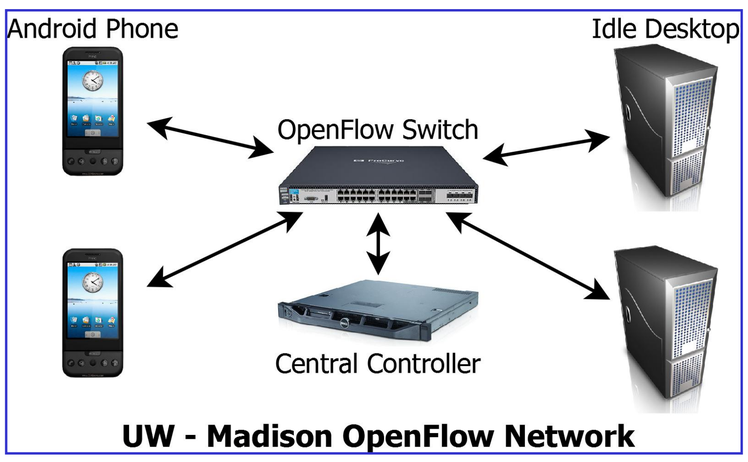

SOON setup for a small enterprise

Phone	Application	Goal	Trusted Resources
P1,P2,P3	Chess	Latency	All
P4,P5,P6	Chess	Energy	All
P7	Speech Recognition	Latency	D1,D2,D6
P8	Speech Recognition	Latency	D2,D3,D6
P9	Speech Recognition	Latency	D1,D3,D6
P10	Speech Recognition	Energy	D1,D2,D6
P11	Speech Recognition	Energy	D2,D3,D6
P12	Speech Recognition	Energy	D1,D3,D6

Configuration for phones

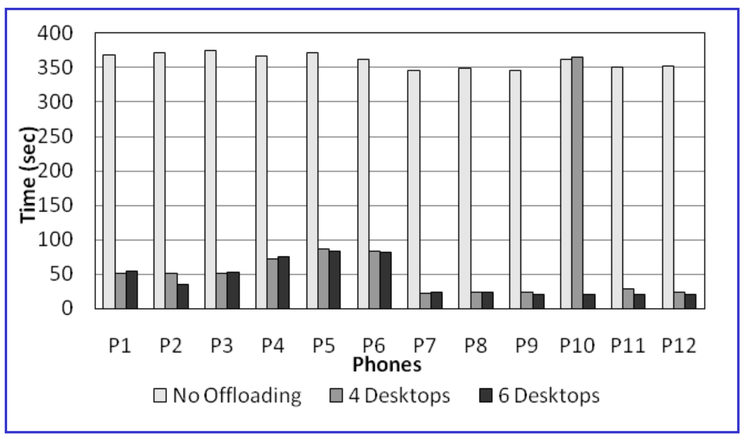

Comparison of execution times

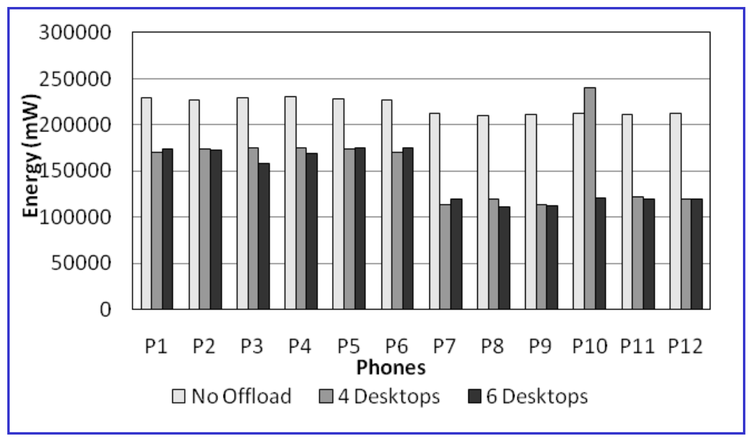

Comparison of energy usage

Future Work

- Measure the latency improvements and energy savings using compute resources in a remote "cloud"
- Analyze the feasibility of SOON's central controller and offloading decision process when smartphone users are external to the enterprise
- Identify enterprise services that could utilize the same model as SOON but offload execution from data center servers to the cloud

Use of Glab/GENI Infrastructure

The completed experiments use the GENI infrastructure at the University of Wisconsin – Madison, namely our OpenFlow network. Future experiments will utilize EmuLab and OpenFlow networks at other campuses, to represent remote enterprise sites and commercial clouds. Furthermore, PlanetLab sites will be used to explore SOONs performance and feasibility when smartphone users are external to the enterprise.

Current and Proposed Publications

Conference / Journal Papers

A. Gember, C. Dragga, A. Akella. *Secure Opportunistic Mobile Application Offloading for Enterprise Networks*. Submitted to MobiSys 2011.