

Whitepaper: The New Clouds on the Block

Karim Habak, Ahmed Saeed, Mostafa Ammar, Ellen Zegura
School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332, USA
{karim.habak, ahmed.saeed, ammar, ewz}@cc.gatech.edu

Context and New Research Directions

We believe that the availability of capable devices at the edge is providing a unique but mostly overlooked opportunity to overcome the traditional cloud computing limitations without the cost of deploying local clouds (cloudlets). We back up this assertion with the four observations. First, while the gap remains between truly mobile devices (handhelds, wearable) and high capacity servers, mobile devices have grown increasingly powerful, especially when laptops are included. Second, the devices at the edge are mostly idle or underutilized and can be clustered together to provide significant compute resource. Third, in many scenarios the aggregate capacities of these idle devices exceed the capabilities of the deployed cloudlet servers. Finally, efficiently utilizing the available devices at the edge as compute resources to assist neighboring devices have a natural scaling property not enjoyed by a cloudlet, which must be provisioned ahead of time.

We include in our vision of “new clouds on the block” those that comprise handheld and laptop user devices as well as the increasing number of capable Internet of Things (IoT) devices that are found in homes and industrial settings and will see future application in Smart City civic initiatives. We see these clouds as present in homes; public spaces such as public transportation, libraries and parks; and commercial spaces such as coffee shops.

To truly leverage the devices at the edge and introduce them as a significant part in the computing ecosystem, We believe further questions need to be addressed including (1) How to provide a stable interface with these clouds, overcome their dynamic nature, and provide their users with performance guarantees? (2) How to efficiently discover and utilize the available resources in devices without affecting their main user experiences or disrupting their main functionalities? (3) How to provide users with proper incentive to share their resources? (4) How mobile device clouds fit in the cloud computing ecosystem and how can they interact with one another and with data center clouds?

Testbed Infrastructure Requirements

Exploration with edge clouds requires infrastructure that comes much closer to real users than research that focuses on core infrastructure. This is both interesting

and significantly challenging. In our own work, we typically use simulation to explore the parameter space, and augment that with small scale testbeds that provide a proof of concept and a sanity check for simulation results. The limits of both approaches are well known. Locally we are beginning to explore using a larger testbed of IoT devices in the Aware Home Research Laboratory, a wired (and wireless) home built in the late 1990s by Georgia Tech for research purposes. No one lives there, however, so the connection to real users is missing, though study participants can be brought in for controlled experiments. It may be time for the networking and user interaction communities to come together to create a living lab of people who agree to large scale data collection about their network and application use, to create traces and testing opportunities for many researchers. This sort of “in the wild” facility could enable significant advances in system development and understanding of the interaction between people and technologies beyond the Internet.

Our Background

Our team has been working in the area of mobile and edge computing for nearly 10 years. Our interest in edge computing, especially involving user devices, grew out of our earlier work in disruption tolerant networking. We have subsequently developed and evaluated systems that realize edge computing entirely from opportunistic interaction between devices (Serendipity), that use an intermediary service to manage mismatches between the needs of intermittently connected devices and stable cloud infrastructure (COSMOS), that provide cloudlet-like infrastructure using a decoupled controller and device pool (FemtoClouds) and that achieve hyper-local, data center like functionality in a network of IoT devices (SymbIoTs). In the interest of space, we do not include references to this work, but it is easy to find on-line using the system names.

Travel Funding and Reimbursement

We can cover travel costs for the faculty members on this white paper. We would appreciate either full or partial reimbursement for our students to attend.